## Visual Numerics

IMSL
Fortran Numerical Library

## User's Guide math/Library

## Visual Numerics

Visual Numerics Corporate Headquarters
2500 Wilcrest Drive
Suite 200
Houston, TX 77042
USA Contact Information
Toll Free:
800.222.4675

Houston, TX: 713.784.3131
Westminster, CO: $\quad 303.379 .3040$
Email:
info@vni.com
Web site: www.vni.com

Visual Numerics has Offices Worldwide
USA • UK • France • Germany •Mexico
Japan •Korea • Taiwan
For contact information, please visit www.vni.com/contact
© 1970-2006 Visual Numerics, Inc. All rights reserved.
Visual Numerics and PV-WAVE are registered trademarks of Visual Numerics, Inc. in the U.S. and other countries. IMSL, JMSL, JWAVE, TS-WAVE and Knowledge in Motion are trademarks of Visual Numerics, Inc. All other company, product or brand names are the property of their respective owners.

IMPORTANT NOTICE: Information contained in this documentation is subject to change without notice. Use of this document is subject to the terms and conditions of a Visual Numerics Software License Agreement, including, without limitation, the Limited Warranty and Limitation of Liability. If you do not accept the terms of the license agreement, you may not use this documentation and should promptly return the product for a full refund. This documentation may not be copied or distributed in any
form without the express written consent of Visual Numerics.

MST $\quad \begin{aligned} & \text { C, C\#, Java } \\ & \text { AM , and Fortran } \\ & \text { Application Development T }\end{aligned}$
Application Development Tools

## Table of Contents

Introduction ..... xix
The IMSL Fortran Numerical Library ..... xix
User Background ..... xix
Getting Started ..... xx
Finding the Right Routine ..... xx
Organization of the Documentation ..... xxi
Naming Conventions ..... xxii
Using Library Subprograms ..... xxiii
Programming Conventions ..... xxiv
Module Usage ..... xxiv
Using MPI Routines ..... xxiv
Programming Tips ..... xxvi
Optional Subprogram Arguments ..... xxvi
Optional Data ..... xxvii
Overloaded $=, /=$, etc., for Derived Types ..... xxviii
Error Handling ..... xxix
Printing Results ..... xxix
Fortran 90 Constructs ..... xxx
Shared-Memory Multiprocessors and Thread Safety ..... xxx
Using Operators and Generic Functions ..... xxx
Using ScaLAPACK, LAPACK, LINPACK, and EISPACK ..... xxxii
Using ScaLAPACK Enhanced Routines. ..... xxxvi
Chapter 1: Linear Systems ..... 1
Routines ..... 1
Usage Notes ..... 5
Matrix Types .....  5
Solution of Linear Systems ..... 6
Multiple Right Sides ..... 7
Determinants ..... 7
Iterative Refinement .....  7
Matrix Inversion .....  7
Singularity ..... 7
Special Linear Systems. ..... 8
Iterative Solution of Linear Systems ..... 8
QR Decomposition ..... 8
LIN_SOL_GEN .....  9
LIN_SOL_SELF ..... 17
LIN_SOL_LSQ ..... 27
LIN_SOL_SVD ..... 36
LIN_SOL_TRI ..... 44
LIN_SVD ..... 57
Parallel Constrained Least-Squares Solvers ..... 66
Solving Constrained Least-Squares Systems ..... 66
PARALLEL NONNEGATIVE LSQ ..... 66
PARALLEL_BOUNDED LSQ ..... 74
LSARG ..... 82
LSLRG ..... 87
LFCRG ..... 93
LFTRG ..... 98
LFSRG ..... 103
LFIRG ..... 107
LFDRG ..... 113
LINRG ..... 114
LSACG. ..... 118
LSLCG. ..... 123
LFCCG ..... 127
LFTCG ..... 133
LFSCG ..... 138
LFICG ..... 142
LFDCG. ..... 148
LINCG ..... 149
LSLRT ..... 154
LFCRT ..... 157
LFDRT ..... 161
LINRT ..... 163
LSLCT ..... 164
LFCCT ..... 168
LFDCT ..... 172
LINCT ..... 174
LSADS ..... 176
LSLDS ..... 180
LFCDS ..... 185
LFTDS ..... 190
LFSDS. ..... 194
LFIDS ..... 199
LFDDS ..... 204
LINDS ..... 205
LSASF ..... 209
LSLSF ..... 212
LFCSF ..... 214
LFTSF ..... 217
LFSSF ..... 219
LFISF ..... 221
LFDSF ..... 224
LSADH ..... 226
LSLDH. ..... 231
LFCDH ..... 236
LFTDH ..... 241
LFSDH ..... 246
LFIDH ..... 251
LFDDH ..... 256
LSAHF ..... 258
LSLHF ..... 261
LFCHF ..... 263
LFTHF ..... 266
LFSHF ..... 269
LFIHF ..... 271
LFDHF ..... 274
LSLTR ..... 275
LSLCR ..... 277
LSARB ..... 280
LSLRB ..... 282
LFCRB ..... 287
LFTRB ..... 290
LFSRB ..... 293
LFIRB ..... 295
LFDRB ..... 298
LSAQS ..... 300
LSLQS ..... 303
LSLPB ..... 305
LFCQS ..... 308
LFTQS ..... 311
LFSQS ..... 313
LFIQS ..... 315
LFDQS ..... 318
LSLTQ ..... 319
LSLCQ ..... 321
LSACB ..... 324
LSLCB ..... 327
LFCCB ..... 330
LFTCB ..... 333
LFSCB ..... 336
LFICB ..... 338
LFDCB ..... 342
LSAQH ..... 344
LSLQH ..... 346
LSLQB ..... 349
LFCQH ..... 352
LFTQH ..... 355
LFSQH ..... 358
LFIQH ..... 360
LFDQH ..... 362
LSLXG ..... 364
LFTXG ..... 369
LFSXG ..... 374
LSLZG ..... 377
LFTZG ..... 382
LFSZG ..... 387
LSLXD ..... 391
LSCXD ..... 395
LNFXD ..... 399
LFSXD ..... 404
LSLZD ..... 408
LNFZD ..... 412
LFSZD ..... 417
LSLTO ..... 421
LSLTC ..... 423
LSLCC ..... 425
PCGRC ..... 427
JCGRC ..... 433
GMRES ..... 436
LSQRR ..... 446
LQRRV ..... 452
LSBRR ..... 458
LCLSQ ..... 462
LQRRR ..... 466
LQERR. ..... 473
LQRSL ..... 478
LUPQR. ..... 484
LCHRG ..... 489
LUPCH. ..... 491
LDNCH ..... 494
LSVRR ..... 498
LSVCR. ..... 504
LSGRR ..... 508
Chapter 2: Eigensystem Analysis ..... 515
Routines ..... 515
Usage Notes ..... 516
Reformulating Generalized Eigenvalue Problems ..... 519
LIN_EIG SELF ..... 520
LIN_EIG_GEN ..... 527
LIN_GEIG_GEN ..... 535
EVLRG ..... 543
EVCRG ..... 545
EPIRG ..... 548
EVLCG ..... 550
EVCCG ..... 552
EPICG ..... 555
EVLSF ..... 557
EVCSF ..... 559
EVASF ..... 561
EVESF ..... 563
EVBSF ..... 566
EVFSF ..... 568
EPISF ..... 571
EVLSB ..... 573
EVCSB ..... 575
EVASB ..... 578
EVESB ..... 581
EVBSB ..... 584
EVFSB ..... 586
EPISB ..... 589
EVLHF ..... 591
EVCHF ..... 593
EVAHF ..... 596
EVEHF ..... 599
EVBHF ..... 602
EVFHF ..... 604
EPIHF ..... 607
EVLRH ..... 609
EVCRH ..... 611
EVLCH ..... 614
EVCCH ..... 616
GVLRG ..... 618
GVCRG ..... 621
GPIRG ..... 625
GVLCG ..... 627
GVCCG ..... 629
GPICG ..... 632
GVLSP ..... 634
GVCSP ..... 636
GPISP ..... 639
Chapter 3: Interpolation and Approximation ..... 643
Routines ..... 643
Usage Notes ..... 645
Piecewise Polynomials ..... 645
Splines and B-splines ..... 646
Cubic Splines ..... 647
Tensor Product Splines ..... 648
Quadratic Interpolation ..... 649
Scattered Data Interpolation ..... 649
Least Squares. ..... 649
Smoothing by Cubic Splines ..... 649
Rational Chebyshev Approximation ..... 649
Using the Univariate Spline Routines. ..... 649
Choosing an Interpolation Routine ..... 651
SPLINE CONSTRAINTS ..... 652
SPLINE VALUES ..... 653
SPLINE FITTING ..... 654
SURFACE CONSTRAINTS ..... 664
SURFACE_VALUES ..... 665
SURFACE_FITTING ..... 666
CSIEZ ..... 677
CSINT ..... 680
CSDEC ..... 682
CSHER ..... 687
CSAKM ..... 690
CSCON ..... 692
CSPER ..... 696
CSVAL ..... 699
CSDER ..... 700
CS1GD ..... 703
CSITG ..... 706
SPLEZ ..... 708
BSINT ..... 711
BSNAK ..... 715
BSOPK ..... 718
BS2IN ..... 720
BS3IN ..... 725
BSVAL ..... 731
BSDER ..... 732
BS1GD ..... 735
BSITG ..... 738
BS2VL ..... 741
BS2DR ..... 742
BS2GD ..... 746
BS2IG ..... 750
BS3VL ..... 754
BS3DR ..... 756
BS3GD ..... 760
BS3IG ..... 766
BSCPP ..... 770
PPVAL ..... 771
PPDER ..... 774
PP1GD ..... 776
PPITG ..... 780
QDVAL ..... 782
QDDER ..... 784
QD2VL ..... 786
QD2DR ..... 789
QD3VL ..... 792
QD3DR ..... 796
SURF. ..... 800
RLINE ..... 803
RCURV ..... 806
FNLSQ ..... 811
BSLSQ ..... 815
BSVLS ..... 819
CONFT ..... 824
BSLS2 ..... 833
BSLS3 ..... 838
CSSED ..... 844
CSSMH ..... 848
CSSCV ..... 851
RATCH ..... 854
Chapter 4: Integration and Differentiation ..... 859
Routines ..... 859
Usage Notes ..... 860
Univariate Quadrature ..... 860
Multivariate Quadrature ..... 861
Gauss rules and three-term recurrences ..... 861
Numerical differentiation ..... 862
QDAGS ..... 862
QDAG ..... 865
QDAGP ..... 869
QDAGI ..... 872
QDAWO ..... 875
QDAWF ..... 879
QDAWS ..... 883
QDAWC ..... 886
QDNG ..... 889
TWODQ ..... 891
QAND ..... 896
QMC ..... 899
GQRUL ..... 901
GQRCF ..... 905
RECCF ..... 908
RECQR ..... 911
FQRUL ..... 914
DERIV ..... 918
Chapter 5: Differential Equations ..... 923
Routines ..... 923
Usage Notes ..... 923
Ordinary Differential Equations ..... 924
Differential-algebraic Equations ..... 924
Partial Differential Equations ..... 925
Summary ..... 926
IVPRK ..... 927
IVMRK ..... 934
IVPAG ..... 944
BVPFD ..... 961
BVPMS ..... 973
DASPG ..... 980
Introduction to Subroutine PDE_1D_MG ..... 1003
PDE_1D_MG ..... 1004
Description ..... 1011
Remarks on the Examples ..... 1013
Example 1 - Electrodynamics Model ..... 1015
Example 3 - Population Dynamics ..... 1020
Example 4 - A Model in Cylindrical Coordinates ..... 1023
Example 5 - A Flame Propagation Model ..... 1024
Example 6 - A 'Hot Spot' Model ..... 1027
Example 7 - Traveling Waves ..... 1029
Example 9 - Electrodynamics, Parameters Studied with MPI ..... 1033
MOLCH ..... 1038
FPS2H ..... 1053
FPS3H ..... 1059
SLEIG ..... 1066
SLCNT ..... 1078
Chapter 6: Transforms ..... 1083
Routines ..... 1083
Usage Notes ..... 1084
Fast Fourier Transforms ..... 1084
Continuous versus Discrete Fourier Transform ..... 1085
Inverse Laplace Transform ..... 1086
FAST DFT ..... 1086
FAST_2DFT ..... 1093
FAST_3DFT ..... 1099
FFTRF ..... 1103
FFTRB ..... 1106
FFTRI. ..... 1109
FFTCF ..... 1111
FFTCB ..... 1113
FFTCI ..... 1116
FSINT ..... 1118
FSINI. ..... 1120
FCOST ..... 1122
FCOSI ..... 1124
QSINF ..... 1126
QSINB ..... 1129
QSINI ..... 1131
QCOSF ..... 1133
QCOSB ..... 1135
QCOSI. ..... 1137
FFT2D ..... 1139
FFT2B ..... 1142
FFT3F ..... 1145
FFT3B ..... 1149
RCONV ..... 1153
CCONV ..... 1158
RCORL ..... 1163
CCORL ..... 1168
INLAP ..... 1172
SINLP ..... 1175
Chapter 7: Nonlinear Equations ..... 1183
Routines ..... 1183
Usage Notes ..... 1183
Zeros of a Polynomial ..... 1183
Zero(s) of a Function ..... 1184
Root of System of Equations ..... 1184
ZPLRC ..... 1184
ZPORC ..... 1186
ZPOCC ..... 1188
ZANLY ..... 1189
ZBREN ..... 1192
ZREAL ..... 1195
NEQNF ..... 1198
NEQNJ ..... 1201
NEQBF ..... 1204
NEQBJ ..... 1210
Chapter 8: Optimization ..... 1217
Routines ..... 1217
Usage Notes ..... 1218
Unconstrained Minimization ..... 1218
Minimization with Simple Bounds ..... 1219
Linearly Constrained Minimization ..... 1219
Nonlinearly Constrained Minimization ..... 1219
Selection of Routines ..... 1220
UVMIF ..... 1222
UVMID ..... 1225
UVMGS ..... 1229
UMINF ..... 1232
UMING ..... 1237
UMIDH ..... 1243
UMIAH ..... 1249
UMCGF ..... 1255
UMCGG ..... 1259
UMPOL ..... 1263
UNLSF ..... 1267
UNLSJ ..... 1273
BCONF ..... 1279
BCONG ..... 1286
BCODH ..... 1293
BCOAH ..... 1299
BCPOL ..... 1306
BCLSF ..... 1310
BCLSJ ..... 1317
BCNLS ..... 1324
READ_MPS ..... 1333
MPS File Format ..... 1337
NAME Section ..... 1338
ROWS Section. ..... 1338
COLUMNS Section ..... 1339
RHS Section ..... 1339
RANGES Section ..... 1340
BOUNDS Section ..... 1340
QUADRATIC Section. ..... 1342
ENDATA Section ..... 1342
MPS FREE ..... 1343
DENSE LP ..... 1346
DLPRS ..... 1351
SLPRS ..... 1355
QPROG ..... 1361
LCONF. ..... 1364
LCONG ..... 1370
NNLPF ..... 1377
NNLPG ..... 1383
CDGRD ..... 1390
FDGRD ..... 1392
FDHES ..... 1394
GDHES ..... 1397
FDJAC ..... 1400
CHGRD ..... 1403
CHHES ..... 1406
CHJAC ..... 1410
GGUES ..... 1414
Chapter 9: Basic Matrix/Vector Operations ..... 1419
Routines ..... 1419
Basic Linear Algebra Subprograms ..... 1422
Programming Notes for Level 1 BLAS ..... 1422
Descriptions of the Level 1 BLAS Subprograms ..... 1423
Programming Notes for Level 2 and Level 3 BLAS ..... 1433
Descriptions of the Level 2 and Level 3 BLAS ..... 1434
Other Matrix/Vector Operations ..... 1443
CRGRG ..... 1444
CCGCG ..... 1445
CRBRB ..... 1447
CCBCB ..... 1448
CRGRB ..... 1450
CRBRG ..... 1452
CCGCB ..... 1453
CCBCG ..... 1455
CRGCG ..... 1457
CRRCR ..... 1458
CRBCB ..... 1460
CSFRG ..... 1462
CHFCG ..... 1463
CSBRB ..... 1465
CHBCB ..... 1467
TRNRR ..... 1469
MXTXF ..... 1470
MXTYF ..... 1472
MXYTF ..... 1474
MRRRR ..... 1476
MCRCR ..... 1479
HRRRR ..... 1481
BLINF ..... 1483
POLRG ..... 1485
MURRV ..... 1487
MURBV ..... 1489
MUCRV ..... 1491
MUCBV ..... 1493
ARBRB ..... 1495
ACBCB ..... 1497
NRIRR ..... 1499
NR1RR ..... 1501
NR2RR ..... 1502
NR1RB ..... 1504
NR1CB ..... 1505
DISL2 ..... 1507
DISL1 ..... 1509
DISLI ..... 1510
VCONR ..... 1512
VCONC ..... 1514
Extended Precision Arithmetic ..... 1516
Chapter 10: Linear Algebra Operators and Generic Functions ..... 1521
Routines ..... 1521
Usage Notes ..... 1522
Matrix Optional Data Changes ..... 1522
Dense Matrix Computations ..... 1524
Dense Matrix Functions ..... 1526
Dense Matrix Parallelism Using MPI ..... 1527
General Remarks ..... 1527
Getting Started with Modules MPI setup int and MPI node int ..... 1527
Using Processors ..... 1529
Sparse Matrix Computations ..... 1530
Introduction ..... 1530
Derived Type Definitions ..... 1532
Type SLU_options ..... 1533
Overloaded Assignments ..... 1534
.X. ..... 1537
.tx ..... 1541
.xt ..... 1544
.hx ..... 1547
.xh ..... 1550
.t. ..... 1553
.h. ..... 1556
.i. ..... 1558
.ix ..... 1561
.xi. ..... 1571
CHOL ..... 1574
COND ..... 1577
DET ..... 1581
DIAG ..... 1584
DIAGONALS ..... 1585
EIG ..... 1586
EYE ..... 1590
FFT ..... 1592
FFT BOX ..... 1594
IFFT ..... 1596
IFFT BOX ..... 1598
isNaN ..... 1600
NaN ..... 1601
NORM ..... 1602
ORTH ..... 1605
RAND ..... 1608
RANK ..... 1610
SVD ..... 1612
UNIT ..... 1614
Chapter 11: Utilities ..... 1617
Routines ..... 1617
Usage Notes for ScaLAPACK Utilities ..... 1619
ScaLAPACK Supporting Modules ..... 1622
ScaLAPACK SETUP ..... 1622
ScaLAPACK GETDIM ..... 1624
ScaLAPACK READ ..... 1625
ScaLAPACK WRITE ..... 1627
ScaLAPACK MAP ..... 1636
ScaLAPACK UNMAP ..... 1637
ScaLAPACK EXIT ..... 1640
ERROR POST ..... 1640
SHOW ..... 1643
WRRRN ..... 1647
WRRRL ..... 1649
WRIRN ..... 1653
WRIRL ..... 1655
WRCRN ..... 1658
WRCRL ..... 1660
WROPT ..... 1664
PGOPT ..... 1671
PERMU ..... 1673
PERMA ..... 1674
SORT_REAL ..... 1677
SVRGN ..... 1679
SVRGP ..... 1681
SVIGN ..... 1683
SVIGP ..... 1684
SVRBN ..... 1685
SVRBP ..... 1687
SVIBN ..... 1688
SVIBP ..... 1690
SRCH ..... 1691
ISRCH ..... 1694
SSRCH ..... 1696
ACHAR ..... 1698
IACHAR ..... 1699
ICASE ..... 1700
IICSR ..... 1701
IIDEX ..... 1703
CVTSI ..... 1704
CPSEC ..... 1705
TIMDY ..... 1706
TDATE ..... 1707
NDAYS ..... 1708
NDYIN. ..... 1710
IDYWK ..... 1711
VERML ..... 1713
RAND GEN ..... 1714
RNGET ..... 1722
RNSET ..... 1723
RNOPT ..... 1724
RNIN32 ..... 1726
RNGE32 ..... 1727
RNSE32 ..... 1729
RNIN64 ..... 1729
RNGE64 ..... 1730
RNSE64 ..... 1732
RNUNF ..... 1732
RNUN ..... 1734
FAURE_INIT ..... 1736
FAURE_FREE ..... 1736
FAURE NEXT ..... 1737
IUMAG ..... 1739
UMAG ..... 1743
SUMAG/DUMAG ..... 1746
PLOTP ..... 1746
PRIME ..... 1749
CONST. ..... 1751
CUNIT ..... 1753
HYPOT ..... 1757
MP_SETUP ..... 1758
Reference Material ..... 1763
Contents ..... 1763
User Errors ..... 1763
What Determines Error Severity ..... 1763
Kinds of Errors and Default Actions ..... 1764
Errors in Lower-Level Routines ..... 1765
Routines for Error Handling ..... 1765
ERSET ..... 1765
IERCD and N1RTY ..... 1766
Examples ..... 1766
Machine-Dependent Constants ..... 1769
IMACH ..... 1769
AMACH ..... 1771
DMACH ..... 1772
IFNAN(X) ..... 1772
UMACH ..... 1774
Matrix Storage Modes ..... 1775
Reserved Names ..... 1784
Deprecated Features and Renamed Routines ..... 1785
Automatic Workspace Allocation ..... 1785
Changing the Amount of Space Allocated. ..... 1786
Character Workspace ..... 1787
Appendix A: GAMS Index ..... A-1
Description ..... A-1
IMSL MATH/LIBRARY ..... A-2
Appendix B: Alphabetical Summary of Routines ..... B-1
Routines ..... B-1
Appendix C: References ..... C-1
Appendix D: Benchmarking or Timing Programs ..... D-1
Scalar Program Descriptions ..... D-1
Parallel Program Descriptions. ..... D-5
Product Support ..... i
Contacting Visual Numerics Support .....  i

Index

## Introduction

## The IMSL Fortran Numerical Library

The IMSL Fortran Numerical Library consists of two separate but coordinated Libraries that allow easy user access. These Libraries are organized as follows:

- MATH/LIBRARY general applied mathematics and special functions

The User's Guide for IMSL MATH/LIBRARY has two parts:

1. MATH/LIBRARY (Volumes 1, 2, and 3)
2. MATH/LIBRARY Special Functions

- STAT/LIBRARY statistics (Volumes 1, and 2)

Most of the routines are available in both single and double precision versions. Many routines for linear solvers and eigensystems are also available for complex and double -complex precision arithmetic. The same user interface is found on the many hardware versions that span the range from personal computer to supercomputer.
This library is the result of a merging of the products: IMSL Fortran Numerical Libraries and IMSL Fortran 90 Library.

## User Background

To use this product you should be familiar with the Fortran 90 language as well as the withdrawn Fortran 77 language, which is, in practice, a subset of Fortran 90. A summary of the ISO and ANSI standard language is found in Metcalf and Reid (1990). A more comprehensive illustration is given in Adams et al. (1992).

Those routines implemented in the IMSL Fortran Numerical Library provide a simpler, more reliable user interface than was possible with Fortran 77. Features of the IMSL Fortran Numerical Library include the use of descriptive names, short required argument lists, packaged userinterface blocks, a suite of testing and benchmark software, and a collection of examples. Source code is provided for the benchmark software and examples.
Some of the routines in the IMSL Fortran Numerical Library can take advantage of a standard (MPI) Message Passing Interface environment but do not require an MPI environment if the user chooses to not take advantage of MPI.

The MPI logo shown below cues the reader when this is the case:

Routines documented with the MPI Capable logo can be called in a scalar or one computer environment.

Other routines in the IMSL Library take advantage of MPI and require that an MPI environment be present in order to use them. The MPI Required logo shown below clues the reader when this is the case:


NOTE: It is recommended that users considering using the MPI capabilities of the product read the following sections of the MATH Library documentation:

Introduction: Using MPI Routines,
Introduction: Using ScaLAPACK Enhanced Routines,
Chapter 10, Linear Algebra Operators and Generic Functions - see Dense Matrix Parallelism Using MPI.

## Getting Started

The IMSL MATH/LIBRARY is a collection of Fortran routines and functions useful in mathematical analysis research and application development. Each routine is designed and documented for use in research activities as well as by technical specialists.

To use any of these routines, you must write a program in Fortran 90 (or possibly some other language) to call the MATH/LIBRARY routine. Each routine conforms to established conventions in programming and documentation. We give first priority in development to efficient algorithms, clear documentation, and accurate results. The uniform design of the routines makes it easy to use more than one routine in a given application. Also, you will find that the design consistency enables you to apply your experience with one MATH/LIBRARY routine to other IMSL routines that you use.

## Finding the Right Routine

The MATH/LIBRARY is organized into chapters; each chapter contains routines with similar computational or analytical capabilities. To locate the right routine for a given problem, you may use either the table of contents located in each chapter introduction, or the alphabetical list of routines. The GAMS index uses GAMS classification (Boisvert, R.F., S.E. Howe, D.K. Kahaner, and J. L. Springmann 1990, Guide to Available Mathematical Software, National Institute of Standards and Technology NISTIR 90-4237). Use the GAMS index to locate which MATH/LIBRARY routines pertain to a particular topic or problem.

Often the quickest way to use the MATH/LIBRARY is to find an example similar to your problem and then to mimic the example. Each routine document has at least one example demonstrating its application. The example for a routine may be created simply for illustration, it may be from a textbook (with reference to the source), or it may be from the mathematical literature.

## Organization of the Documentation

This manual contains a concise description of each routine, with at least one demonstrated example of each routine, including sample input and results. You will find all information pertaining to the MATH/LIBRARY in this manual. Moreover, all information pertaining to a particular routine is in one place within a chapter.

Each chapter begins with an introduction followed by a table of contents that lists the routines included in the chapter. Documentation of the routines consists of the following information:

- IMSL Routine's Generic Name
- Purpose: a statement of the purpose of the routine. If the routine is a function rather than a subroutine the purpose statement will reflect this fact.
- Function Return Value: a description of the return value (for functions only).
- Required Arguments: a description of the required arguments in the order of their occurrence. Input arguments usually occur first, followed by input/output arguments, with output arguments described last. Futhermore, the following terms apply to arguments:

Input Argument must be initialized; it is not changed by the routine.
Input/Output Argument must be initialized; the routine returns output through this argument; cannot be a constant or an expression.

Input[/Output] Argument must be initialized; the routine may return output through this argument based on other optional data the user may choose to pass to this routine; cannot be a constant or an expression.

Input or Output Select appropriate option to define the argument as either input or output. See individual routines for further instructions.

Output No initialization is necessary; cannot be a constant or an expression. The routine returns output through this argument.

- Optional Arguments: a description of the optional arguments in the order of their occurrence.
- Fortran 90 Interface: a section that describes the generic and specific interfaces to the routine.
- Fortran 77 Style Interface: an optional section, which describes Fortran 77 style interfaces, is supplied for backwards compatibility with previous versions of the Library.
- ScaLAPACK Interface: an optional section, which describes an interface to a ScaLAPACK based version of this routine.
- Description: a description of the algorithm and references to detailed information. In many cases, other IMSL routines with similar or complementary functions are noted.
- Comments: details pertaining to code usage.
- Programming notes: an optional section that contains programming details not covered elsewhere.
- Example: at least one application of this routine showing input and required dimension and type statements.
- Output: results from the example(s).
- Additional Examples: an optional section with additional applications of this routine showing input and required dimension and type statements.


## Naming Conventions

The names of the routines are mnemonic and unique. Most routines are available in both a single precision and a double precision version, with names of the two versions sharing a common root. The root name is also the generic interface name. The name of the double precision specific version begins with a "D_." The single precision specific version begins with an "s_". For example, the following pairs are precision specific names of routines in the two different precisions: S_GQRUL/D_GQRUL (the root is "GQRUL ," for "Gauss quadrature rule") and S_RECCF/D_RECCF (the root is "RECCF," for "recurrence coefficient"). The precision specific names of the IMSL routines that return or accept the type complex data begin with the letter "C_" or "z_" for complex or double complex, respectively. Of course, the generic name can be used as an entry point for all precisions supported.

When this convention is not followed the generic and specific interfaces are noted in the documentation. For example, in the case of the BLAS and trigonometric intrinsic functions where standard names are already established, the standard names are used as the precision specific names. There may also be other interfaces supplied to the routine to provide for backwards compatibility with previous versions of the IMSL Fortran Numerical Library. These alternate interfaces are noted in the documentation when they are available.

Except when expressly stated otherwise, the names of the variables in the argument lists follow the Fortran default type for integer and floating point. In other words, a variable whose name begins with one of the letters " $I$ " through " $N$ " is of type INTEGER, and otherwise is of type REAL or DOUBLE PRECISION, depending on the precision of the routine.

An assumed-size array with more than one dimension that is used as a Fortran argument can have an assumed-size declarator for the last dimension only. In the MATH/LIBRARY routines, the information about the first dimension is passed by a variable with the prefix "LD" and with the array name as the root. For example, the argument LDA contains the leading dimension of array $A$. In most cases, information about the dimensions of arrays is obtained from the array through the use of Fortran 90 's size function. Therefore, arguments carrying this type of information are usually defined as optional arguments.
Where appropriate, the same variable name is used consistently throughout a chapter in the MATH/LIBRARY. For example, in the routines for random number generation, NR denotes the number of random numbers to be generated, and $R$ or IR denotes the array that stores the numbers.

When writing programs accessing the MATH/LIBRARY, the user should choose Fortran names that do not conflict with names of IMSL subroutines, functions, or named common blocks. The
careful user can avoid any conflicts with IMSL names if, in choosing names, the following rules are observed:

- Do not choose a name that appears in the Alphabetical Summary of Routines, at the end of the User's Manual, nor one of these names preceded by a D, S_, D_, C_, or Z_.
- Do not choose a name consisting of more than three characters with a numeral in the second or third position.
For further details, see the section on "Reserved Names" in the Reference Material.


## Using Library Subprograms

The documentation for the routines uses the generic name and omits the prefix, and hence the entire suite of routines for that subject is documented under the generic name.

Examples that appear in the documentation also use the generic name. To further illustrate this principle, note the LIN_SOL_GEN documentation (see Chapter 1, Linear Systems), for solving general systems of linear algebraic equations. A description is provided for just one data type. There are four documented routines in this subject area: s_lin_sol_gen, d_lin_sol_gen, c_lin_sol_gen, and z_lin_sol_gen.

These routines constitute single-precision, double-precision, complex, and double-complex precision versions of the code.

The Fortran 90 compiler identifies the appropriate routine. Use of a module is required with the routines. The naming convention for modules joins the suffix "_int" to the generic routine name. Thus, the line "use lin_sol_gen_int" is inserted near the top of any routine that calls the subprogram "lin_sol_gen". More inclusive modules are also available, such as imsl_libraries and numerical libraries. To avoid name conflicts, Fortran 90 permits relabeling names defined in modules so they do not conflict with names of routines or variables in the user's program. The user can also restrict access to names defined in IMSL Library modules by use of the ": ONLY, <list of names>" qualifier.

When dealing with a complex matrix, all references to the transpose of a matrix, $A^{T}$, are replaced by the adjoint matrix

$$
\bar{A}^{T} \equiv A^{*}=A^{H}
$$

where the overstrike denotes complex conjugation. IMSL Fortran Numerical Library linear algebra software uses this convention to conserve the utility of generic documentation for that code subject. All references to orthogonal matrices are to be replaced by their complex counterparts, unitary matrices. Thus, an $n \times n$ orthogonal matrix $Q$ satisfies the condition $Q^{T} Q=I_{n}$. An $n \times n$ unitary matrix $V$ satisfies the analogous condition for complex matrices, $V^{*} V=I_{n}$.

## Programming Conventions

In general, the IMSL MATH/LIBRARY codes are written so that computations are not affected by underflow, provided the system (hardware or software) places a zero value in the register. In this case, system error messages indicating underflow should be ignored.
IMSL codes are also written to avoid overflow. A program that produces system error messages indicating overflow should be examined for programming errors such as incorrect input data, mismatch of argument types, or improper dimensioning.

In many cases, the documentation for a routine points out common pitfalls that can lead to failure of the algorithm.
Library routines detect error conditions, classify them as to severity, and treat them accordingly. This error-handling capability provides automatic protection for the user without requiring the user to make any specific provisions for the treatment of error conditions. See the section on "User Errors" in the Reference Material for further details.

## Module Usage

Users are required to incorporate a "use" statement near the top of their program for the IMSL routine being called when writing new code that uses this library. However, legacy code which calls routines in the previous version of the library without the use of a "use" statement will continue to work as before. Also, code that employed the "use numerical_libraries" statement from the previous version of the library will continue to work properly with this version of the library.

Users wishing to update existing programs so as to call other routines from this library should incorporate a use statement for the specific new routine being called. (Here, the term "new routine" implies any routine in the library, only "new" to the user's program.) Use of the more encompassing "imsl_libraries" module in this case could result in argument mismatches for the "old" routine(s) being called. (The compiler would catch this.)

Users wishing to update existing programs to call the new generic versions of the routines must change their calls to the existing routines to match the new calling sequences and use either the routine specific interface modules or the all-encompassing "imsl_libraries" module.

## Using MPI Routines

REQUIRED
Users of the IMSL Fortran Numerical Library benefit by having a standard (MPI) Message Passing Interface environment. This is needed to accomplish parallel computing within parts of the Library. Either of the icons above clues the reader when this is the case. If parallel computing is not required, then the IMSL Library suite of dummy MPI routines can be substituted for standard MPI routines. All requested MPI routines called by the IMSL Library are in this dummy suite. Warning messages will appear if a code or example requires more than one process to execute. Typically users need not be aware of the parallel codes.

NOTE: that a standard MPI environment is not part of the IMSL Fortran Numerical Library. The standard includes a library of MPI Fortran and C routines, MPI "include" files, usage documentation, and other run-time utilities.

NOTE: Details on linking to the appropriate libraries are explained in the online README file of the product distribution.

There are three situations of MPI usage in the IMSL Fortran Numerical Library:

1. There are some computations that are performed with the 'box' data type that benefit from the use of parallel processing. For computations involving a single array or a single problem, there is no IMSL use of parallel processing or MPI codes. The box type data type implies that several problems of the same size and type are to be computed and solved. Each rack of the box is an independent problem. This means that each problem could potentially be solved in parallel. The default for computing a box data type calculation is that a single processor will do all of the problems, one after the other. If this is acceptable there should be no further concern about which version of the libraries is used for linking. If the problems are to be solved in parallel, then the user must link with a working version of an MPI Library and the appropriate IMSL Library. Examples demonstrating the use of box type data may be found in Chapter 10, "Linear Algebra Operators and Generic Functions".

NOTE: Box data type routines are marked with the MPI Capable icon.
2. Various routines in Chapter 1, "Linear Systems" allow the user to interface with the ScaLAPACK Library routines. If the user chooses to run on only one processor then these routines will utilize either IMSL Library code or LAPACK Library code based on the libraries the user chooses to use during linking. If the user chooses to run on multiple processors then working versions of MPI, ScaLAPACK, PBLAS, and Blacs will need to be present. These routines are marked with the MPI Capable icon.
3. There are some routines or operators in the Library that require that a working MPI Library be present in order for them to run. Examples are the large-scale parallel solvers and the ScaLAPACK utilities. Routines of this type are marked with the MPI Required icon. For these routines, the user must link with a working version of an MPI Library and the appropriate IMSL Library.

In all cases described above it is the user's responsibility to supply working versions of the aforementioned third party libraries when those libraries are required.

Table A below lists the chapters and IMSL routines calling MPI routines or the replacement nonparallel package.

| Chapter Name and Number | Routine with MPI Utilized |
| :--- | :--- |
| Linear Systems, 1 | PARALLEL_NONNEGATIVE_LSQ |
| Linear Systems, 1 | PARALLEL_BOUNDED_LSQ |
| Linear Systems, 1 | Those routines which utilize ScaLAPACK <br> listed in Table D below. |


| Chapter Name and Number | $\quad$ Routine with MPI Utilized |
| :--- | :--- |
| Linear Algebra and Generic Functions, 10 | See entire following <br> Table B.1 - Defined Operators and Generic Functions |
| Utilities, 11 | ScaLAPACK_SETUP |
| Utilities, 11 | ScaLAPACK_GETDIM |
| Utilities, 11 | ScaLAPACK_READ |
| Utilities, 11 | ScaLAPACK_WRITE |
| Utilities, 11 | ScaLAPACK_MAP |
| Utilities, 11 | ScaLAPACK_UNMAP |
| Utilities, 11 | ScaLAPACK_EXIT |
| Reference Material | Entire Error Processor Package for IMSL <br> Library, if MPI is utilized |

Table A - IMSL Routines Calling MPI Routines or Replacement Non-Parallel Package

## Programming Tips

Each subject routine called or otherwise referenced requires the "use" statement for an interface block designed for that subject routine. The contents of this interface block are the interfaces to the separate routines available for that subject. Packaged descriptive names for option numbers that modify documented optional data or internal parameters might also be provided in the interface block. Although this seems like an additional complication, many errors are avoided at an early stage in development through the use of these interface blocks. The "use" statement is required for each routine called in the user's program. As illustrated in Examples 3 and 4 in routine lin_geig_gen, the "use" statement is required for defining the secondary option flags.

The function subprogram for s_NaN() or d_NaN() does not require an interface block because it has only a single "required" dummy argument. Also, if one is only using the Fortran 77 interfaces supplied for backwards compatibility then the "use" statements are not required.

## Optional Subprogram Arguments

IMSL Fortran Numerical Library routines have required arguments and may have optional arguments. All arguments are documented for each routine. For example, consider the routine lin_sol_gen that solves the linear algebraic matrix equation $A x=b$. The required arguments are three rank-2 Fortran 90 arrays: $A, b$, and $x$. The input data for the problem are the $A$ and $b$ arrays; the solution output is the $x$ array. Often there are other arguments for this linear solver that are closely connected with the computation but are not as compelling as the primary problem. The inverse matrix $A^{-1}$ may be needed as part of a larger application. To output this parameter, use the optional argument given by the "ainv=" keyword. The rank-2 output array argument used on the right-hand side of the equal sign contains the inverse matrix. See Example 2 in Chapter 1, "Linear Systems" of LIN_SOL_GEN for an example of computing the inverse matrix.

For compatibility with previous versions of the IMSL Libraries, the NUMERICAL_LIBRARIES interface module includes backwards-compatible positional argument interfaces to all routines that
existed in the Fortran 77 version of the Library. Note that it is not necessary to include "use" statements when calling these routines by themselves. Existing programs that called these routines will continue to work in the same manner as before.

Some of the primary routines have arguments "epack=" and "iopt=". As noted the "epack=" argument is of derived type s_error or d_error. The prefix "s_" or "d_" is chosen depending on the precision of the data type for that routine. These optional arguments are part of the interface to certain routines, and are used to modify internal algorithm choices or other parameters.

## Optional Data

This additional optional argument (available for some routines) is further distinguished-a derived type array that contains a number of parameters to modify the internal algorithm of a routine. This derived type has the name ? options, where "? " is either "s_" or "d_". The choice depends on the precision of the data type. The declaration of this derived type is packaged within the modules for these codes.

The definition of the derived types is:

```
type ?_options
    integer idummy; real(kind(?)) rdummy
end type
```

where the "? " is either "s_" or "d_", and the kind value matches the desired data type indicated by the choice of " $s$ " or " $d$ ".
Example 3 in Chapter 1, "Linear Systems" of LIN_SOL_GEN illustrates the use of iterative refinement to compute a double-precision solution $\bar{b}$ based on a single-precision factorization of the matrix. This is communicated to the routine using an optional argument with optional data. For efficiency of iterative refinement, perform the factorization step once, and then save the factored matrix in the array $A$ and the pivoting information in the rank- 1 integer array, ipivots. By default, the factorization is normally discarded. To enable the routine to be re-entered with a previously computed factorization of the matrix, optional data are used as array entries in the "iopt=" optional argument. The packaging of LIN_SOL_GEN includes the definitions of the selfdocumenting integer parameters lin_sol_gen_save_Lu and lin_sol_gen_solve_A. These parameters have the values 2 and 3 , but the programmer usually does not need to be aware of it. The following rules apply to the "iopt=iopt" optional argument:

1. Define a relative index, for example IO, for placing option numbers and data into the array argument iopt. Initially, set $I O=1$. Before a call to the IMSL Library routine, follow Steps 2 through 4.
2. The data structure for the optional data array has the following form:
iopt (IO) = ?_options (Option_number, Optional_data)
[iopt $(\mathrm{IO}+1)=$ ? _options (Option_number, Optional_data)]
The length of the data set is specified by the documentation for an individual routine. (The Optional_data is output in some cases and may not be used in other cases.) The square braces [. . .] denote optional items.

Illustration: Example 3 in Chapter 2, "Singular Value and Eigenvalue Decomposition" of

LIN_EIG_SELF, a new definition for a small diagonal term is passed to LIN_SOL_SELF. There is one line of code required for the change and the new tolerance:

```
iopt (1) = d_options(d_lin_sol_self_set_small,
epsilon(one)}\mp@subsup{}{}{*}\mathrm{ *abs (d(i)}
```

3. The internal processing of option numbers stops when Option_number $=0$ or when IO > SIZE(iopt). This signals each routine having this optional argument that all desired changes to default values of internal parameters have been made. This implies that the last option number is the value zero or the value of SIZE (iopt) matches the last optional value changed.
4. To add more options, replace $I O$ with $I O+n$, where $n$ is the number of items required for the previous option. Go to Step 2.

Option numbers can be written in any order, and any selected set of options can be changed from the defaults. They may be repeated. Example 3 in Chapter 1, "Linear Solvers" of LIN_SOL_SELF uses three and then four option numbers for purposes of computing an eigenvector associated with a known eigenvalue.

## Overloaded =, l=, etc., for Derived Types

To assist users in writing compact and readable code, the IMSL Fortran Numerical Library provides overloaded assignment and logical operations for the derived types s_options, d_options, s_error, and d_error. Each of these derived types has an individual record consisting of an integer and a floating-point number. The components of the derived types, in all cases, are named idummy followed by rdummy. In many cases, the item referenced is the component idummy. This integer value can be used exactly as any integer by use of the component selector character (\%). Thus, a program could assign a value and test after calling a routine:

```
s_epack(1)%idummy = 0
call lin_sol_gen(A,b,x,epack=s_epack)
if (s_epack(1)%idummy > 0) call error_post(s_epack)
```

Using the overloaded assignment and logical operations, this code fragment can be written in the equivalent and more readable form:

```
s_epack(1) = 0
call lin_sol_gen(A,b,x,epack=s_epack)
if (s_epack(1) > 0) call error_post(s_epack)
```

Generally the assignments and logical operations refer only to component idummy. The assignment "s_epack (1)=0" is equivalent to "s_epack (1)=s_error (0,0E0)". Thus, the floating-point component rdummy is assigned the value 0 E 0 . The assignment statement "I=s_epack (1)", for I an integer type, is equivalent to "I=s_epack (1) \%idummy". The value of component rdummy is ignored in this assignment. For the logical operators, a single element of any of the IMSL Fortran Numerical Library derived types can be in either the first or second operand.

| Derived Type | Overloaded Assignments and Tests |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| s_options | I=s_options (1) ; s_options (1)=I | $=$ | /= | < | <= | > | >= |
| s_options | I=d_options (1);d_options (1) = I | $=$ | /= | < | <= | $>$ | >= |



In the examples, operator_ex $01, \ldots$, ex 37 , the overloaded assignments and tests have been used whenever they improve the readability of the code.

## Error Handling



The routines in the IMSL MATH/LIBRARY attempt to detect and report errors and invalid input. Errors are classified and are assigned a code number. By default, errors of moderate or worse severity result in messages being automatically printed by the routine. Moreover, errors of worse severity cause program execution to stop. The severity level and the general nature of the error are designated by an "error type" ranging from 0 to 5 . An error type 0 is no error; types 1 through 5 are progressively more severe. In most cases, you need not be concerned with our method of handling errors. For those interested, a complete description of the error-handling system is given in the Reference Material, which also describes how you can change the default actions and access the error code numbers.

A separate error handler is provided to allow users to handle errors of differing types being reported from several nodes without danger of "jumbling" or mixing error messages. The design of this error handler is described more fully in Hanson (1992). The primary feature of the design is the use of a separate array for each parallel call to a routine. This allows the user to summarize errors using the routine error_post in a non-parallel part of an application. For a more detailed discussion of the use of this error handler in applications which use MPI for distributed computing, see the Reference Material.

## Printing Results

Most of the routines in the IMSL MATH/LIBRARY (except the line printer routines and special utility routines) do not print any of the results. The output is returned in Fortran variables, and you can print these yourself. See Chapter 11, "Utilities," for detailed descriptions of these routines.

A commonly used routine in the examples is the IMSL routine UMACH (see the Reference Material), which retrieves the Fortran device unit number for printing the results. Because this routine obtains device unit numbers, it can be used to redirect the input or output. The section on "MachineDependent Constants" in the Reference Material contains a description of the routine UMACH.

## Fortran 90 Constructs

The IMSL Fortran Numerical Library contains routines which take advantage of Fortran 90 language constructs, including Fortran 90 array data types. One feature of the design is that the default use may be as simple as the problem statement. Complicated, professional-quality mathematical software is hidden from the casual or beginning user.

In addition, high-level operators and functions are provided in the Library. They are described in Chapter 10, "Linear Algebra Operators and Generic Functions".

## Shared-Memory Multiprocessors and Thread Safety

The IMSL Fortran Numerical Library allows users to leverage the high-performance technology of shared memory parallelism (SMP) when their environment supports it. Support for SMP systems within the IMSL Library is delivered through various means, depending upon the availability of technologies such as OpenMP, high performance LAPACK and BLAS, and hardware-specific IMSL algorithms. Use of the IMSL Fortran Numerical Library on SMP systems can be achieved by using the appropriate link environment variable when building your application. Details on the available link environment variables for your installation of the IMSL Fortran Numerical Library can be found in the online README file of the product distribution.
The IMSL Fortran Numerical Library is thread-safe in those environments that support OpenMP 2.0. This was achieved by using OpenMP directives that define global variables located in the code so they are private to the individual threads. Thread safety allows users to create instances of routines running on multiple threads and to include any routine in the IMSL Fortran Numerical Library in these threads.

## Using Operators and Generic Functions

For users who are primarily interested in easy-to-use software for numerical linear algebra, see Chapter 10, "Linear Algebra Operators and Generic Functions". This compact notation for writing Fortran 90 programs, when it applies, results in code that is easier to read and maintain than traditional subprogram usage.

Users may begin their code development using operators and generic functions. If a more efficient executable code is required, a user may need to switch to equivalent subroutine calls using IMSL Fortran Numerical Library routines.
Table B contain lists of the defined operators and some of their generic functions.

| Defined Array Operation | Matrix Operation |
| :---: | :---: |
| A .x. B | $A B$ |
| .i. A | $A^{-1}$ |
| .t. A, h. A | $A^{T}, A^{*}$ |
| A.ix. B | $A^{-1} B$ |
| B .xi. A | $B A^{-1}$ |
| $\begin{aligned} & \text { A.tx. B, or (.t. A) } \\ & \text { A.x. B } \\ & \text { A. B, or (.h. A) } . x . B \end{aligned}$ | $A^{T} B, A^{*} B$ |
| $\begin{array}{ll} \text { B .xt. A, or B .x. } & \text { (.t. A) } \\ \text { B .xh. A, or B .x. } & \text { (.h. A) } \end{array}$ | $B A^{T}, B A^{*}$ |
| $\mathrm{S}=\mathrm{SVD}(\mathrm{A} \quad[, \mathrm{U}=\mathrm{U}, \mathrm{V}=\mathrm{V}])$ | $A=U S V^{T}$ |
| $\mathrm{E}=\mathrm{EIG}(\mathrm{A} \quad[\mathrm{l}, \mathrm{B}=\mathrm{B}, \mathrm{D}=\mathrm{D}], \mathrm{V}=\mathrm{V}, \mathrm{W}=\mathrm{W}]$ ) | $(A V=V E), A V D=B V E,(A W=W E), A W D=B W E$ |
| $\mathrm{R}=\mathrm{CHOL}$ ( A ) | $A=R^{T} R$ |
| $\mathrm{Q}=$ ORTH ( $\mathrm{A} \quad[, \mathrm{R}=\mathrm{R}]$ ) | $(A=Q R), Q^{T} Q=I$ |
| U=UNIT ( A ) | $\left[u_{1}, \ldots\right]=\left[a_{1} /\left\\|a_{1}\right\\|, \ldots\right]$ |
| $\mathrm{F}=\mathrm{DET}$ ( A ) | $\operatorname{det}(A)=$ determinant |
| K=RANK (A) | $\operatorname{rank}(A)=\operatorname{rank}$ |
| $\mathrm{P}=\mathrm{NORM}(\mathrm{A}$ [,[type=]i]) | $\begin{aligned} & p=\\|A\\|_{1}=\max _{j}\left(\sum_{i=1}^{m}\left\|a_{i j}\right\|\right) \\ & p=\\|A\\|_{2}=s_{1}=\text { largest singular value } \\ & p=\\|A\\|_{\infty \leftrightarrow \text { huge(1) }}=\max _{i}\left(\sum_{j=1}^{n}\left\|a_{i j}\right\|\right) \end{aligned}$ |
| $\mathrm{C}=\mathrm{COND}$ ( A ) | $\left\\|A^{-1}\right\\| \cdot\\|A\\|$ |
| $\mathrm{Z}=\mathrm{EYE}$ ( N ) | $Z=I_{N}$ |
| A=DIAG (X) | $A=\operatorname{diag}\left(x_{1}, \ldots\right)$ |
| X=DIAGONALS (A) | $x=\left(a_{11}, \ldots\right)$ |
| W=FFT (Z) ; Z=IFFT (W) | Discrete Fourier Transform, Inverse |
| A=RAND (A) | random numbers, $0<A<1$ |
| L=isNaN (A) | test for NaN, if (l) then... |

Table B. 1 - Defined Operators and Generic Functions for Dense Arrays

| Defined Operation | Matrix Operation |
| :--- | :--- |
| Data Management | Define entries of sparse matrices |
| A.x. B | $A B$ |


| Defined Operation | Matrix Operation |
| :---: | :---: |
| .t. A, .h. A | $A^{T}, A^{*}$ |
| A .ix. B | $A^{-1} B$ |
| B . xi. A | $B A^{-1}$ |
| $\begin{array}{\|lll} \hline \text { A .tx. B, or (.t. A) } & \text { x. } & \text { ( } \\ \text { A .h. } & \text { B, or (.h. A) } & \text { x. } \end{array}$ | $A^{T} B, A^{*} B$ |
| $\begin{array}{\|l\|l\|} \hline \text { B .xt. A, or B .x. (.t. A) } \\ \text { B .xh. A, or B .x. } & \text { (.h. A) } \\ \hline \end{array}$ | $B A^{T}, B A^{*}$ |
| A+B | Sum of two sparse matrices |
| $\mathrm{C}=\mathrm{COND}(\mathrm{A})$ | $\left\\|A^{-1}\right\\| .\\|A\\|$ |

Table B. 2 - Defined Operators and Generic Functions for Harwell-Boeing Sparse Matrices

## Using ScaLAPACK, LAPACK, LINPACK, and EISPACK

Many of the codes in the IMSL Library are based on LINPACK, Dongarra et al. (1979), and EISPACK, Smith et al. (1976), collections of subroutines designed in the 1970s and early 1980s. LAPACK, Anderson et al. (1999), was designed to make the linear solvers and eigensystem routines run more efficiently on high performance computers. For a number of IMSL routines, the user of the IMSL Fortran Numerical Library has the option of linking to code which is based on either the legacy routines or the more efficient LAPACK routines.

Table C below lists the IMSL routines that make use of LAPACK codes. The intent is to obtain improved performance for IMSL codes by using LAPACK codes that have good performance by virtue of using BLAS with good performance. To obtain improved performance we recommend linking with High Performance versions of LAPACK and BLAS, if available. The LAPACK, codes are listed where they are used. Details on linking to the appropriate IMSL Library and alternate libraries for LAPACK and BLAS are explained in the online README file of the product distribution.

| Generic Name <br> of <br> IMSL Routine | LAPACK Routines <br> used when Linking with High <br> Performance Libraries |
| :--- | :---: |
| LSARG | ?GERFS, ?GETRF, ?GECON, ?=S /D |
| LSLRG | ?GETRF, ?GETRS, ?=S /D |
| LFCRG | ?GETRF, ?GECON, ? $=$ S/D |
| LFTRG | ?GETRF, ?=S /D |
| LFSRG | ?GETRS, ?=S /D |
| LFIRG | ?GETRS, ?=S /D |


| Generic Name of <br> IMSL Routine | LAPACK Routines used when Linking with High Performance Libraries |
| :---: | :---: |
| LINRG | ?GETRF, ?GETRI ?=S/D |
| LSACG | ?GETRF, GETRS, ?GECON, ?=C/Z |
| LSLCG | ?GETRF, ?GETRS, ?=C/Z |
| LFCCG | ?GETRF, ?GECON, ?=C/Z |
| LFTCG | ?GETRF, ?C/Z |
| LFSCG | ?GETRS, ?C/Z |
| LFICG | ?GERFS, ?GETRS, ?=C/Z |
| LINCG | ?GETRF, ?GETRI, ?=C/Z |
| LSLRT | ?TRTRS, ? $=$ / D |
| LFCRT | ?TRCON, ? $=$ S/D |
| LSLCT | ?TRTRS, ? $=\mathrm{C} / \mathrm{Z}$ |
| LFCCT | ?TRCON, ? $=$ C/Z |
| LSADS | ?PORFS, ?POTRS, ?=S/D |
| LSLDS | ?POTRF, ?POTRS, ?=S/D |
| LFCDS | ?POTRF, ?POCON, ?=S/D |
| LFTDS | ?POTRF, ?-S/D |
| LFSDS | ?POTRS, ?-S/D |
| LFIDS | ?PORFS, ?POTRS, ?=S/D |
| LINDS | ?POTRF, ?=S/D |
| LSASF | ?SYRFS, ?SYTRF, ?SYTRS, ?=S/D |
| LSLSF | ?SYTRF, ?SYTRS, ?=S/D |
| LFCSF | ?SYTRF, ?SYCON, ?=S/D |
| LFTSF | ?SYTRF, ?=S/D |
| LFSSF | ?SYTRF, ?=S/D |
| LFISF | ?SYRFS, ?=S/D |
| LSADH | ?POCON, ?POTRF, ?POTRS, ?=C/Z |
| LSLDH | ?TRTRS, ?POTRF, ?=C/Z |
| LFCDH | ?POTRF, ?POCON, ?=C/Z |
| LFTDH | ? POTRF, ? $=$ C/Z |
| LFSDH | ?TRTRS, ? $=\mathrm{C} / \mathrm{Z}$ |
| LFIDH | ?PORFS, ?POTRS, ?=C/Z |
| LSAHF | ? HECON, ? HERFS, ? ${ }^{\text {HETRF, }}$ ? HETRS, ? $=$ C/Z |
| LSLHF | ?HECON, ?HETRF, ?HETRS, ?=C/Z |


| Generic Name of IMSL Routine | LAPACK Routines used when Linking with High Performance Libraries |
| :---: | :---: |
| LFCHF | ?HETRF, ?HECON, ?=C/Z |
| LFTHF | ?HETRF, ? $=$ C/Z |
| LFSHF | ?HETRS, ? $=$ C/Z |
| LFIHF | ? HERFS , ? HETRS, ?=C/Z |
| LSARB | ?GBTRF, ?GBTRS, ?GBRFS ?=S/D |
| LSLRB | ?GBTRF, ?GBTRS, ?=S/D |
| LFCRB | ?GBTRF, ?GBCON, ?=S/D |
| LFTRB | ?GBTRF, ? $=$ S/D |
| LFSRB | ?GBTRS, ? $=$ S/D |
| LFIRB | ?GBTRS, ?GBRFS, ?=S/D |
| LSQRR | ?GEQP3, ?GEQRF, ?ORMQR, ?TRTRS. ?=S/D |
| LQRRV | ?GEQP3, ?GEQRF, ?ORMQR, ?=S/D |
| LSBRR | ?GEQRF, ? $=$ S/D |
| LQRRR | ?GEQRE, ?=S/D |
| LSVRR | ?GESVD, ?-S/D |
| LSVCR | ?GESVD, ?=C/Z |
| LSGRR | ?GESVD, ?=S/D |
| LQRSL | ?TRTRS, ? ${ }^{\text {armQR, }}$ ?=S/D |
| LQERR | ?ORGQR, ?=S/D |
| EVLRG | ?GEBAL, ?GEHRD, ?HSEQR ?=S/D |
| EVCRG | ?GEEVX, ?=S/D |
| EVLCG | ?HSEQR, ?GEBAL, ?GEHRD, ?=C/Z |
| EVCCG | ? GEEV, ? $=$ C/Z |
| EVLSF | ?SYEV, ?=S/D |
| EVCSF | ?SYEV, ?=S/D |
| EVLHF | ? HEEV , ? $=$ C/Z |
| EVCHF | ? HEEV, ?=C/Z |
| GVLRG |  |
| GVCRG | ?GEQRF, ?ORMQR, ?GGHRD, ?HGEQZ, ?TGEVC, ?=S/D |
| GVLCG | ?GEQRF, ?UMMQR, ?GGHRD, ?HGEQZ, ?=C/Z |
| GVCCG | ?GEQRF, ?UMMQR, ?GGHRD, <br> ? HGEQZ, ?TGEVC, ?=C/Z |
| GCLSP | ?SYGV, ?=S/D |


| Generic Name <br> of <br> IMSL Routine | LAPACK Routines <br> used when Linking with High <br> Performance Libraries |
| :--- | :---: |
| GCCSP | ?SYGV, ?=S/D |

Table C - IMSL Routines and LAPACK Routines Utilized Within
ScaLAPACK, Blackford et al. (1997), includes a subset of LAPACK codes redesigned for use on distributed memory MIMD parallel computers. A number of IMSL Library routines make use of a subset of the ScaLAPACK library.

Table D below lists the IMSL routines that make use of ScaLAPACK codes. The intent is to provide access to the ScaLAPACK codes through the familiar IMSL routine interface. The IMSL routines that utilize ScaLAPACK codes have a ScaLAPACK Interface documented in addition to the FORTRAN 90 Interface. Like the LAPACK codes, access to the ScaLAPACK codes is made by linking to the appropriate library. Details on linking to the appropriate IMSL Library and alternate libraries for ScaLAPACK and BLAS are explained in the online README file of the product distribution.

| Generic Name of <br> IMSL Routine | ScaLAPACK Routines used when Linking with High Performance Libraries |
| :---: | :---: |
| LSARG | P?GERFS, P? GETRF, P?GETRS, ?=S/D |
| LSLRG | P?GETRF, P?GETRS, ?=S/D |
| LFCRG | P?GETRF, P?GECON, ?=S/D |
| LFTRG | P?GETRF, ?=S/D |
| LFSRG | P?GETRS, ?=S/D |
| LFIRG | P?GETRS, P?GERFS, ?=S/D |
| LINRG | P?GETRE, P?GETRI ?=S/D |
| LSACG | P?GETRF, P?GETRS, P?GERFS, ?=C/Z |
| LSLCG | P?GETRF, P?GETRS, ?=C/Z |
| LFCCG | P?GETRF, P?GECON, ?=C/Z |
| LFTCG | P?GETRF, ?C/Z |
| LFSCG | P?GETRS, ?C/Z |
| LFICG | P?GERFS, P?GETRS, ?=C/Z |
| LINCG | P?GETRF, P?GETRI, ?=C/Z |
| LSLRT | P?TRTRS, ? $=$ / D |
| LFCRT | P?TRCON, ?=S/D |
| LSLCT | P?TRTRS, ?=C/Z |
| LFCCT | P?TRCON, ?=C/Z |
| LSADS | P?PORFS, P?POTRF, P?POTRS, ?=S/D |


| Generic Name <br> of <br> IMSL Routine | ScaLAPACK Routines <br> used when Linking with High Performance <br> Libraries |
| :--- | :---: |
| LSLDS | P?POTRF, P?POTRS, ? |

Table D-IMSL Routines and ScaLAPACK Routines Utilized Within

## Using ScaLAPACK Enhanced Routines

## General Remarks

Use of the ScaLAPACK enhanced routines allows a user to solve large linear systems of algebraic equations at a performance level that might not be achievable on one computer by performing the work in parallel across multiple computers. One might also use these routines on linear systems that prove to be too large for the address space of the target computer. Visual Numerics has tried to facilitate the use of parallel computing in these situations by providing interfaces to ScaLAPACK routines which accomplish the task. The IMSL Library solver interface has the same look and feel whether one is using the routine on a single computer or across multiple computers.

The basic steps required to utilize the IMSL routines which interface with ScaLAPACK routines are:

1. Initialize MPI
2. Initialize the processor grid
3. Define any necessary array descriptors
4. Allocate space for the local arrays
5. Set up local matrices across the processor grid
6. Call the IMSL routine which interfaces with ScaLAPACK
7. Gather the results from across the processor grid
8. Release the processor grid
9. Exit MPI

Utilities are provided in the IMSL Library that facilitate these steps for the user. Each of these utilities is documented in Chapter 11, "Utilities". We visit the steps briefly here:

## 1. Initialize MPI

The user should call MP_SETUP () at this step. This function is described in detail in "Getting Started with Modules MPI_setup int and MPI_node_int" in Chapter 10, Linear Algebra Operators and Generic Functions of this manual. For ScaLAPACK usage, suffice it to say that following a call to the function MP_SETUP (), the module MPI_node_int will contain information about the number of processors, the rank of a processor, and the communicator for the application. A call to this function will return the number of processors available to the program. Since the module MPI_node_int is used by MPI_setup_int, it is not necessary to explicitly use the module MPI_node_int. If MP_SETUP () is not called, then the program will compute entirely on one node. No routine from MPI will be called.

## 2. Initialize the processor grid

SCALAPACK_SETUP (see Chapter 11, "Utilities") is called at this step. This call will set up the processor grid for the user, define the context ID variable, MP_ICTXT, for the processor grid, and place MP_ICTXT into the module GRIDINFO_INT. Use of SCALAPACK_SUPPORT will make the information in MPI_NODE_INT and GRIDINFO_INT available to the user's program.

## 3. Define any necessary array descriptors

Consider the generic matrix A which is to be carved up and distributed across the processors in the processor grid. In ScaLAPACK parlance, we refer to A as being the "global" array A which is to be distributed across the processor grid in 2D block cyclic fashion (Chapter 11, "Utilities"). Each processor in the grid will then have access to a subset of the global array A. We refer to the subset array to which the individual processor has access as the "local" array A0. Just as it is sometimes necessary for a program to be aware of the leading dimension of the global array A, it is also necessary for the program to be aware of other critical information about the local array A0. This information can be obtained by calling the IMSL utility SCALAPACK_GETDIM
(Chapter 11, "Utilities"). ScaLAPACK Library utility DESCINIT (See the Usage Notes section of Chapter 11, "Utilities") is then used to store this information in a vector.

## 4. Allocate space for the local arrays

The array dimensions, obtained in the previous step, are used at this point to allocate space for any local arrays that will be used in the call to the IMSL routine.

## 5. Set up local matrices across the processor grid

If the matrices to be used by the solvers have not been distributed across the processor grid, IMSL provides utility routines SCALAPACK_READ and SCALAPACK_MAP to help in the distribution of global arrays across processors. SCALAPACK_READ will read data from a file while SCALAPACK_MAP will map a global array to the processor grid. Users may choose to distribute the arrays themselves as long as they distribute the arrays in 2D block cyclic fashion consistent with the array descriptors that have been defined.

## 6. Call the IMSL routine which interfaces with ScaLAPACK

The IMSL routines which interface with ScaLAPACK are listed in Table D.

## 7. Gather the results from across the processor grid

IMSL provides utility routines SCALAPACK_WRITE and SCALAPACK_UNMAP to help in the gathering of results from across processors to a global array or file. SCALAPACK_WRITE will write data to a file while SCALAPACK_UNMAP will map local arrays from the processor grid to a global array.

## 8. Release the processor grid

This is accomplished by a call to SCALAPACK_EXIT.

## 9. Exit MPI

A call to MP_SETUP with the argument 'FINAL' will shut down MPI and set the value of MP_NPROCS $=0$. This flags that MPI has been initialized and terminated. It cannot be initialized again in the same program unit execution. No MPI routine is defined when MP_NPROCS has this value.

## Chapter 1: Linear Systems

## Routines

1.1. Linear Solvers
1.1.1 Solves a general system of linear equationsAx = b..................................................................LIN_SOL_GEN9
1.1.2 Solves a system of linear equations $A x=b$, where $A$ is a self-adjoint matrix LIN_SOL_SELF ..... 17
1.1.3 Solves a rectangular system of linear equations $A x \cong b$, in a least-squares sense LIN_SOL_LSQ ..... 27
1.1.4 Solves a rectangular least-squares system of linear equations $A x \cong b$ using singular value decomposition

$\qquad$ LIN_SOL_SVD ..... 36
1.1.5 Solves multiple systems of linear equations LIN_SOL_TRI ..... 44
1.1.6 Computes the singular value decomposition (SVD) of a rectangular matrix, $A$. .LIN_SVD ..... 57
1.2. Large-Scale Parallel Solvers
1.2.1 Parallel Constrained Least-Squares Solvers ..... 66
1.2.2 Solves a linear, non-negative constrained least-squares system PARALLEL_NONNEGATIVE_LSQ ..... 66
1.2.3 Solves a linear least-squares system with bounds on the unknowns PARALLEL_BOUNDED_LSQ ..... 74
1.3. Solution of Linear Systems, Matrix Inversion, and $Q$ Determinant Evaluation
1.3.1 Real General MatricesHigh accuracy linear system solution ................................LSARG82
Solves a linear system ..... LSLRG ..... 87
Factors and computes condition number LFCRG ..... 93
Factors ..... 98
Solves after factoring ..... 103
LFSRG
High accuracy linear system solution after factoring ..... 107
Computes determinant after factoring LFDRG ..... 113
Inverts LINRG ..... 114
1.3.2 Complex General Matrices High accuracy linear system solution LSACG ..... 118
Solves a linear system LSLCG ..... 123
Factors and computes condition number ..... LFCCG ..... 127
Factors LFTCG ..... 133
Solves a linear system after factoring LFSCG ..... 138
High accuracy linear system solution after factoring ..... LFICG ..... 142
Computes determinant after factoring LFDCG ..... 148
Inverts LINCG ..... 149
1.3.3 Real Triangular Matrices
Solves a linear system LSLRT ..... 154
Computes condition number LFCRT ..... 157
Computes determinant after factoring LFDRT ..... 161
Inverts ..... 163
1.3.4 Complex Triangular Matrices
Solves a linear system LSLCT ..... 164
Computes condition number ..... LFCCT ..... 168
Computes determinant after factoring ..... LFDCT ..... 172
Inverts LINCT ..... 174
1.3.5 Real Positive Definite Matrices High accuracy linear system solution LSADS ..... 176
Solves a linear system ..... 180
Factors and computes condition number ..... 185
Factors. ..... 190
Solve a linear system after factoring ..... 194
High accuracy linear system solution after factoring ..... 199
Computes determinant after factoring ..... 204
Inverts ..... 205
1.3.6 Real Symmetric Matrices
High accuracy linear system solution ..... LSASF ..... 209
Solves a linear system ..... LSLSF ..... 212
Factors and computes condition number ..... LFCSF ..... 214
Factors. ..... LFTSF ..... 217
Solves a linear system after factoring ..... LFSSF ..... 219
High accuracy linear system solution after factoring ..... LFISF ..... 221
Computes determinant after factoring LFDSF ..... 224
1.3.7 Complex Hermitian Positive Definite Matrices High accuracy linear system solution ..... LSADH ..... 226
Solves a linear system ..... 231
Factors and computes condition number LFCDH ..... 236
Factors LFTDH ..... 241
Solves a linear system after factoring ..... LFSDH ..... 246
High accuracy linear system solution after factoring LFIDH ..... 251
Computes determinant after factoring LFDDH ..... 256
1.3.8 Complex Hermitian Matrices
High accuracy linear system solution .................................. LSAHF ..... 258
Solves a linear system ..... LSLHF ..... 261
Factors and computes condition number LFCHF ..... 263
Factors ..... LFTHF ..... 266
Solves a linear system after factoring ..... LFSHF ..... 269
High accuracy linear system solution after factoring ..... LFIHF ..... 271
Computes determinant after factoring LFDHF ..... 274
1.3.9 Real Band Matrices in Band Storage Solves a tridiagonal system LSLTR ..... 275
Solves a tridiagonal system: Cyclic Reduction LSLCR ..... 277
High accuracy linear system solution ..... LSARB ..... 280
Solves a linear system ..... LSLRB ..... 282
Factors and compute condition number ..... LFCRB ..... 287
Factors ..... LFTRB ..... 290
Solves a linear system after factoring LFSRB ..... 293
High accuracy linear system solution after factoring LFIRB ..... 295
Computes determinant after factoring LFDRB ..... 298
1.3.10 Real Band Symmetric Positive Definite Matrices in Band Storage High accuracy linear system solution LSAQS ..... 300
Solves a linear system LSLQS ..... 303
Solves a linear system ..... 305
Factors and computes condition number LFCQS ..... 308
Factors LFTQS ..... 311
Solves a linear system after factoring LFSQS ..... 313
High accuracy linear system solution after factoring ..... LFIQS ..... 315
Computes determinant after factoring LFDQS ..... 318
1.3.11 Complex Band Matrices in Band Storage
Solves a tridiagonal system ..... LSLTQ ..... 319
Solves a tridiagonal system: Cyclic Reduction LSLCQ ..... 321
High accuracy linear system solution LSACB ..... 324
Solves a linear system ..... LSLCB ..... 327
Factors and computes condition number LFCCB ..... 330
Factors LFTCB ..... 333
Solves a linear system after factoring LFSCB ..... 336
High accuracy linear system solution after factoring ..... LFICB ..... 338
Computes determinant after factoring LFDCB ..... 342
1.3.12 Complex Band Positive Definite Matrices in Band Storage High accuracy linear system solution LSAQH ..... 344
Solves a linear system LSLQH ..... 346
Solves a linear system LSLQB ..... 349
Factors and compute condition number LFCQH ..... 352
Factors LFTQH ..... 355
Solves a linear system after factoring LFSQH ..... 358
High accuracy linear system solution after factoring LFIQH ..... 360
Computes determinant after factoring LFDQH ..... 362
1.3.13 Real Sparse Linear Equation Solvers Solves a sparse linear system LSLXG ..... 364
Factors. LFTXG ..... 369
Solves a linear system after factoring LFSXG ..... 374
1.3.14 Complex Sparse Linear Equation Solvers Solves a sparse linear system LSLZG ..... 377
Factors. LFTZG ..... 382
Solves a linear system after factoring LFSZG ..... 387
1.3.15 Real Sparse Symmetric Positive Definite Linear Equation Solvers Solves a sparse linear system LSLXD ..... 391
Symbolic Factor LSCXD ..... 395
Computes Factor LNFXD ..... 399
Solves a linear system after factoring LFSXD ..... 404
1.3.16 Complex Sparse Hermitian Positive Definite Linear Equation Solvers Solves a sparse linear system ..... LSLZD ..... 408
Computes Factor LNFZD ..... 412
Solves a linear system after factoring LFSZD ..... 417
1.3.17 Real Toeplitz Matrices in Toeplitz Storage Solves a linear system ..... LSLTO ..... 421
1.3.18 Complex Toeplitz Matrices in Toeplitz Storage Solves a linear system LSLTC ..... 423
1.3.19 Complex Circulant Matrices in Circulant Storage Solves a linear system LSLCC ..... 425
1.3.20 Iterative Methods
Preconditioned conjugate gradient. PCGRC ..... 427
Jacobi conjugate gradient JCGRC ..... 433
Generalized minimum residual GMRES ..... 436
1.4. Linear Least Squares and Matrix Factorization
1.4.1 Least Squares, QR Decomposition and Generalized Inverse Solves a Least-squares system LSQRR ..... 446
Solves a Least-squares system ..... LQRRV ..... 452
High accuracy Least squares ..... LSBRR ..... 458
Linearly constrained Least squares ..... LCLSQ ..... 462
QR decomposition ..... 466
Accumulation of QR decomposition LQERR ..... 473
QR decomposition Utilities ..... LQRSL ..... 478
QR factor update LUPQR ..... 484
1.4.2 Cholesky Factorization Cholesky factoring for rank deficient matrices ..... 489
Cholesky factor update. LUPCH ..... 491
Cholesky factor down-date LDNCH ..... 494
1.4.3 Singular Value Decomposition (SVD) Real singular value decomposition LSVRR ..... 498
Complex singular value decomposition ..... LSVCR ..... 504
Generalized inverse LSGRR ..... 508

## Usage Notes

Section 1.1 describes routines for solving systems of linear algebraic equations by direct matrix factorization methods, for computing only the matrix factorizations, and for computing linear least-squares solutions.
Section 1.2 describes routines for solving systems of parallel constrained least-squares.
Many of the routines described in sections 1.3 and 1.4 are for matrices with special properties or structure. Computer time and storage requirements for solving systems with coefficient matrices of these types can often be drastically reduced, using the appropriate routine, compared with using a routine for solving a general complex system.
The appropriate matrix property and corresponding routine can be located in the "Routines" section. Many of the linear equation solver routines in this chapter are derived from subroutines from LINPACK, Dongarra et al. (1979). Other routines have been developed by Visual Numerics, derived from draft versions of LAPACK subprograms, Bischof et al. (1988), or were obtained from alternate sources.

A system of linear equations is represented by $A x=b$ where $A$ is the $n \times n$ coefficient data matrix, $b$ is the known right-hand-side $n$-vector, and $x$ is the unknown or solution $n$-vector. Figure 1-1 summarizes the relationships among the subroutines. Routine names are in boxes and input/output data are in ovals. The suffix ** in the subroutine names depend on the matrix type. For example, to compute the determinant of $A$ use LFC** or LFT** followed by LFD**.

The paths using LSA** or LFI** use iterative refinement for a more accurate solution. The path using LSA** is the same as using LFC** followed by LFI**. The path using LSL** is the same as the path using LFC** followed by LFS**. The matrix inversion routines LIN** are available only for certain matrix types.

## Matrix Types

The two letter codes for the form of coefficient matrix, indicated by ** in Figure 1-1, are as follows:

| RG | Real general (square) matrix. |
| :--- | :--- |
| CG | Complex general (square) matrix. |
| TR or $C R$ | Real tridiagonal matrix. |
| RB | Real band matrix. |
| $T Q$ or $C Q$ | Complex tridiagonal matrix. |
| CB | Complex band matrix. |
| SF | Real symmetric matrix stored in the upper half of a square matrix. |
| DS | Real symmetric positive definite matrix stored in the upper half of a square matrix. |
| DH | Complex Hermitian positive definite matrix stored in the upper half of a complex <br> square matrix. |
| HF | Complex Hermitian matrix stored in the upper half of a complex square matrix. |


| QS or $P B$ | Real symmetric positive definite band matrix. |
| :--- | :--- |
| QH or $Q B$ | Complex Hermitian positive definite band matrix. |
| $X G$ | Real general sparse matrix. |
| $Z G$ | Complex general sparse matrix. |
| $X D$ | Real symmetric positive definite sparse matrix. |
| $Z D$ | Complex Hermitian positive definite sparse matrix. |



Figure 1-1 Solution and Factorization of Linear Systems

## Solution of Linear Systems

The simplest routines to use for solving linear equations are LSL** and LSA**. For example, the mnemonic for matrices of real general form is RG. So, the routines LSARG and LSLRG are appropriate to use for solving linear systems when the coefficient matrix is of real general form. The routine LSARG uses iterative refinement, and more time than LSLRG, to determine a high accuracy solution.
The high accuracy solvers provide maximum protection against extraneous computational errors. They do not protect the results from instability in the mathematical approximation. For a more complete discussion of this and other important topics about solving linear equations, see Rice (1983), Stewart (1973), or Golub and van Loan (1989).

## Multiple Right Sides

There are situations where the LSL** and LSA** routines are not appropriate. For example, if the linear system has more than one right-hand-side vector, it is most economical to solve the system by first calling a factoring routine and then calling a solver routine that uses the factors. After the coefficient matrix has been factored, the routine LFS** or LFI** can be used to solve for one right-hand side at a time. Routines LFI** uses iterative refinement to determine a high accuracy solution but requires more computer time and storage than routines LFS**.

## Determinants

The routines for evaluating determinants are named LFD**. As indicated in Figure 1-1, these routines require the factors of the matrix as input. The values of determinants are often badly scaled. Additional complications in structures for evaluating them result from this fact. See Rice (1983) for comments on determinant evaluation.

## Iterative Refinement

Iterative refinement can often improve the accuracy of a well-posed numerical solution. The iterative refinement algorithm used is as follows:
$x_{0}=A^{-1} b$
For $i=1,50$
$r_{i}=A x_{i-1}-b$ computed in higher precision
$p_{i}=A^{-1} r_{i}$
$x_{i}=x_{i_{-1}}-p_{i}$
if $\left(\left\|p_{i}\right\|_{\infty} \leq \varepsilon\left\|x_{i}\right\|_{\infty}\right)$ Exit
End for
Error - Matrix is too ill-conditioned
If the matrix $A$ is in single precision, then the residual $r_{i}=A x_{i_{-1}}-b$ is computed in double precision. If $A$ is in double precision, then quadruple-precision arithmetic routines are used.

The use of the value 50 is arbitrary. In fact a single correction is usually sufficient. It is also helpful even when $r_{i}$ is computed in the same precision as the data.

## Matrix Inversion

An inverse of the coefficient matrix can be computed directly by one of the routines named LIN**. These routines are provided for general matrix forms and some special matrix forms. When they do not exist, or when it is desirable to compute a high accuracy inverse, the two-step technique of calling the factoring routine followed by the solver routine can be used. The inverse is the solution of the matrix system $A X=I$ where $I$ denotes the $n \times n$ identity matrix, and the solution is $X=A^{-1}$.

## Singularity

The numerical and mathematical notions of singularity are not the same. A matrix is considered numerically singular if it is sufficiently close to a mathematically singular matrix. If error
messages are issued regarding an exact singularity then specific error message level reset actions must be taken to handle the error condition. By default, the routines in this chapter stop. The solvers require that the coefficient matrix be numerically nonsingular. There are some tests to determine if this condition is met. When the matrix is factored, using routines LFC**, the condition number is computed. If the condition number is large compared to the working precision, a warning message is issued and the computations are continued. In this case, the user needs to verify the usability of the output. If the matrix is determined to be mathematically singular, or ill-conditioned, a least-squares routine or the singular value decomposition routine may be used for further analysis.

## Special Linear Systems

Toeplitz matrices have entries which are constant along each diagonal, for example:

$$
A=\left[\begin{array}{llll}
p_{0} & p_{1} & p_{2} & p_{3} \\
p_{-1} & p_{0} & p_{1} & p_{2} \\
p_{-2} & p_{-1} & p_{0} & p_{1} \\
p_{-3} & p_{-2} & p_{-1} & p_{0}
\end{array}\right]
$$

Real Toeplitz systems can be solved using LSLTO. Complex Toeplitz systems can be solved using LSLTC.

Circulant matrices have the property that each row is obtained by shifting the row above it one place to the right. Entries that are shifted off at the right reenter at the left. For example:

$$
A=\left[\begin{array}{llll}
p_{1} & p_{2} & p_{3} & p_{4} \\
p_{4} & p_{1} & p_{2} & p_{3} \\
p_{3} & p_{4} & p_{1} & p_{2} \\
p_{2} & p_{3} & p_{4} & p_{1}
\end{array}\right]
$$

Complex circulant systems can be solved using LSLCC.

## Iterative Solution of Linear Systems

The preconditioned conjugate gradient routines PCGRC and JCGRC can be used to solve symmetric positive definite systems. The routines are particularly useful if the system is large and sparse.
These routines use reverse communication, so $A$ can be in any storage scheme. For general linear systems, use GMRES.

## QR Decomposition

The $Q R$ decomposition of a matrix $A$ consists of finding an orthogonal matrix $Q$, a permutation matrix $P$, and an upper trapezoidal matrix $R$ with diagonal elements of nonincreasing magnitude, such that $A P=Q R$. This decomposition is determined by the routines $L Q R R R$ or $L Q R R V$. It returns $R$ and the information needed to compute $Q$. To actually compute $Q$ use LQERR. Figure 1-2 summarizes the relationships among the subroutines.

The $Q R$ decomposition can be used to solve the linear system $A x=b$. This is equivalent to $R x=Q^{T} P b$. The routine LQRSL, can be used to find $Q^{T} P b$ from the information computed by

LQRRR. Then $x$ can be computed by solving a triangular system using LSLRT. If the system $A x=b$ is overdetermined, then this procedure solves the least-squares problem, i.e., it finds an $x$ for which

$$
\|A x-b\|_{2}^{2}
$$

is a minimum.
If the matrix $A$ is changed by a rank-1 update, $A \rightarrow A+\alpha x y^{T}$, the QR decomposition of $A$ can be updated/down-dated using the routine $\angle U P Q R$. In some applications a series of linear systems which differ by rank-1 updates must be solved. Computing the QR decomposition once and then updating or down-dating it usually faster than newly solving each system.


Figure 1-2 Least-Squares Routine

## LIN_SOL_GEN

Solves a general system of linear equations $A x=b$. Using optional arguments, any of several related computations can be performed. These extra tasks include computing the $L U$ factorization of $A$ using partial pivoting, representing the determinant of $A$, computing the inverse matrix $A^{-1}$, and solving $A^{T} x=b$ or $A x=b$ given the $L U$ factorization of $A$.

## Required Arguments

A- Array of size $n \times n$ containing the matrix. (Input [/Output])
If the packaged option lin_sol_gen_save_LU is used then the $L U$ factorization of A is saved in A. For solving efficiency, the diagonal reciprocals of the matrix $U$ are saved in the diagonal entries of $A$.
$\boldsymbol{B}$ - Array of size $n \times n b$ containing the right-hand side matrix. (Input [/Output])
If the packaged option lin_sol_gen_save_LU is used then input $B$ is used as work storage and is not saved.
$\boldsymbol{X}$ - Array of size $n \times n b$ containing the solution matrix.(Output)

## Optional Arguments

NROWS $=\mathrm{n}$ (Input)
Uses array $A(1: n, 1: n)$ for the input matrix.
Default: $\mathrm{n}=\operatorname{size}(\mathrm{A}, 1)$
NRHS = nb (Input)
Uses array $\mathrm{b}(1: \mathrm{n}, 1: \mathrm{nb})$ for the input right-hand side matrix.
Default: $\mathrm{nb}=\operatorname{size}(\mathrm{b}, 2)$
Note that b must be a rank-2 array.
pivots = pivots(:) (Output [/Input])
Integer array of size $n$ that contains the individual row interchanges. To construct the permuted order so that no pivoting is required, define an integer array ip(n). Initialize $\mathrm{ip}(i)=i, i=1, \mathrm{n}$ and then execute the loop, after calling lin_sol_gen,
k=pivots(i)
interchange ip(i) and ip(k), $i=1, n$
The matrix defined by the array assignment that permutes the rows, $A(1: n, 1: n)=A(i p(1: n), 1: n)$, requires no pivoting for maintaining numerical stability. Now, the optional argument "iopt=" and the packaged option number ?_lin_sol_gen_no_pivoting can be safely used for increased efficiency during the $L \bar{U}$ factorization of $A$.
$\operatorname{det}=\operatorname{det}(1: 2)$ (Output)
Array of size 2 of the same type and kind as A for representing the determinant of the input matrix. The determinant is represented by two numbers. The first is the base with the sign or complex angle of the result. The second is the exponent. When $\operatorname{det}(2)$ is within exponent range, the value of this expression is given by $\operatorname{abs}(\operatorname{det}(1)) * * \operatorname{det}(2) *(\operatorname{det}(1)) / \mathrm{abs}(\operatorname{det}(1))$. If the matrix is not singular, $\operatorname{abs}(\operatorname{det}(1))=\operatorname{radix}(\operatorname{det})$; otherwise, $\operatorname{det}(1)=0$., and $\operatorname{det}(2)=-\operatorname{huge}(\operatorname{abs}(\operatorname{det}(1)))$.
ainv = ainv(:,:) (Output)
Array of the same type and kind as $\mathrm{A}(1: \mathrm{n}, 1: \mathrm{n})$. It contains the inverse matrix, $A^{-1}$, when the input matrix is nonsingular.
iopt $=$ iopt (: ) (Input)
Derived type array with the same precision as the input matrix; used for passing optional data to the routine. The options are as follows:

| Packaged Options for lin_sol_gen |  |  |
| :---: | :---: | :---: |
| Option Prefix = ? | Option Name | Option Value |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_sol_gen_set_small | 1 |
| $S_{-}, d_{-}, \mathrm{C}_{-}, z_{-}$ | lin_sol_gen_save_LU | 2 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_sol_gen_solve_A | 3 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_sol_gen_solve_ADJ | 4 |
| $S_{-}, d_{-}, \mathrm{c}_{-}, z_{-}$ | lin_sol_gen_no_pivoting | 5 |
| $S_{-}, d_{-}, \mathrm{c}_{-}, z_{-}$ | lin_sol_gen_scan_for_NaN | 6 |
| $S_{-}, d_{-}, \mathrm{c}_{-}, z_{-}$ | lin_sol_gen_no_sing_mess | 7 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_sol_gen_A_is_sparse | 8 |

iopt(IO) $=$ ?_options(?_lin_sol_gen_set_small, Small)
Replaces a diagonal term of the matrix $U$ if it is smaller in magnitude than the value Small using the same sign or complex direction as the diagonal. The system is declared singular. A solution is approximated based on this replacement if no overflow results. Default: the smallest number that can be reciprocated safely

```
iopt(IO) = ?_options(?_lin_sol_gen_save_LU, ?_dummy)
```

Saves the $L U$ factorization of A. Requires the optional argument "pivots=" if the routine will be used later for solving systems with the same matrix. This is the only case where the input arrays A and b are not saved. For solving efficiency, the diagonal reciprocals of the matrix $U$ are saved in the diagonal entries of A .
iopt(IO) = ?_options(?_lin_sol_gen_solve_A, ?_dummy)
Uses the $\bar{L} U$ factorization of $\bar{A}$ computed and saved to solve $A x=b$.
iopt(IO) $=$ ?_options(?_lin_sol_gen_solve_ADJ,?_dummy)
Uses the $L U$ factorization of $A$ computed and saved to solve $A^{T} x=b$.
iopt(IO) = ?_options(?_lin_sol_gen_no_pivoting, ?_dummy)
Does no row pivoting. The $\overline{\operatorname{arrray}} \mathrm{pivots}(:)$, if present, are output as pivots $(i)=i$, for $i=1, \ldots, \mathrm{n}$.
iopt(IO) = ?_options(?_lin_sol_gen_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(a(i,j)) .or. isNan(b(i,j)) ==.true.
See the isNaN() function, Chapter 10.
Default: Does not scan for NaNs.
iopt(IO) = ?_options(?_lin_sol_gen_no_sing_mess,?_dummy)
Do not point an error message when the matrix $A$ is singular.

```
iopt(IO) = ?_options(?_lin_sol_gen_A_is_sparse,?_dummy)
```

Uses an indirect updating loop for the LU factorization that is efficient for sparse matrices where all matrix entries are stored.

## FORTRAN 90 Interface

Generic: CALL LIN_SOL_GEN (A, B, X [, ...])
Specific: The specific interface names are S_LIN_SOL_GEN, D_LIN_SOL_GEN, C_LIN_SOL_GEN, and Z_LIN_SOL_GEN.

## Description

Routine LIN_SOL_GEN solves a system of linear algebraic equations with a nonsingular coefficient matrix $A$. It first computes the $L U$ factorization of $A$ with partial pivoting such that $L U=A$. The matrix $U$ is upper triangular, while the following is true:

$$
L^{-1} A \equiv L_{n} P_{n} L_{n-1} P_{n-1} \cdots L_{1} P_{1} A \equiv U
$$

The factors $P_{i}$ and $L_{i}$ are defined by the partial pivoting. Each $P_{i}$ is an interchange of row $i$ with row $j \geq i$. Thus, $P_{i}$ is defined by that value of $j$. Every

$$
L_{i}=I+m_{i} e_{i}^{T}
$$

is an elementary elimination matrix. The vector $m_{i}$ is zero in entries $1, \ldots, i$. This vector is stored as column $i$ in the strictly lower-triangular part of the working array containing the decomposition information. The reciprocals of the diagonals of the matrix $U$ are saved in the diagonal of the working array. The solution of the linear system $A x=b$ is found by solving two simpler systems,

$$
y=L^{-1} b \text { and } x=U^{-1} y
$$

More mathematical details are found in Golub and Van Loan (1989, Chapter 3).

## Fatal and Terminal Error Messages

See the messages.gls file for error messages for LIN_SOL_GEN. The messages are numbered 161-175; 181-195; 201-215; 221-235.

## Example 1: Solving a Linear System of Equations

This example solves a linear system of equations. This is the simplest use of lin_sol_gen. The equations are generated using a matrix of random numbers, and a solution is obtained corresponding to a random right-hand side matrix. Also, see operator_ex01, supplied with the product examples, for this example using the operator notation.

```
use lin_sol_gen_int
use ran\overline{d}ge\overline{n}in\overline{t}
use error_option_packet
implicit none
```

```
! This is Example 1 for LIN_SOL_GEN.
    integer, parameter :: n=32
    real(kind(1e0)), parameter :: one=1e0
    real(kind(1e0)) err
    real(kind(le0)) A(n,n), b(n,n), x(n,n), res(n,n), y(n**2)
! Generate a random matrix.
    call rand_gen(y)
    A = reshape(y,(/n,n/))
! Generate random right-hand sides.
    call rand_gen(y)
    b = reshape (y, (/n,n/))
! Compute the solution matrix of Ax=b.
    call lin_sol_gen(A, b, x)
! Check the results for small residuals.
    res = b - matmul (A,x)
    err = maxval(abs(res))/sum(abs(A)+abs(b))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_SOL_GEN is correct.'
    end if
    end
```


## Output

Example 1 for LIN_SOL_GEN is correct.

## Additional Examples

## Example 2: Matrix Inversion and Determinant

This example computes the inverse and determinant of $A$, a random matrix. Tests are made on the conditions

$$
A A^{-1}=I
$$

and

$$
\operatorname{det}\left(A^{-1}\right)=\operatorname{det}(A)^{-1}
$$

Also, see operator_ex02.
use lin_sol_gen_int
use ran $\bar{d} \_g e \bar{n} \_i n \bar{t}$
implicit none
! This is Example 2 for LIN_SOL_GEN.
integer i

```
    integer, parameter :: n=32
    real(kind(1e0)), parameter :: one=1.0e0, zero=0.0e0
    real(kind(1e0)) err
    real(kind(1e0)) A(n,n), b(n,0), inv(n,n), x(n,0), res(n,n), &
        y(n**2), determinant(2), inv_determinant(2)
! Generate a random matrix.
    call rand_gen(y)
    A = reshape(y,(/n,n/))
! Compute the matrix inverse and its determinant.
    call lin_sol_gen(A, b, x, nrhs=0, &
    ain}v=inv, det=determinant
! Compute the determinant for the inverse matrix.
    call lin_sol_gen(inv, b, x, nrhs=0, &
            de\overline{t}=inv_determinant)
! Check residuals, A times inverse = Identity.
    res = matmul(A,inv)
    do i=1, n
        res(i,i) = res(i,i) - one
    end do
    err = sum(abs(res)) / sum(abs(a))
    if (err <= sqrt(epsilon(one))) then
        if (determinant(1) == inv_determinant(1) .and. &
            (abs(determinant(2)+inv_determinant(2)) &
            <= abs(determinant(2))*s
            write (*,*) 'Example 2 for LIN_SOL_GEN is correct.'
        end if
    end if
    end
```


## Output

Example 2 for LIN_SOL_GEN is correct.

## Example 3: Solving a System with Iterative Refinement

This example computes a factorization of a random matrix using single-precision arithmetic. The double-precision solution is corrected using iterative refinement. The corrections are added to the developing solution until they are no longer decreasing in size. The initialization of the derived type array iopti $(1: 2)=$ s_option $(0,0.0 e 0)$ leaves the integer part of the second element of iopti (:) at the value zero. This stops the internal processing of options inside lin_sol_gen. It results in the $L U$ factorization being saved after exit. The next time the routine is entered the integer entry of the second element of iopt (: ) results in a solve step only. Since the $L U$ factorization is saved in arrays $A(:,:)$ andipivots (:), at the final step, solve only steps can occur in subsequent entries to lin_sol_gen. Also, see operator_ex03, Chapter 10.

```
    use lin_sol_gen_int
    use rand_gen_int
    implicit none
! This is Example 3 for LIN_SOL_GEN.
    integer, parameter :: n=32
    real(kind(1e0)), parameter :: one=1.0e0, zero=0.0e0
    real(kind(1d0)), parameter : : d_zero=0.0d0
    integer ipivots(n)
    real(kind(le0)) a(n,n), b(n,1), x(n,1), w(n**2)
    real(kind(1e0)) change_new, change_old
    real(kind(1d0)) c(n,1), d(n,n), y(n,1)
    type(s_options) :: iopti(2)=s_options(0, zero)
! Generate a random matrix.
    call rand_gen(w)
    a = reshape(w, (/n,n/))
! Generate a random right hand side.
    call rand_gen(b(1:n,1))
! Save double precision copies of the matrix and right hand side.
    d = a
    c}=\textrm{b
! Start solution at zero.
    y = d_zero
    changè_old = huge (one)
! Use packaged option to save the factorization.
    iopti(1) = s_options(s_lin_sol_gen_save_LU,zero)
    iterative_refinement: do
        b = c - matmul (d,y)
        call lin_sol_gen(a, b, x, &
                pivots=ipivots, iopt=iopti)
        y = x + y
        change_new = sum(abs(x))
! Exit when changes are no longer decreasing.
        if (change_new >= change_old) &
        exit iterative_refinement
        change_old = change_new
! Use option to re-enter code with factorization saved; solve only.
```

```
    iopti(2) = s_options(s_lin_sol_gen_solve_A,zero)
end do iterative_refinement
write (*,*) 'Example 3 for LIN_SOL_GEN is correct.'
end
```


## Output

Example 3 for LIN_SOL_GEN is correct.

## Example 4: Evaluating the Matrix Exponential

This example computes the solution of the ordinary differential equation problem

$$
\frac{d y}{d t}=A y
$$

with initial values $y(0)=y_{0}$. For this example, the matrix $A$ is real and constant with respect to $\dot{t}$. The unique solution is given by the matrix exponential:

$$
y(t)=e^{A t} y_{0}
$$

This method of solution uses an eigenvalue-eigenvector decomposition of the matrix

$$
A=X D X^{-1}
$$

to evaluate the solution with the equivalent formula

$$
y(t)=X e^{D t} z_{0}
$$

where

$$
z_{0}=X^{-1} y_{0}
$$

is computed using the complex arithmetic version of lin_sol_gen. The results for $y(t)$ are real quantities, but the evaluation uses intermediate complex-valued calculations. Note that the computation of the complex matrix $X$ and the diagonal matrix $D$ is performed using the IMSL MATH/LIBRARY FORTRAN 77 interface to routine EVCRG. This is an illustration of intermixing interfaces of Fortran 77 and Fortran 90 code. The information is made available to the Fortran 90 compiler by using the FORTRAN 77 interface for EVCRG. Also, see operator_ex 04 , supplied with the product examples, where the Fortran 90 function EIG () has replaced the call to EVCRG.

```
use lin_sol_gen_int
use rand_gen_int
use Nume\overline{rical_Libraries}
implicit none
! This is Example 4 for LIN_SOL_GEN.
integer, parameter :: n=32, k=128
real(kind(1e0)), parameter :: one=1.0e0, t_max=1, delta_t=t_max/(k-1)
real(kind(1e0)) err, A(n,n), atemp(n,n), y\overline{temp(n**2)}
real(kind(le0)) t(k), y(n,k), Y_prime(n,k)
complex(kind(1e0)) EVAL(n), EVEC(n,n)
```

```
    complex(kind(1e0)) x(n,n), z_0(n,1), y_0(n,1), d(n)
    integer i
! Generate a random matrix in an F90 array.
    call rand gen (ytemp)
    atemp = reshape(ytemp,(/n,n/))
! Assign data to an F77 array.
    A = atemp
! Use IMSL Numerical Libraries F77 subroutine for the
! eigenvalue-eigenvector calculation.
    CALL EVCRG(N, A, N, EVAL, EVEC, N)
! Generate a random initial value for the ODE system.
    call rand_gen(ytemp (1:n))
    y_0(1:n,1) = ytemp(1:n)
! Assign the eigenvalue-eigenvector data to F90 arrays.
    d = EVAL; x = EVEC
! Solve complex data system that transforms the initial values, Xz_0=y_0.
    call lin_sol_gen(x, y_0, z_0)
    t = (/(i\stardel\overline{ta_t,i=0, k}-1)/)
! Compute y and y' at the values t(1:k).
    Y = matmul(x, exp (spread (d, 2,k)*spread(t,1,n))* &
        spread(z_0(1:n,1),2,k))
    y_prime = matmul(x, spread(d,2,k)* &
                                    exp (spread(d,2,k)*spread(t,1,n))* &
                                    spread(z_0(1:n,1),2,k))
! Check results. Is }\mp@subsup{Y}{}{\prime}-Ay=0\mathrm{ ?
    err = sum(abs(y_prime-matmul(atemp,y))) / &
            (sum(abs (atemp))*sum(abs(y)))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 4 for LIN_SOL_GEN is correct.'
    end if
    end
```


## Output

Example 4 for LIN_SOL_GEN is correct.

## LIN_SOL SELF

Solves a system of linear equations $A x=b$, where $A$ is a self-adjoint matrix. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of $A$ using symmetric pivoting, representing the determinant of $A$, computing the inverse matrix $A^{-1}$, or computing the solution of $A x=b$ given the
factorization of $A$. An optional argument is provided indicating that $A$ is positive definite so that the Cholesky decomposition can be used.

## Required Arguments

A- Array of size $n \times n$ containing the self-adjoint matrix. (Input [/Output]
If the packaged option lin_sol_self_save_factors is used then the factorization of A is saved in A. For solving efficiency, the diagonal reciprocals of the matrix $R$ are saved in the diagonal entries of A when the Cholesky method is used.
$\boldsymbol{B}$ - Array of size $n \times n b$ containing the right-hand side matrix. (Input [/Output]
If the packaged option lin_sol_self_save_factors is used then input $B$ is used as work storage and is not saved.
$\boldsymbol{X}-\quad$ Array of size $n \times n b$ containing the solution matrix. (Output)

## Optional Arguments

NROWS $=\mathrm{n}$ (Input)
Uses array $A(1: n, 1: n)$ for the input matrix.
Default: $\mathrm{n}=\operatorname{size}(\mathrm{A}, 1)$

NRHS = nb (Input)
Uses the array $b(1: n, 1: n b)$ for the input right-hand side matrix.
Default: nb $=\operatorname{size}(\mathrm{b}, 2)$
Note that b must be a rank-2 array.
pivots = pivots(:) (Output [/Input])
Integer array of size $n+1$ that contains the individual row interchanges in the first $n$ locations. Applied in order, these yield the permutation matrix $P$. Location $n+1$ contains the number of the first diagonal term no larger than Small, which is defined on the next page of this chapter.
$\operatorname{det}=\operatorname{det}(1: 2)$ (Output)
Array of size 2 of the same type and kind as A for representing the determinant of the input matrix. The determinant is represented by two numbers. The first is the base with the sign or complex angle of the result. The second is the exponent. When $\operatorname{det}(2)$ is within exponent range, the value of the determinant is given by the expression $\operatorname{abs}(\operatorname{det}(1)))^{\star *} \operatorname{det}(2) *(\operatorname{det}(1)) / a b s(\operatorname{det}(1))$. If the matrix is not singular, $\operatorname{abs}(\operatorname{det}(1))=\operatorname{radix}(\operatorname{det})$; otherwise, $\operatorname{det}(1)=0$., and $\operatorname{det}(2)=-\operatorname{huge}(\operatorname{abs}(\operatorname{det}(1)))$.
ainv = ainv(:,:) (Output)
Array of the same type and kind as $\mathrm{A}(1: \mathrm{n}, 1: \mathrm{n})$. It contains the inverse matrix, $A^{-1}$ when the input matrix is nonsingular.
iopt = iopt (: ) (Input)
Derived type array with the same precision as the input matrix; used for passing optional data to the routine. The options are as follows:

| Packaged Options for lin_sol_self |  |  |
| :---: | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| s_, d_, c_, z_ | lin_sol_self_set_small | 1 |
| s_, d_, c_, z_ $^{\|c\|}$ | lin_sol_self_save_factors | 2 |
| s_, d_, c_, z_ | lin_sol_self_no_pivoting | 3 |
| s_, d_, c_, z_ | lin_sol_self_use_Cholesky | 4 |
| s_, d_, c_, z_ | lin_sol_self_solve_A | 5 |
| s_, d_, c_, z_ | lin_sol_self_scan_for_NaN | 6 |
| s_, d_, c_, z_ | lin_sol_self_no_sing_mess | 7 |

iopt(IO) = ?_options(?_lin_sol_self_set_small, Small)
When Aasen's method is used, the tridiagonal system $T u=v$ is solved using $L U$ factorization with partial pivoting. If a diagonal term of the matrix $U$ is smaller in magnitude than the value Small, it is replaced by Small. The system is declared singular. When the Cholesky method is used, the upper-triangular matrix $R$, (see "Description"), is obtained. If a diagonal term of the matrix $R$ is smaller in magnitude than the value Small, it is replaced by Small. A solution is approximated based on this replacement in either case.
Default: the smallest number that can be reciprocated safely
iopt(IO) = ?_options(?_lin_sol_self_save_factors, ?_dummy)
Saves the factorization of A. Requires the optional argument "pivots=" if the routine will be used for solving further systems with the same matrix. This is the only case where the input arrays $A$ and b are not saved. For solving efficiency, the diagonal reciprocals of the matrix $R$ are saved in the diagonal entries of $A$ when the Cholesky method is used.
iopt(IO) = ?_options(?_lin_sol_self_no_pivoting, ?_dummy)
Does no row pivoting. The array pivots(:), if present, satisfies pivots $(i)=i+1$ for $i=1, \ldots, \mathrm{n}-1$ when using Aasen's method. When using the Cholesky method, pivots $(i)=i$ for $i=1, \ldots$, n.
iopt(IO) = ?_options(?_lin_sol_self_use_Cholesky, ?_dummy)
The Cholesky decomposition $P A \overline{P^{T}}=R^{\bar{T} R}$ is used instead of the Aasen method.
iopt(IO) = ?_options(?_lin_sol_self_solve_A, ?_dummy)
Uses the factorization of A computed and saved to solve $A x=b$.

```
iopt(IO) = ?_options(?_lin_sol_self_scan_for_NaN, ?_dummy)
```

Examines each input array entry to find the first value such that
isNaN(a(i,j)) .or. isNan(b(i,j)) ==.true.
See the isNan() function, Chapter 10.
Default: Does not scan for NaNs

```
iopt(IO) = ?_options(?_lin_sol_self_no_sing_mess,?_dummy)
```

Do not print an error message when the matrix $A$ is singular.

## FORTRAN 90 Interface

Generic: CALL LIN_SOL_SELF (A, B, X $[, \ldots]$ )
Specific: The specific interface names are S_LIN_SOL_SELF, D_LIN_SOL_SELF, C_LIN_SOL_SELF, and Z_LIN_SOL_SELF.

## Description

Routine LIN_SOL_SELF routine solves a system of linear algebraic equations with a nonsingular coefficient matrix $A$. By default, the routine computes the factorization of $A$ using Aasen's method. This decomposition has the form

$$
P A P^{T}=L T L^{T}
$$

where $P$ is a permutation matrix, $L$ is a unit lower-triangular matrix, and $T$ is a tridiagonal self-adjoint matrix. The solution of the linear system $A x=b$ is found by solving simpler systems,

$$
\begin{gathered}
u=L^{-1} P b \\
T v=u
\end{gathered}
$$

and

$$
x=P^{T} L^{-T} v
$$

More mathematical details for real matrices are found in Golub and Van Loan (1989, Chapter 4).
When the optional Cholesky algorithm is used with a positive definite, self-adjoint matrix, the factorization has the alternate form

$$
P A P^{T}=R^{T} R
$$

where $P$ is a permutation matrix and $R$ is an upper-triangular matrix. The solution of the linear system $A x=b$ is computed by solving the systems

$$
u=R^{-T} P b
$$

and

$$
x=P^{T} R^{-1} u
$$

The permutation is chosen so that the diagonal term is maximized at each step of the decomposition. The individual interchanges are optionally available in the argument "pivots".

## Fatal and Terminal Error Messages

See the messages.gls file for error messages for LIN_SOL_SELF. These error messages are numbered 321-336; 341-356; 361-376; 381-396.

## Example 1: Solving a Linear Least-squares System

This example solves a linear least-squares system $C x \cong d$, where $\mathrm{C}_{m x n}$ is a real matrix with $m \geq n$. The least-squares solution is computed using the self-adjoint matrix

$$
A=C^{T} C
$$

and the right-hand side

$$
b=A^{T} d
$$

The $n \times n$ self-adjoint system $A x=b$ is solved for $x$. This solution method is not as satisfactory, in terms of numerical accuracy, as solving the system $C x \cong d$ directly by using the routine lin_sol_lsq. Also, see operator_ex05, Chapter 10.

> use lin_sol_self_int
use rand_gen_int
implicit none
! This is Example 1 for LIN_SOL_SELF.
integer, parameter :: m=64, $\mathrm{n}=32$
real(kind(1e0)), parameter : : one=1e0
real(kind(1e0)) err
real(kind(1e0)), dimension(n,n) : A, b, x, res, y(m*n), \& $C(m, n), d(m, n)$
! Generate two rectangular random matrices.
call rand_gen(y)
C = reshape (y, (/m,n/))
call rand_gen (y)
d = reshape (y, (/m,n/))
! Form the normal equations for the rectangular system.
$A=$ matmul (transpose (C), C)
b = matmul(transpose (C), d)
! Compute the solution for $\mathrm{Ax}=\mathrm{b}$.
call lin_sol_self(A, b, x)
! Check the results for small residuals.
res $=\mathrm{b}-\operatorname{matmul}(\mathrm{A}, \mathrm{x})$
err = maxval(abs(res))/sum(abs(A)+abs (b))
if (err <= sqrt(epsilon(one))) then write (*,*) 'Example 1 for LIN_SOL_SELF is correct.'
end if
end

## Output

Example 1 for LIN_SOL_SELF is correct.

## Additional Examples

## Example 2: System Solving with Cholesky Method

This example solves the same form of the system as Example 1. The optional argument "iopt=" is used to note that the Cholesky algorithm is used since the matrix $A$ is positive definite and selfadjoint. In addition, the sample covariance matrix

$$
\Gamma=\sigma^{2} A^{-1}
$$

is computed, where

$$
\sigma^{2}=\frac{\|d-C x\|^{2}}{m-n}
$$

the inverse matrix is returned as the "ainv=" optional argument. The scale factor $\sigma^{2}$ and $\Gamma$ are computed after returning from the routine. Also, see operator_ex0 6, Chapter 10 .

```
    use lin_sol_self_int
    use rand_gen_int
    use error_option_packet
    implicit none
! This is Example 2 for LIN_SOL_SELF.
    integer, parameter :: m=64, n=32
    real(kind(1e0)), parameter :: one=1.0e0, zero=0.0e0
    real(kind(1e0)) err
    real(kind(le0)) a(n,n), b (n,1), c(m,n), d(m,1), cov(n,n), x(n,1), &
        res(n,1), y(m*n)
    type(s_options) :: iopti(1)=s options(0,zero)
! Generate a random rectangular matrix and a random right hand side.
    call rand_gen(y)
    c = reshape(y,(/m,n/))
    call rand_gen(d(1:n,1))
! Form the normal equations for the rectangular system.
    a = matmul(transpose(c),c)
    b = matmul(transpose(c),d)
! Use packaged option to use Cholesky decomposition.
    iopti(1) = s_options(s_lin_sol_self_Use_Cholesky,zero)
! Compute the solution of Ax=b with optional inverse obtained.
    call lin_sol_self(a, b, x, ainv=cov, &
    iopt=iopti)
```

```
! Compute residuals, x - (inverse)*b, for consistency check.
    res = x - matmul (cov,b)
! Scale the inverse to obtain the covariance matrix.
    cov = (sum((d-matmul (c,x))**2)/(m-n)) * cov
! Check the results.
    err = sum(abs(res))/sum(abs(cov))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 2 for LIN_SOL_SELF is correct.'
    end if
    end
```


## Output

Example 2 for LIN_SOL_SELF is correct.

## Example 3: Using Inverse Iteration for an Eigenvector

This example illustrates the use of the optional argument "iopt=" to reset the value of a Small diagonal term encountered during the factorization. Eigenvalues of the self-adjoint matrix

$$
A=C^{T} C
$$

are computed using the routine lin_eig_self. An eigenvector, corresponding to one of these eigenvalues, $\lambda$, is computed using inverse iteration. This solves the near singular system $(A-\lambda I) x=b$ for an eigenvector, $x$. Following the computation of a normalized eigenvector

$$
y=\frac{x}{\|x\|}
$$

the consistency condition

$$
\lambda=y^{T} A y
$$

is checked. Since a singular system is expected, suppress the fatal error message that normally prints when the error post-processor routine error_post is called within the routine lin_sol_self. Also, see operator_ex07, Chapter 10.
use lin_sol_self_int
use lin_eig_self_int
use rand_gen_int
use error_option_packet
implicit none
! This is Example 3 for LIN_SOL_SELF.
integer i, tries

```
    integer, parameter :: m=8, n=4, k=2
    integer ipivots(n+1)
    real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
    real(kind(1dO)) err
    real(kind(1d0)) a(n,n), b(n,1), c(m,n), x(n,1), y(m*n), &
        e(n), atemp (n,n)
    type(d_options) :: iopti(4)
! Generate a random rectangular matrix.
    call rand_gen(y)
    c = reshape(y,(/m,n/))
! Generate a random right hand side for use in the inverse
! iteration.
    call rand_gen(y(1:n))
    b = reshape(y, (/n,1/))
! Compute the positive definite matrix.
    a = matmul(transpose(c),c)
! Obtain just the eigenvalues.
    call lin_eig_self(a, e)
! Use packaged option to reset the value of a small diagonal.
    iopti = d_options(0,zero)
    iopti(1) = d_options(d_lin_sol_self_set_small,&
            epsilon(one) * abs(e(1)))
! Use packaged option to save the factorization.
    iopti(2) = d_options(d_lin_sol_self_save_factors,zero)
! Suppress error messages and stopping due to singularity
! of the matrix, which is expected.
    iopti(3) = d_options(d_lin_sol_self_no_sing_mess,zero)
    atemp = a
    do i=1, n
        a(i,i) = a(i,i) - e(k)
    end do
! Compute A-eigenvalue*I as the coefficient matrix.
    do tries=1, 2
        call lin_sol_self(a, b, x, &
            pivots=ipivots, iopt=iopti)
! When code is re-entered, the already computed factorization
! is used.
            iopti(4) = d_options(d_lin_sol_self_solve_A,zero)
! Reset right-hand si\overline{de nearly in the d\overline{irecti}}\overline{i}on of the eigenvector.
            b = x/sqrt(sum(x**2))
    end do
! Normalize the eigenvector.
    x = x/sqrt(sum(x**2))
```

```
! Check the results.
    err = dot_product(x(1:n,1),matmul(atemp(1:n,1:n),x(1:n,1))) - &
        e(k)
! If any result is not accurate, quit with no summary printing.
    if (abs(err) <= sqrt(epsilon(one))*e(1)) then
        write (*,*) 'Example 3 for LIN_SOL_SELF is correct.'
    end if
    end
```


## Output

Example 3 for LIN_SOL_SELF is correct.

## Example 4: Accurate Least-squares Solution with Iterative Refinement

This example illustrates the accurate solution of the self-adjoint linear system

$$
\left[\begin{array}{cc}
I & A \\
A^{T} & 0
\end{array}\right]\left[\begin{array}{l}
r \\
x
\end{array}\right]=\left[\begin{array}{l}
b \\
0
\end{array}\right]
$$

computed using iterative refinement. This solution method is appropriate for least-squares problems when an accurate solution is required. The solution and residuals are accumulated in double precision, while the decomposition is computed in single precision. Also, see operator_ex08, supplied with the product examples.

```
    use lin_sol_self_int
    use rand_gen_int
    implicit none
! This is Example 4 for LIN_SOL_SELF.
```

    integer i
    integer, parameter : : m=8, \(n=4\)
    real(kind(1e0)), parameter : : one=1.0e0, zero=0.0e0
    real(kind(1d0)), parameter : : d_zero=0.0d0
    integer ipivots \(((n+m)+1)\)
    real (kind (1e0)) a \((m, n), b(m, 1), w(m * n), f(n+m, n+m), \quad \&\)
        \(g(n+m, 1), h(n+m, 1)\)
    real(kind(1e0)) change_new, change_old
    real (kind(1d0)) \(c(m, 1), d(m, n), y(n+m, 1)\)
    type (s_options) : iopti(2)=s_options(0, zero)
    ! Generate a random matrix.
call rand_gen(w)
$\mathrm{a}=\operatorname{reshape}(\mathrm{w},(/ \mathrm{m}, \mathrm{n} /))$
! Generate a random right hand side.

```
    call rand_gen(b(1:m,1))
! Save double precision copies of the matrix and right hand side.
    d = a
    c}=\textrm{b
! Fill in augmented system for accurately solving the least-squares
! problem.
    f = zero
    do i=1, m
        f(i,i) = one
    end do
    f(1:m,m+1:) = a
    f(m+1:,1:m) = transpose(a)
! Start solution at zero.
    y = d_zero
    change_old = huge(one)
! Use packaged option to save the factorization.
    iopti(1) = s_options(s_lin_sol_self_save_factors,zero)
    iterative refinement: do
        g(1:m,\overline{1})=c(1:m,1) - y(1:m,1) - matmul(d,y(m+1:m+n,1))
        g(m+1:m+n,1) = - matmul(transpose(d),y(1:m,1))
        call lin_sol_self(f, g, h, &
                pivots=ipivots, iopt=iopti)
        y=h + y
        change_new = sum(abs(h))
! Exit when changes are no longer decreasing.
        if (change_new >= change_old) &
            exit iterative refinement
        change_old = change__new
! Use option to re-enter code with factorization saved; solve only.
        iopti(2) = s_options(s_lin_sol_self_solve_A,zero)
    end do iterative_refinement
    write (*,*) 'Example 4 for LIN_SOL_SELF is correct.'
    end
```


## Output

Example 4 for LIN_SOL_SELF is correct.

## LIN_SOL_LSQ

Solves a rectangular system of linear equations $A x \cong b$, in a least-squares sense. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of $A$ using column and row pivoting, representing the determinant of $A$, computing the generalized inverse matrix $A^{\dagger}$, or computing the least-squares solution of

$$
A x \cong b
$$

or

$$
A^{T} y \cong b
$$

given the factorization of $A$. An optional argument is provided for computing the following unscaled covariance matrix

$$
C=\left(A^{T} A\right)^{-1}
$$

Least-squares solutions, where the unknowns are non-negative or have simple bounds, can be computed with PARALLEL_NONNEGATIVE_LSQ and PARALLEL_BOUNDED_LSQ. These codes can be restricted to execute without MPI.

## Required Arguments

A- Array of size $m \times n$ containing the matrix. (Input [/Output])
If the packaged option lin_sol_lsq_save_ $Q$ R is used then the factorization of $A$ is saved in A. For efficiency, the diagonal reciprocals of the matrix $R$ are saved in the diagonal entries of $A$.

B - Array of size $m \times n b$ containing the right-hand side matrix. When using the option to solve adjoint systems $A^{T} x \cong b$, the size of $b$ is $n \times n b$. (Input [/Output])
If the packaged option lin_sol_lsq_save_QR is used then input $B$ is used as work storage and is not saved.
$\boldsymbol{X}$ - Array of size $m \times n b$ containing the right-hand side matrix. When using the option to solve adjoint systems $A^{T} x \cong b$, the size of $x$ is $m \times n b$. (Output)

## Optional Arguments

```
MROWS = m (Input)
    Uses array }A(1:m,1:n) for the input matrix.
    Default: m= size(A, 1)
NCOLS = n (Input)
    Uses array A(1:m, 1:n) for the input matrix.
    Default: n = size(A, 2)
```

NRHS = nb (Input)
Uses the array $b(1:, 1: n b)$ for the input right-hand side matrix.
Default: nb $=\operatorname{size}(\mathrm{b}, 2)$
Note that b must be a rank-2 array.
pivots = pivots(:) (Output [/Input])
Integer array of size $2 * \min (m, n)+1$ that contains the individual row followed by the column interchanges. The last array entry contains the approximate rank of A.
trans = trans (: ) (Output [/Input])
Array of size $2 * \min (m, n)$ that contains data for the construction of the orthogonal decomposition.
$\operatorname{det}=\operatorname{det}(1: 2)$ (Output)
Array of size 2 of the same type and kind as A for representing the products of the determinants of the matrices $Q, P$, and $R$. The determinant is represented by two numbers. The first is the base with the sign or complex angle of the result. The second is the exponent. When $\operatorname{det}(2)$ is within exponent range, the value of this expression is given by abs $(\operatorname{det}(1)) * * \operatorname{det}(2) *(\operatorname{det}(1)) / \operatorname{abs}(\operatorname{det}(1))$. If the matrix is not singular, $\operatorname{abs}(\operatorname{det}(1))=\operatorname{radix}(\operatorname{det})$; otherwise, $\operatorname{det}(1)=0$., and $\operatorname{det}(2)=-\operatorname{huge}(\operatorname{abs}(\operatorname{det}(1)))$.
ainv = ainv(:,:) (Output)
Array with size $n \times m$ of the same type and kind as $A(1: m, 1: n)$. It contains the generalized inverse matrix, $A^{\dagger}$.
cov $=\operatorname{cov}(:,:)$ (Output)
Array with size $n \times n$ of the same type and kind as $A(1: m, 1: n)$. It contains the unscaled covariance matrix, $C=\left(A^{T} A\right)^{-1}$.
iopt = iopt(:) (Input)
Derived type array with the same precision as the input matrix; used for passing optional data to the routine. The options are as follows:

| Packaged Options for lin_sol_lsq |  |  |
| :--- | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| s_, d_, c_, z_ | lin_sol_lsq_set_small | 1 |
| s_, d_, c_, z_ | lin_sol_lsq_save_QR | 2 |
| s_, d_, c_, z_ | lin_sol_lsq_solve_A | 3 |
| s_, d_, c_, z_ | lin_sol_lsq_solve_ADJ | 4 |
| s_, d_, c_, z_ | lin_sol_lsq_no_row_pivoting | 5 |
| s_, d_, c_, z_ | lin_sol_lsq_no_col_pivoting | 6 |
| s_, d_, c_, z_ | lin_sol_lsq_scan_for_NaN | 7 |
| s_, d_, c_, z_ | lin_sol_lsq_no_sing_mess | 8 |

```
iopt(IO) = ?_options(?_lin_sol_lsq_set_small, Small)
```

Replaces with Small if a diagonal term of the matrix $R$ is smaller in magnitude than the value Small. A solution is approximated based on this replacement in either case.
Default: the smallest number that can be reciprocated safely

```
iopt(IO) = ?_options(?_lin_sol_lsq_save_QR, ?_dummy)
```

Saves the factorization of A. Requires the optional arguments "pivots=" and "trans=" if the routine is used for solving further systems with the same matrix. This is the only case where the input arrays $A$ and $b$ are not saved. For efficiency, the diagonal reciprocals of the matrix $R$ are saved in the diagonal entries of $A$.

```
iopt(IO) = ?_options(?_lin_sol_lsq_solve_A, ?_dummy)
```

Uses the factorization of A computed and saved to solve $A x=b$.
iopt(IO) = ?_options(?_lin_sol_lsq_solve_ADJ, ?_dummy)
Uses the factorization of A computed and saved to solve $\bar{A}^{T} x=b$.
iopt(IO) = ?_options(?_lin_sol_lsq_no_row_pivoting, ?_dummy)
Does no row pivoting. The $\overline{\operatorname{arrray}} \overline{\mathrm{p} i v o t s}(:)$, if present, satisfies $\overline{\mathrm{p}} \mathrm{ivots}(i)=i$ for $i=1$, $\ldots, \min (m, n)$.
iopt(IO) = ?_options(?_lin_sol_lsq_no_col_pivoting, ?_dummy)
Does no column pivoting. The array pivots(:), if present, satisfies pivots $(i+\min (m$, $\mathrm{n})$ ) $=i$ for $i=1, \ldots, \min (\mathrm{~m}, \mathrm{n})$.
iopt(IO) = ?_options(?_lin_sol_lsq_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(a(i,j)) .or. isNan(b(i,j)) ==.true.
See the isNaN() function, Chapter 10.
Default: Does not scan for NaNs
iopt(IO) = ?_options(?_lin_sol_lsq_no_sing_mess,?_dummy)
Do not print an error message when $A$ is singular or $k<\min (\mathrm{m}, \mathrm{n})$.

## FORTRAN 90 Interface

Generic: CALL LIN_SOL_LSQ (A, B, X [, ...])
Specific: The specific interface names are S_LIN_SOL_LSQ, D_LIN_SOL_LSQ, C_LIN_SOL_LSQ, and Z_LIN_SOL_LSQ.

## Description

Routine LIN_SOL_LSQ solves a rectangular system of linear algebraic equations in a least-squares sense. It computes the decomposition of $A$ using an orthogonal factorization. This decomposition has the form

$$
Q A P=\left[\begin{array}{cc}
R_{k \times k} & 0 \\
0 & 0
\end{array}\right]
$$

where the matrices $Q$ and $P$ are products of elementary orthogonal and permutation matrices. The matrix $R$ is $k \times k$, where $k$ is the approximate rank of $A$. This value is determined by the value of the parameter Small. See Golub and Van Loan (1989, Chapter 5.4) for further details. Note that the use of both row and column pivoting is nonstandard, but the routine defaults to this choice for enhanced reliability.

## Fatal and Terminal Error Messages

See the messages.gls file for error messages for LIN_SOL_LSQ. These error messages are numbered 241-256; 261-276; 281-296; 301-316.

## Example 1: Solving a Linear Least-squares System

This example solves a linear least-squares system $C x \cong d$, where

$$
C_{m \times n}
$$

is a real matrix with $m>n$. The least-squares problem is derived from polynomial data fitting to the function

$$
y(x)=e^{x}+\cos \left(\pi \frac{x}{2}\right)
$$

using a discrete set of values in the interval $-1 \leq x \leq 1$. The polynomial is represented as the series

$$
u(x)=\sum_{i=0}^{N} c_{i} T_{i}(x)
$$

where the $T_{i}(x)$ are Chebyshev polynomials. It is natural for the problem matrix and solution to have a column or entry corresponding to the subscript zero, which is used in this code. Also, see operator_ex09, supplied with the product examples.
use lin_sol_lsq_int
use rand_gen_int
use error_option_packet
implicit none
! This is Example 1 for LIN_SOL_LSQ.
integer i
integer, parameter : : m=128, $n=8$
real(kind(1d0)), parameter : : one=1d0, zero=0d0
real (kind (1d0)) A(m,0:n), c(0:n,1), pi_over_2, x(m), y(m,1), \&
$u(m), v(m), w(m), d e l t a \_x$
! Generate a random grid of points.
call rand_gen (x)
! Transform points to the interval -1,1.

```
    x = x*2 - one
! Compute the constant 'PI/2'.
    pi_over_2 = atan(one)*2
! Generate known function data on the grid.
    y(1:m,1) = exp(x) + cos(pi_over_2*x)
! Fill in the least-squares matrix for the Chebyshev polynomials.
    A(:,0) = one; A(:,1) = x
    do i=2, n
        A(:,i) = 2*X*A(:,i-1) - A(:,i-2)
    end do
! Solve for the series coefficients.
    call lin_sol_lsq(A, y, c)
! Generate an equally spaced grid on the interval.
    delta_x = 2/real(m-1,kind(one))
    do i=1, m
        x(i) = -one + (i-1)*delta_x
    end do
! Evaluate residuals using backward recurrence formulas.
    u = zero
    v = zero
    do i=n, 0, -1
        w = 2*x*u - v + c(i,1)
        v}=
        u = w
    end do
    y(1:m,1) = exp(x) + cos(pi_over_2*x) - (u-x*v)
! Check that n+1 sign changes in the residual curve occur.
    x = one
    x = sign(x,y(1:m,1))
    if (count(x(1:m-1) /= x(2:m)) >= n+1) then
        write (*,*) 'Example 1 for LIN_SOL_LSQ is correct.'
    end if
    end
```


## Output

Example 1 for LIN_SOL_LSQ is correct.

## Additional Examples

## Example 2: System Solving with the Generalized Inverse

This example solves the same form of the system as Example 1. In this case, the grid of evaluation points is equally spaced. The coefficients are computed using the "smoothing formulas" by rows of the generalized inverse matrix, $A^{\dagger}$, computed using the optional argument "ainv=". Thus, the coefficients are given by the matrix-vector product $c=\left(A^{\dagger}\right) y$, where $y$ is the vector of values of the function $y(x)$ evaluated at the grid of points. Also, see operator_ex10, supplied with the product examples.

```
    use lin_sol_lsq_int
    implicit none
! This is Example 2 for LIN_SOL_LSQ.
    integer i
    integer, parameter :: m=128, n=8
    real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
    real(kind(1d0)) a(m,0:n), c(0:n,1), pi_over_2, x(m), y(m,1), &
    u(m), v(m), w(m), delta x, inv(0:n, m)
! Generate an array of equally spaced points on the interval -1,1.
    delta_x = 2/real(m-1,kind(one))
    do i=1, m
        x(i) = -one + (i-1)*delta_x
    end do
! Compute the constant 'PI/2'.
    pi over 2 = atan(one)*2
! Compute data values on the grid.
    y(1:m,1) = exp(x) + cos(pi_over_2*x)
! Fill in the least-squares matrix for the Chebyshev polynomials.
    a(:,0) = one
    a(:,1) = x
    do i=2, n
        a(:,i) = 2*x*a(:,i-1) - a(:,i-2)
    end do
! Compute the generalized inverse of the least-squares matrix.
    call lin_sol_lsq(a, y, c, nrhs=0, ainv=inv)
! Compute the series coefficients using the generalized inverse
! as 'smoothing formulas.'
```

```
    c(0:n,1) = matmul(inv(0:n,1:m),y(1:m,1))
! Evaluate residuals using backward recurrence formulas.
    u = zero
    v = zero
    do i=n, 0, -1
        w = 2*x*u - v + c(i,1)
        v = u
        u = w
    end do
    y(1:m,1) = exp(x) + cos(pi_over_2*x) - (u-x*v)
! Check that n+2 sign changes in the residual curve occur.
! (This test will fail when n is larger.)
    x = one
    x = sign(x,y(1:m,1))
    if (count(x(1:m-1) /= x(2:m)) == n+2) then
        write (*,*) 'Example 2 for LIN_SOL_LSQ is correct.'
    end if
    end
```


## Output

Example 2 for LIN_SOL_LSQ is correct.

## Example 3: Two-Dimensional Data Fitting

This example illustrates the use of radial-basis functions to least-squares fit arbitrarily spaced data points. Let $m$ data values $\left\{y_{i}\right\}$ be given at points in the unit square, $\left\{p_{i}\right\}$. Each $p_{i}$ is a pair of real values. Then, $n$ points $\left\{q_{j}\right\}$ are chosen on the unit square. A series of radial-basis functions is used to represent the data,

$$
f(p)=\sum_{j=1}^{n} c_{j}\left(\left\|p-q_{j}\right\|^{2}+\delta^{2}\right)^{1 / 2}
$$

where $\delta^{2}$ is a parameter. This example uses $\delta^{2}=1$, but either larger or smaller values can give a better approximation for user problems. The coefficients $\left\{c_{j}\right\}$ are obtained by solving the following $m \times n$ linear least-squares problem:

$$
f\left(p_{j}\right)=y_{j}
$$

This example illustrates an effective use of Fortran 90 array operations to eliminate many details required to build the matrix and right-hand side for the $\left\{c_{j}\right\}$. For this example, the two sets of points $\left\{p_{i}\right\}$ and $\left\{q_{j}\right\}$ are chosen randomly. The values $\left\{y_{j}\right\}$ are computed from the following formula:

$$
y_{j}=e^{-\left\|p_{j}\right\|^{2}}
$$

The residual function

$$
r(p)=e^{-\|p\|^{2}}-f(p)
$$

is computed at an $N \times N$ square grid of equally spaced points on the unit square. The magnitude of $r(p)$ may be larger at certain points on this grid than the residuals at the given points, $\left\{p_{i}\right\}$. Also, see operator_ex11, supplied with the product examples.
use lin_sol_lsq_int
use rand_gen_int
implicit none
! This is Example 3 for LIN_SOL_LSQ.
integer i, j
integer, parameter : : m=128, $n=32, k=2$, $n$ eval=16
real(kind(1d0)), parameter : : one=1.0d0, delta_sqr=1.0d0
real(kind(1d0)) $a(m, n), b(m, 1), c(n, 1), p(k, m), q(k, n), \&$ $x(k * m), y(k * n), t(k, m, n), r e s\left(n \_e v a l, n \_e v a l\right), ~ \& ~$ w(n_eval), delta
! Generate a random set of data points in $k=2$ space.
call rand gen (x)
$p=\operatorname{reshape}(x,(/ k, m /))$
! Generate a random set of center points in k-space.
call rand_gen (y)
$q=\operatorname{reshape}(y,(/ k, n /))$
! Compute the coefficient matrix for the least-squares system.
$t=\operatorname{spread}(p, 3, n)$
do $j=1$, $n$
$t(1:,:, j)=t(1:,:, j)-\operatorname{spread}(q(1:, j), 2, m)$
end do
$a=\operatorname{sqrt}(\operatorname{sum}(t * * 2, d i m=1)+$ delta_sqr)
! Compute the right hand side of data values.
$b(1:, 1)=\exp (-\operatorname{sum}(p * * 2, \operatorname{dim}=1))$
! Compute the solution.
call lin_sol_lsq(a, b, c)
! Check the results.
if (sum(abs(matmul(transpose (a),b-matmul(a,c))))/sum(abs(a)) \& $<=$ sqrt(epsilon(one))) then write (*,*) 'Example 3 for LIN_SOL_LSQ is correct.'

```
    end if
! Evaluate residuals, known function - approximation at a square
! grid of points. (This evaluation is only for k=2.)
    delta = one/real(n_eval-1,kind(one))
    do i=1, n_eval
        w(i) = (i-1)*delta
    end do
    res = exp(-(spread(w,1,n_eval)**2 + spread(w,2,n_eval)**2))
    do j=1, n
        res = res - c(j,1)*sqrt((spread(w,1,n_eval) - q(1,j))**2 + &
        (spread(w,2,n_eval) - q(2,j))**2 + delta_sqr)
    end do
    end
```


## Output

```
Example 3 for LIN_SOL_LSQ is correct.
```


## Example 4: Least-squares with an Equality Constraint

This example solves a least-squares system $A x \cong b$ with the constraint that the solution values have a sum equal to the value 1 . To solve this system, one heavily weighted row vector and right-hand side component is added to the system corresponding to this constraint. Note that the weight used is

$$
\varepsilon^{-1 / 2}
$$

where $\varepsilon$ is the machine precision, but any larger value can be used. The fact that lin_sol_lsq performs row pivoting in this case is critical for obtaining an accurate solution to the constrained problem solved using weighting. See Golub and Van Loan (1989, Chapter 12) for more information about this method. Also, see operator_ex12, supplied with the product examples.

```
        use lin_sol_lsq_int
        use rand_gen_int
        implicit none
! This is Example 4 for LIN_SOL_LSQ.
    integer, parameter :: m=64, n=32
    real(kind(1e0)), parameter :: one=1.0e0
    real(kind(1e0)) : : a (m+1,n), b (m+1,1), x(n,1), y(m*n)
! Generate a random matrix.
    call rand gen(y)
    a(1:m,1:n) = reshape(y,(/m,n/))
! Generate a random right hand side.
```

```
    call rand_gen(b(1:m,1))
! Heavily weight desired constraint. All variables sum to one.
    a(m+1,1:n) = one/sqrt(epsilon(one))
    b(m+1,1) = one/sqrt(epsilon(one))
    call lin_sol_lsq(a, b, x)
    if (abs(sum(x) - one)/sum(abs(x)) <= &
                sqrt(epsilon(one))) then
        write (*,*) 'Example 4 for LIN_SOL_LSQ is correct.'
    end if
    end
```


## Output

Example 4 for LIN_SOL_LSQ is correct.

## LIN_SOL_SVD

Solves a rectangular least-squares system of linear equations $A x \cong b$ using singular value decomposition

$$
A=U S V^{T}
$$

With optional arguments, any of several related computations can be performed. These extra tasks include computing the rank of $A$, the orthogonal $m \times m$ and $n \times n$ matrices $U$ and $V$, and the $m \times n$ diagonal matrix of singular values, $S$.

## Required Arguments

A- Array of size $m \times n$ containing the matrix. (Input [/Output])
If the packaged option lin_sol_svd_overwrite_input is used, this array is not saved on output.
$\boldsymbol{B}$ - Array of size $m \times n b$ containing the right-hand side matrix. (Input [/Output] If the packaged option lin_sol_svd_overwrite_input is used, this array is not saved on output.
$\boldsymbol{X}$ - Array of size $n \times n b$ containing the solution matrix. (Output)

## Optional Arguments

```
MROWS = m (Input)
    Uses array A(1:m, 1:n) for the input matrix.
    Default: m= size (A, 1)
```

$\mathrm{NCOLS}=\mathrm{n}$ (Input)
Uses array $A(1: m, 1: n)$ for the input matrix.
Default: $\mathrm{n}=\operatorname{size}(\mathrm{A}, 2)$

NRHS = nb (Input)
Uses the array $b(1:, 1: n b)$ for the input right-hand side matrix.
Default: nb $=\operatorname{size}(\mathrm{b}, 2)$
Note that b must be a rank-2 array.

RANK $=k$ (Output)
Number of singular values that are at least as large as the value Small. It will satisfy k
$<=\min (m, n)$.
$u=u(:,:)$ (Output)
Array of the same type and kind as $A(1: m, 1: n)$. It contains the $m \times m$ orthogonal matrix $U$ of the singular value decomposition.
$s=s(:)$ (Output)
Array of the same precision as $A(1: m, 1: n)$. This array is real even when the matrix data is complex. It contains the $m \times n$ diagonal matrix $S$ in a rank-1 array. The singular values are nonnegative and ordered non-increasing.
$\mathrm{v}=\mathrm{v}(:,:$ ) (Output)
Array of the same type and kind as $\mathrm{A}(1: \mathrm{m}, 1: \mathrm{n})$. It contains the $n \times n$ orthogonal matrix V.

```
iopt = iopt(:) (Input)
```

Derived type array with the same precision as the input matrix. Used for passing optional data to the routine. The options are as follows:

| Packaged Options for lin_sol_svd |  |  |
| :---: | :---: | :---: |
| Option Prefix $=$ ? | Option Name | Option Value |
| s_, d_, c_, z_ | lin_sol_svd_set_small | 1 |
| s_, d_, c_, z_ | lin_sol_svd_overwrite_input | 2 |
| s_, d_, c_, z_ | lin_sol_svd_safe_reciprocal | 3 |
| s_, d_, c_, z_ | lin_sol_svd_scan_for_NaN | 4 |

iopt(IO) = ?_options(?_lin_sol_svd_set_small, Small)
Replaces with zero a diagonal term of the matrix $S$ if it is smaller in magnitude than the value Small. This determines the approximate rank of the matrix, which is returned as the "rank=" optional argument. A solution is approximated based on this replacement.
Default: the smallest number that can be safely reciprocated

```
iopt(IO) = ?_options(?_lin_sol_svd_overwrite_input,?_dummy)
```

Does not save the input arrays $A(:,:)$ and $b(:,:)$.
iopt(IO) = ?_options(?_lin_sol_svd_safe_reciprocal, safe)
Replaces a denominator term with safe if it is smaller in magnitude than the value safe.
Default: the smallest number that can be safely reciprocated
iopt(IO) = ?_options(?_lin_sol_svd_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN( $\mathrm{a}(\mathrm{i}, \mathrm{j}))$.or. isNan $(\mathrm{b}(\mathrm{i}, \mathrm{j}))==$. true.
See the isNaN() function, Chapter 10.
Default: Does not scan for NaNs

## FORTRAN 90 Interface

Generic: CALL LIN_SOL_SVD (A, B, X [, ...])
Specific: The specific interface names are S_LIN_SOL_SVD, D_LIN_SOL_SVD, C_LIN_SOL_SVD, and Z_LIN_SOL_SVD.

## Description

Routine LIN_SOL_SVD solves a rectangular system of linear algebraic equations in a least-squares sense. It computes the factorization of $A$ known as the singular value decomposition. This decomposition has the following form:

$$
A=U S V^{T}
$$

The matrices $U$ and $V$ are orthogonal. The matrix $S$ is diagonal with the diagonal terms non-increasing. See Golub and Van Loan (1989, Chapters 5.4 and 5.5) for further details.

## Fatal, Terminal, and Warning Error Messages

See the messages.gls file for error messages for LIN_SOL_SVD. These error messages are numbered 401-412; 421-432; 441-452; 461-472.

## Example 1: Least-squares solution of a Rectangular System

The least-squares solution of a rectangular $m \times n$ system $A x \cong b$ is obtained. The use of lin_sol_lsq is more efficient in this case since the matrix is of full rank. This example anticipates a problem where the matrix $A$ is poorly conditioned or not of full rank; thus, lin_sol_svd is the appropriate routine. Also, see operator_ex13, Chapter 10.

```
use lin_sol_svd_int
```

use rand_gen_int
implicit none
! This is Example 1 for LIN_SOL_SVD.

```
    integer, parameter : : m=128, n=32
    real(kind(1d0)), parameter :: one=1d0
    real(kind(1d0)) A(m,n), b(m,1), x(n,1), y(m*n), err
! Generate a random matrix and right-hand side.
    call rand_gen(y)
    A = reshape(y,(/m,n/))
    call rand_gen(b(1:m,1))
! Compute the least-squares solution matrix of Ax=b.
    call lin_sol_svd(A, b, x)
! Check that the residuals are orthogonal to the
! column vectors of A.
    err = sum(abs(matmul(transpose(A),b-matmul(A,x))))/sum(abs(A))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_SOL_SVD is correct.'
    end if
    end
```


## Output

Example 1 for LIN_SOL_SVD is correct.

## Additional Examples

## Example 2: Polar Decomposition of a Square Matrix

A polar decomposition of an $n \times n$ random matrix is obtained. This decomposition satisfies $A=P Q$, where $P$ is orthogonal and $Q$ is self-adjoint and positive definite.

Given the singular value decomposition

$$
A=U S V^{T}
$$

the polar decomposition follows from the matrix products

$$
P=U V^{T} \text { and } Q=V S V^{T}
$$

This example uses the optional arguments " $u="$, " $s=$ ", and " $v=$ ", then array intrinsic functions to calculate $P$ and $Q$. Also, see operator_ex14, Chapter 10.

```
    use lin_sol_svd_int
    use ran\overline{d}_gen_in\overline{t}
    implicit none
! This is Example 2 for LIN_SOL_SVD.
integer i
integer, parameter :: n=32
real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
```

```
    real(kind(1d0)) a(n,n), b(n,0), ident(n,n), p(n,n), q(n,n), &
        s_d(n), u_d(n,n), v_d(n,n), x(n,0), y(n*n)
! Generate a random matrix.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
! Compute the singular value decomposition.
    call lin_sol_svd(a, b, x, nrhs=0, s=s_d, &
    u=u_d, v=v_d)
! Compute the (left) orthogonal factor.
    p = matmul(u_d,transpose(v_d))
! Compute the (right) self-adjoint factor.
    q = matmul(v_d*spread(s_d,1,n),transpose(v_d))
    ident=zero
    do i=1, n
        ident(i,i) = one
    end do
! Check the results.
    if (sum(abs(matmul(p,transpose(p)) - ident))/sum(abs(p)) &
                <= sqrt(epsilon(one))) then
        if (sum(abs(a - matmul(p,q)))/sum(abs(a)) &
                <= sqrt(epsilon(one))) then
            write (*,*) 'Example 2 for LIN_SOL_SVD is correct.'
        end if
    end if
    end
```


## Output

Example 2 for LIN_SOL_SVD is correct.

## Example 3: Reduction of an Array of Black and White

An $n \times n$ array $A$ contains entries that are either 0 or 1 . The entry is chosen so that as a twodimensional object with origin at the point $(1,1)$, the array appears as a black circle of radius $n / 4$ centered at the point $(n / 2, n / 2)$.
A singular value decomposition

$$
A=U S V^{T}
$$

is computed, where $S$ is of low rank. Approximations using fewer of these nonzero singular values and vectors suffice to reconstruct $A$. Also, see operator_ex15, supplied with the product examples.

```
    use lin_sol_svd_int
    use ran\overline{d}ge\overline{n}_in\overline{t}
    use error_option_packet
    implicit none
! This is Example 3 for LIN_SOL_SVD.
    integer i, j, k
    integer, parameter :: n=32
    real(kind(1e0)), parameter :: half=0.5e0, one=1e0, zero=0e0
    real(kind(1e0)) a (n,n), b (n,0), x(n,0), s(n), u(n,n), &
    v(n,n), c(n,n)
! Fill in value one for points inside the circle.
    a = zero
    do i=1, n
        do j=1, n
            if ((i-n/2)**2 + (j-n/2)**2 <= (n/4)**2) a(i,j) = one
        end do
    end do
! Compute the singular value decomposition.
    call lin_sol_svd(a, b, x, nrhs=0, &
        s=\overline{s},u=u,v=v)
! How many terms, to the nearest integer, exactly
! match the circle?
            c = zero; k = count(s > half)
    do i=1, k
        c = c + spread(u(1:n,i),2,n)*spread(v(1:n,i),1,n)*s(i)
        if (count(int(c-a) /= 0) == 0) exit
    end do
    if (i < k) then
        write (*,*) 'Example 3 for LIN_SOL_SVD is correct.'
    end if
    end
```


## Output

Example 3 for LIN_SOL_SVD is correct.

## Example 4: Laplace Transform Solution

This example illustrates the solution of a linear least-squares system where the matrix is poorly conditioned. The problem comes from solving the integral equation:

$$
\int_{0}^{1} e^{-s t} f(t) d t=s^{-1}\left(1-e^{-s}\right)=g(s)
$$

The unknown function $f(t)=1$ is computed. This problem is equivalent to the numerical inversion of the Laplace Transform of the function $g(s)$ using real values of $t$ and $s$, solving for a function
that is nonzero only on the unit interval. The evaluation of the integral uses the following approximate integration rule:

$$
\int_{0}^{1} f(t) e^{-s t} d t=\sum_{j=1}^{n} f\left(t_{j}\right) \int_{t_{j}}^{t_{j+1}} e^{-s t} d t
$$

The points $\left\{t_{j}\right\}$ are chosen equally spaced by using the following:

$$
t_{j}=\frac{j-1}{n}
$$

The points $\left\{s_{j}\right\}$ are computed so that the range of $g(s)$ is uniformly sampled. This requires the solution of $m$ equations

$$
g\left(s_{i}\right)=g_{i}=\frac{i}{m+1}
$$

for $j=1, \ldots, n$ and $i=1, \ldots, m$. Fortran 90 array operations are used to solve for the collocation points $\left\{s_{i}\right\}$ as a single series of steps. Newton's method,

$$
s \leftarrow s-\frac{h}{h^{\prime}}
$$

is applied to the array function

$$
h(s)=e^{-s}+s g-1
$$

where the following is true:

$$
g=\left[g_{1}, \ldots, g_{m}\right]^{T}
$$

Note the coefficient matrix for the solution values

$$
f=\left[f\left(t_{1}\right), \ldots, f\left(t_{n}\right)\right]^{T}
$$

whose entry at the intersection of row $i$ and column $j$ is equal to the value

$$
\int_{t_{j}}^{t_{j+1}} e^{-s_{i} t} d t
$$

is explicitly integrated and evaluated as an array operation. The solution analysis of the resulting linear least-squares system

$$
A f \cong g
$$

is obtained by computing the singular value decomposition

$$
A=U S V^{T}
$$

An approximate solution is computed with the transformed right-hand side

$$
b=U^{T} g
$$

followed by using as few of the largest singular values as possible to minimize the following squared error residual:

$$
\sum_{j=1}^{n}\left(1-f_{j}\right)^{2}
$$

This determines an optimal value $k$ to use in the approximate solution

$$
f=\sum_{j=1}^{k} b_{j} \frac{v_{j}}{s_{j}}
$$

Also, see operator_ex16, supplied with the product examples.

```
    use lin_sol_svd_int
    use rand_gen_int
    use error_option_packet
    implicit none
! This is Example 4 for LIN_SOL_SVD.
    integer i, j, k
    integer, parameter :: m=64, n=16
    real(kind(le0)), parameter :: one=1e0, zero=0.0e0
    real(kind(le0)) :: g(m), s(m), t(n+1), a(m,n), b(m,1), &
        f(n,1), U_S (m,m), V_S (n,n), S_S(n), &
        rms, oldrms
    real(kind(1e0)) :: delta_g, delta_t
    delta_g = one/real(m+1,kind(one))
! Compute which collocation equations to solve.
    do i=1,m
        g(i)=i*delta_g
    end do
! Compute equally spaced quadrature points.
    delta t =one/real(n,kind(one))
    do j=\overline{1},n+1
        t(j)=(j-1)*delta_t
    end do
! Compute collocation points.
    s=m
    solve_equations: do
        s=s=
        if (sum(abs((one-exp(-s))/s - g)) <= &
            epsilon(one)*sum(g)) &
            exit solve equations
    end do solve_equātions
! Evaluate the integrals over the quadrature points.
```

```
        a = (exp(-spread(t(1:n),1,m)*spread(s,2,n)) &
            - exp(-spread(t(2:n+1),1,m)*spread(s,2,n))) / &
        spread(s,2,n)
    b (1:,1)=g
! Compute the singular value decomposition.
    call lin_sol_svd(a, b, f, nrhs=0, &
        rank=k, u=U_S, v=V_S, s=S_S)
! Singular values that are larger than epsilon determine
! the rank=k.
    k = count(S_S > epsilon(one))
    oldrms = huge(one)
    g = matmul(transpose(U_S), b(1:m,1))
! Find the minimum number of singular values that gives a good
! approximation to f(t) = 1.
    do i=1,k
        f(1:n,1) = matmul(V_S(1:,1:i), g(1:i)/S_S(1:i))
        f = f - one
        rms = sum(f**2)/n
        if (rms > oldrms) exit
        oldrms = rms
    end do
    write (*,"( ' Using this number of singular values, ', &
        &i4 / ' the approximate R.M.S. error is ', 1pe12.4)") &
    i-1, oldrms
    if (sqrt(oldrms) <= delta t**2) then
        write (*,*) 'Example 4 for LIN_SOL_SVD is correct.'
    end if
    end
```


## Output

Example 4 for LIN_SOL_SVD is correct.

## LIN_SOL_TRI

Solves multiple systems of linear equations

$$
A_{j} x_{j}=y_{j}, j=1, \ldots, k
$$

Each matrix $A_{j}$ is tridiagonal with the same dimension, $n$. The default solution method is based on $L U$ factorization computed using cyclic reduction or, optionally, Gaussian elimination with partial pivoting.

## Required Arguments

$\boldsymbol{C}$ - Array of size $2 n \times k$ containing the upper diagonals of the matrices $A_{j}$. Each upper diagonal is entered in array locations $\mathrm{c}(1: n-1, j)$. The data $C(n, 1: k)$ are not used. (Input [/Output])
The input data is overwritten. See note below.
$\boldsymbol{D}$ - Array of size $2 n \times k$ containing the diagonals of the matrices $A_{j}$. Each diagonal is entered in array locations $D(1: n, j)$. (Input [/Output])
The input data is overwritten. See note below.
$\boldsymbol{B}$ - Array of size $2 n \times k$ containing the lower diagonals of the matrices $A_{j}$. Each lower diagonal is entered in array locations $\mathrm{B}(2: n, j)$. The data
$B(1,1: k)$ are not used. (Input [/Output])
The input data is overwritten. See note below.
$\boldsymbol{Y}$ - Array of size $2 n \times k$ containing the right-hand sides, $y_{j}$. Each right-hand side is entered in array locations $Y(1: n, j)$. The computed solution $x_{j}$ is returned in locations $Y(1: n, j)$. (Input [/Output])

NOTE: The required arguments have the Input data overwritten. If these quantities are used later, they must be saved in user-defined arrays. The routine uses each array's locations ( $\mathrm{n}+1: 2 * \mathrm{n}, 1: \mathrm{k}$ ) for scratch storage and intermediate data in the LU factorization. The default values for problem dimensions are $n=(\operatorname{size}(D, 1)) / 2$ and $\mathrm{k}=\operatorname{size}(\mathrm{D}, 2)$.

## Optional Arguments

NCOLS $=\mathrm{n}$ (Input)
Uses arrays $C(1: n-1,1: k), D(1: n, 1: k)$, and $B(2: n, 1: k)$ as the upper, main and lower diagonals for the input tridiagonal matrices. The right-hand sides and solutions are in array $\mathrm{Y}(1: \mathrm{n}, 1: \mathrm{k})$. Note that each of these arrays are rank-2.
Default: $\mathrm{n}=(\operatorname{size}(\mathrm{D}, 1)) / 2$

NPROB $=k$ (Input)
The number of systems solved.
Default: $k=\operatorname{size}(\mathrm{D}, 2)$
iopt = iopt(:) (Input)
Derived type array with the same precision as the input matrix. Used for passing optional data to the routine. The options are as follows:

## Packaged Options for LIN_SOL_TRI

| Option Prefix =? | Option Name | Option Value |
| :---: | :---: | :---: |
| s_, $d_{-}, c_{-}, z_{-}$ | lin_sol_tri_set_small | 1 |


| Packaged Options for LIN_SOL_TRI |  |  |
| :---: | :---: | :---: |
| s_, d_, c_, z_ | lin_sol_tri_set_jolt | 2 |
| s_, d_, c_, z_ | lin_sol_tri_scan_for_NaN | 3 |
| s_, d_, c_, $z_{-}$ | lin_sol_tri_factor_only | 4 |
| s_, d_, c_, z_ | lin_sol_tri_solve_only | 5 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_sol_tri_use_Gauss_elim | 6 |

iopt(IO) $=$ ?_options(?_lin_sol_tri_set_small, Small)
Whenever a reciprocation is performed on a quantity smaller than Small, it is replaced by that value plus $2 \times$ jolt.
Default: $0.25 \times$ epsilon ()
iopt(IO) = ?_options(?_lin_sol_tri_set_jolt, jolt)
Default: epsilon(), machine precision
iopt(IO) = ?_options(?_lin_sol_tri_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(C(i,j)) .or.
isNaN(D(i,j)) .or.
isNaN(B(i,j)) .or.
isNaN(Y(i,j)) == .true.
See the isnan () function, Chapter 10.
Default: Does not scan for NaNs.
iopt(IO) = ?_options(?_lin_sol_tri_factor_only, ?_dummy)
Obtain the $L U$ factorization of the matrices $A_{j}$. Does not solve for a solution.
Default: Factor the matrices and solve the systems.
iopt(IO) = ?_options(?_lin_sol_tri_solve_only, ?_dummy)
Solve the systems $A_{j} x_{j}=y_{j}$ using the previously computed $\bar{L} U$ factorization.
Default: Factor the matrices and solve the systems.
iopt (IO) = ?_options(?_lin_sol_tri_use_Gauss_elim, ?_dummy)
The accuracy, numerical stability or efficiency of the cyclic reduction algorithm may be inferior to the use of $L U$ factorization with partial pivoting.
Default: Use cyclic reduction to compute the factorization.

## FORTRAN 90 Interface

Generic: CALL LIN_SOL_TRI (C, D, B, Y [ , ...])
Specific: The specific interface names are S_LIN_SOL_TRI, D_LIN_SOL_TRI, C_LIN_SOL_TRI, and Z_LIN_SOL_TRI.

## Description

Routine lin_sol_tri solves $k$ systems of tridiagonal linear algebraic equations, each problem of dimension $n \times n$. No relation between $k$ and $n$ is required. See Kershaw, pages 86-88 in Rodrigue (1982) for further details. To deal with poorly conditioned or singular systems, a specific regularizing term is added to each reciprocated value. This technique keeps the factorization process efficient and avoids exceptions from overflow or division by zero. Each occurrence of an array reciprocal $a^{-1}$ is replaced by the expression $(a+t)^{-1}$, where the array temporary $t$ has the value 0 whenever the corresponding entry satisfies $|a|>$ Small. Alternately, $t$ has the value $2 \times$ jolt. (Every small denominator gives rise to a finite "jolt".) Since this tridiagonal solver is used in the routines lin_svd and lin_eig_self for inverse iteration, regularization is required. Users can reset the values of Small and jolt for their own needs. Using the default values for these parameters, it is generally necessary to scale the tridiagonal matrix so that the maximum magnitude has value approximately one. This is normally not an issue when the systems are nonsingular.

The routine is designed to use cyclic reduction as the default method for computing the $L U$ factorization. Using an optional parameter, standard elimination and partial pivoting will be used to compute the factorization. Partial pivoting is numerically stable but is likely to be less efficient than cyclic reduction.

## Fatal, Terminal, and Warning Error Messages

See the messages.gls file for error messages for LIN_SOL_TRI. These error messages are numbered 1081-1086; 1101-1106; 1121-1126; 1141-1146.

## Example 1: Solution of Multiple Tridiagonal Systems

The upper, main and lower diagonals of $n$ systems of size $n \times n$ are generated randomly. A scalar is added to the main diagonal so that the systems are positive definite. A random vector $x_{j}$ is generated and right-hand sides $y_{j}=A_{j} y_{j}$ are computed. The routine is used to compute the solution, using the $A_{j}$ and $y_{j}$. The results should compare closely with the $x_{j}$ used to generate the right-hand sides. Also, see operator_ex17, supplied with the product examples.

```
        use lin sol tri int
        use ran\overline{d_gen_in}\overline{t}
        use error_option_packet
        implicit none
! This is Example 1 for LIN_SOL_TRI.
    integer i
    integer, parameter :: n=128
    real(kind(1d0)), parameter :: one=1d0, zero=0d0
    real(kind(1d0)) err
    real(kind(1d0)), dimension(2*n,n) : : d, b, c, res(n,n), &
        t(n), x, y
! Generate the upper, main, and lower diagonals of the
! n matrices A_i. For each system a random vector x is used
```

```
! to construct the right-hand side, Ax = y. The lower part
! of each array remains zero as a result.
    c = zero; d=zero; b=zero; x=zero
    do i = 1, n
        call rand_gen (c(1:n,i))
        call rand_gen (d(1:n,i))
        call rand_gen (b(1:n,i))
        call rand_gen (x(1:n,i))
    end do
! Add scalars to the main diagonal of each system so that
! all systems are positive definite.
    t = sum(c+d+b,DIM=1)
    d(1:n,1:n) = d(1:n,1:n) + spread(t,DIM=1,NCOPIES=n)
! Set Ax = y. The vector x generates y. Note the use
! of EOSHIFT and array operations to compute the matrix
! product, n distinct ones as one array operation.
    y(1:n,1:n)=d(1:n,1:n)*x(1:n,1:n) + &
        c(1:n,1:n)*EOSHIFT(x (1:n,1:n),SHIFT=+1,DIM=1) + &
        b (1:n,1:n)*EOSHIFT(x(1:n,1:n),SHIFT=-1,DIM=1)
! Compute the solution returned in y. (The input values of c,
! d, b, and y are overwritten by lin_sol_tri.) Check for any
! error messages.
    call lin_sol_tri (c, d, b, y)
! Check the size of the residuals, y-x. They should be small,
! relative to the size of values in x.
    res = x(1:n,1:n) - y(1:n,1:n)
    err = sum(abs(res)) / sum(abs(x(1:n,1:n)))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_SOL_TRI is correct.'
    end if
    end
```


## Output

Example 1 for LIN SOL_TRI is correct.

## Additional Examples

## Example 2: Iterative Refinement and Use of Partial Pivoting

This program unit shows usage that typically gives acceptable accuracy for a large class of problems. Our goal is to use the efficient cyclic reduction algorithm when possible, and keep on using it unless it will fail. In exceptional cases our program switches to the $L U$ factorization with partial pivoting. This use of both factorization and solution methods enhances reliability and maintains efficiency on the average. Also, see operator_ex18, supplied with the product examples.

```
    use lin_sol_tri_int
    use ran\overline{d}_gen_int
    implicit none
! This is Example 2 for LIN_SOL_TRI.
    integer i, nopt
    integer, parameter : : n=128
    real(kind(1e0)), parameter :: s one=1e0, s zero=0e0
    real(kind(1d0)), parameter :: d_one=1d0, d_zero=0d0
    real(kind(1e0)), dimension(2*n,n) : : d, b, c, res(n,n), &
        x, y
    real(kind(1e0)) change_new, change_old, err
    type(s_options) : : iopt(2) = s_options(0,s_zero)
```



```
        x_save, y_save, x_sol
    logical solve_only
    c = s_zero; d=s_zero; b=s_zero; x=s_zero
! Generate the upper, main, and lower diagonals of the
! matrices A. A random vector x is used to construct the
! right-hand sides: y=A*x.
    do i = 1, n
        call rand_gen (c(1:n,i))
        call rand_gen (d(1:n,i))
        call rand gen (b(1:n,i))
        call rand_gen (x(1:n,i))
    end do
! Save double precision copies of the diagonals and the
! right-hand side.
    c_save = c(1:n,1:n); d_save = d(1:n,1:n)
    b_save = b(1:n,1:n); x_save = x(1:n,1:n)
    Y_save(1:n,1:n) = d(1:n,1:n)*x_save + &
                c(1:n,1:n)*EOSHIFT(x_S_save,SHIFT=+1,DIM=1) + &
                b(1:n,1:n)*EOSHIFT(x_save,SHIFT=-1,DIM=1)
! Iterative refinement loop.
    factorization_choice: do nopt=0, 1
! Set the logical to flag the first time through.
    solve_only = .false.
    x_sol = d_zero
    c\overline{hange_ol\overline{d}}=\mathrm{ huge(s_one)}
    iterative_refinement: do
! This flag causes a copy of data to be moved to work arrays
! and a factorization and solve step to be performed.
```

```
    if (.not. solve_only) then
        c(1:n,1:n)=c_save; d(1:n,1:n)=d_save
        b (1:n,1:n)=b_save
    end if
! Compute current residuals, y - A*x, using current x.
    y(1:n,1:n) = -y_save + &
        d_save*x_sol + &
        c_save*EO-SHIFT(x_sol,SHIFT=+1,DIM=1) + &
        b_save*EOSHIFT(x_sol,SHIFT=-1,DIM=1)
    call lin_sol_tri (c, d, b, y, iopt=iopt)
    x_sol = x_sol + y(1:n,1:n)
    change_new = sum(abs(y(1:n,1:n)))
! If size of change is not decreasing, stop the iteration.
            if (change_new >= change_old) exit iterative_refinement
            change_old = change_new
            iopt(n\overline{ppt+1) = s_options(s_lin_sol_tri_solve_only,s_zero)}
            solve_only = .true.
    end do iterative_refinement
! Use Gaussian Elimination if Cyclic Reduction did not get an
! accurate solution.
! It is an exceptional event when Gaussian Elimination is required.
    if (sum(abs(x_sol - x_save)) / sum(abs(x_save)) &
            <= sqrt(epsīion(d_oñe))) exit factoriz\overline{ation_choice}
        iopt = s_options(0,s_zero)
        iopt(nopt+1) = s_options(s_lin_sol_tri_use_Gauss_elim,s_zero)
    end do factorization_choice
! Check on accuracy of solution.
    res = x(1:n,1:n)- x_save
    err = sum(abs(res)) / sum(abs(x save))
    if (err <= sqrt(epsilon(d_one))) then
        write (*,*) 'Example 2 for LIN_SOL_TRI is correct.'
    end if
    end
```


## Output

Example 2 for LIN_SOL_TRI is correct.
Example 3: Selected Eigenvectors of Tridiagonal Matrices
The eigenvalues

$$
\lambda_{1}, \ldots, \lambda_{n}
$$

of a tridiagonal real, self-adjoint matrix are computed. Note that the computation is performed using the IMSL MATH/LIBRARY FORTRAN 77 interface to routine EVASB. The user may write this interface based on documentation of the arguments (IMSL 2003, p. 480), or use the module Numerical_Libraries as we have done here. The eigenvectors corresponding to $k<n$ of the eigenvalues are required. These vectors are computed using inverse iteration for all the eigenvalues at one step. See Golub and Van Loan (1989, Chapter 7). The eigenvectors are then orthogonalized. Also, see operator_ex19, supplied with the product examples.

```
    use lin_sol_tri_int
    use rand_gen_int
    use Numerical Libraries
    implicit none
! This is Example 3 for LIN_SOL_TRI.
    integer i, j, nopt
    integer, parameter :: n=128, k=n/4, ncoda=1, lda=2
    real(kind(1e0)), parameter :: s one=1e0, s_zero=0e0
    real(kind(le0)) A(lda,n), EVAL(k}
    type(s_options) :: iopt(2)=s_options(0,s_zero)
    real(kínd(le0)) d(n), b(n), \overline{d}t(2*n,k), \overline{c}t(2*n,k), perf_ratio, &
        b_t(2*n,k), y_t (2*n,k), eval_t(k), res(n,k), temp
    logical small
! This flag is used to get the k largest eigenvalues.
    small = .false.
! Generate the main diagonal and the co-diagonal of the
! tridiagonal matrix.
    call rand gen (b)
    call rand_gen (d)
    A(1,1:)=b; A (2,1:)=d
! Use Numerical Libraries routine for the calculation of k
! largest eigenvalues.
    CALL EVASB (N, K, A, LDA, NCODA, SMALL, EVAL)
    EVAL_T = EVAL
! Use DNFL tridiagonal solver for inverse iteration
! calculation of eigenvectors.
    factorization_choice: do nopt=0,1
! Create k tridiagonal problems, one for each inverse
! iteration system.
    b_t(1:n,1:k) = spread(b,DIM=2,NCOPIES=k)
    c t(1:n,1:k) = EOSHIFT(b t(1:n,1:k),SHIFT=1,DIM=1)
    d_t (1:n,1:k) = spread(d,\overline{D}IM=2,NCOPIES=k) - &
```

```
    spread(EVAL_T,DIM=1,NCOPIES=n)
! Start the right-hand side at random values, scaled downward
! to account for the expected 'blowup' in the solution.
    do i=1, k
        call rand gen (y t(1:n,i))
    end do
! Do two iterations for the eigenvectors.
    do i=1, 2
        y_t(1:n,1:k) = y_t(1:n,1:k)*epsilon(s_one)
        c\overline{all lin_sol_tri`}(c_t, d_t, b_t, y_t, \overline{&}
            iopt=iopt)
            iopt(nopt+1) = s_options(s_lin_sol_tri_solve_only,s_zero)
        end do
! Orthogonalize the eigenvectors. (This is the most
! intensive part of the computing.)
    do j=1,k-1 ! Forward sweep of HMGS orthogonalization.
            temp=s_one/sqrt(sum(y_t(1:n,j)**2))
            y_t(1:n,j)=y_t(1:n,j) *temp
            y_t(1:n,j+1:k)=y_t(1:n,j+1:k)+ &
            spread(-matmul(y_t(1:n,j),y_t(1:n,j+1:k)), &
        DIM=1,NCOPIES=n) * sp
        end do
        temp=s one/sqrt(sum(y_t(1:n,k)**2))
        y_t (1:\overline{n},k)=y_t (1:n,k) *temp
        do j=k-1,1,-1 ! Backward sweep of HMGS.
            y_t(1:n,j+1:k)=y_t(1:n,j+1:k) + &
            spread(-matmul(y_t(1:n,j),y_t(1:n,j+1:k)), &
        DIM=1,NCOPIES=n) * spread(y_t (1:n,j),DIM=2,NCOPIES=k-j)
        end do
! See if the performance ratio is smaller than the value one.
! If it is not the code will re-solve the systems using Gaussian
! Elimination. This is an exceptional event. It is a necessary
! complication for achieving reliable results.
```

```
    \(\operatorname{res}(1: n, 1: k)=\operatorname{spread}(d, D I M=2, N C O P I E S=k) * y t(1: n, 1: k)+\&\)
```

    \(\operatorname{res}(1: n, 1: k)=\operatorname{spread}(d, D I M=2, N C O P I E S=k) * y t(1: n, 1: k)+\&\)
    spread (b, DIM=2, NCOPIES=k) * \&
    spread (b, DIM=2, NCOPIES=k) * \&
    \(\operatorname{EOSHIFT}\left(y \_t(1: n, 1: k), \operatorname{SHIFT}=-1, D I M=1\right)+\&\)
    \(\operatorname{EOSHIFT}\left(y \_t(1: n, 1: k), \operatorname{SHIFT}=-1, D I M=1\right)+\&\)
    EOSHIFT (spread (b, DIM=2, NCOPIES=k) *y t (1:n, 1:k), SHIFT=1) \&
    EOSHIFT (spread (b, DIM=2, NCOPIES=k) *y t (1:n, 1:k), SHIFT=1) \&
    -y_t (1:n, 1:k) *spread (EVAL_T (1:k), DIM=1,NCOPIES=n)
    -y_t (1:n, 1:k) *spread (EVAL_T (1:k), DIM=1,NCOPIES=n)
    ! If the factorization method is Cyclic Reduction and perf ratio is
! larger than one, re-solve using Gaussian Elimination. If the
! method is already Gaussian Elimination, the loop exits
! and perf_ratio is checked at the end.
perf_ratio = sum(abs(res(1:n,1:k))) / \&

```
```

end do factorization_choice
if (perf_ratio <= s_one) then
write (*,*) 'Example 3 for LIN_SOL_TRI is correct.'
end if
end

```

\section*{Output}

Example 3 for LIN_SOL_TRI is correct.

\section*{Example 4: Tridiagonal Matrix Solving within Diffusion Equations}

The normalized partial differential equation
\[
u_{t} \equiv \frac{\partial u}{\partial t}=\frac{\partial^{2} u}{\partial x^{2}} \equiv u_{x x}
\]
is solved for values of \(0 \leq x \leq \pi\) and \(t>0\). A boundary value problem consists of choosing the value
\[
u(0, t)=u_{0}
\]
such that the equation
\[
u\left(x_{1}, t_{1}\right)=u_{1}
\]
is satisfied. Arbitrary values
\[
x_{1}=\frac{\pi}{2}, u_{1}=\frac{1}{2}
\]
and
\[
t_{1}=1
\]
are used for illustration of the solution process. The one-parameter equation
\[
u\left(x_{1}, t_{1}\right)-u_{1}=0
\]

The variables are changed to
\[
v(x, t)=u(x, t)-u_{0}
\]
that \(v(0, t)=0\). The function \(v(x, t)\) satisfies the differential equation. The one-parameter equation solved is therefore
\[
v\left(x_{1}, t_{1}\right)-\left(u_{1}-u_{0}\right)=0
\]

To solve this equation for \(u_{0}\), use the standard technique of the variational equation,
\[
w \equiv \frac{\partial v}{\partial u_{0}}
\]

Thus
\[
\frac{\partial w}{\partial t}=\frac{\partial^{2} w}{\partial x^{2}}
\]

Since the initial data for
\[
v(x, 0)=-u_{0}
\]
the variational equation initial condition is
\[
w(x, 0)=-1
\]

This model problem illustrates the method of lines and Galerkin principle implemented with the differential-algebraic solver, D2SPG (IMSL 2003, pp. 889-911). We use the integrator in "reverse communication" mode for evaluating the required functions, derivatives, and solving linear algebraic equations. See Example 4 of routine DASPG for a problem that uses reverse communication. Next see Example 4 of routine IVPAG for the development of the piecewiselinear Galerkin discretization method to solve the differential equation. This present example extends parts of both previous examples and illustrates Fortran 90 constructs. It further illustrates how a user can deal with a defect of an integrator that normally functions using only dense linear algebra factorization methods for solving the corrector equations. See the comments in Brenan et al. (1989, esp. p. 137). Also, see operator_ex20, supplied with the product examples.
```

    use lin_sol_tri_int
    use ran\overline{d}ge\overline{n}in\overline{t}
    use Numerical_Libraries
    implicit none
    ! This is Example 4 for LIN_SOL_TRI.
integer, parameter :: n=1000, ichap=5, iget=1, iput=2, \&
inum=6, irnum=7
real(kind(1e0)), parameter :: zero=0e0, one = 1e0
integer i, ido, in(50), inr(20), iopt(6), ival(7), \&
iwk(35+n)
real(kind(le0)) hx, pi_value, t, u_0, u_1, atol, rtol, sval(2), \&
tend, wk(41+11*n), y(n), ypr(\overline{n}), a_diag(n), \&
a_off(n), r_diag(n), r_off(n), t_y(n), t_ypr(n), \&
t_g(n), t_diag(2*n,1), t_upper (2*n,1), \&
t_lower (2*}n,1), t_sol (2*\overline{n},1
type(s_optiōns) :: iopti(2)=s_options(0, zero)
character(2) :: pi(1) = 'pi'
! Define initial data.
t = 0.0e0
u_0 = 1
u_1 = 0.5
tend = one

```
```

! Initial values for the variational equation.
y = -one; ypr= zero
pi value = const(pi)
hx = pi_value/(n+1)
a_diag = 2*hx/3
a_off = hx/6
r_diag = -2/hx
r_off = 1/hx
! Get integer option numbers.
iopt(1) = inum
call iumag ('math', ichap, iget, 1, iopt, in)
! Get floating point option numbers.
iopt(1) = irnum
call iumag ('math', ichap, iget, 1, iopt, inr)
! Set for reverse communication evaluation of the DAE.
iopt(1) = in(26)
ival(1) = 0
! Set for use of explicit partial derivatives.
iopt(2) = in(5)
ival(2) = 1
! Set for reverse communication evaluation of partials.
iopt(3) = in(29)
ival(3) = 0
! Set for reverse communication solution of linear equations.
iopt(4) = in(31)
ival(4) = 0
! Storage for the partial derivative array are not allocated or
! required in the integrator.
iopt(5) = in(34)
ival(5) = 1
! Set the sizes of iwk, wk for internal checking.
iopt(6) = in(35)
ival(6) = 35 + n
ival(7) = 41 + 11*n
! Set integer options:
call iumag ('math', ichap, iput, 6, iopt, ival)
! Reset tolerances for integrator:
atol = 1e-3; rtol= 1e-3
sval(1) = atol; sval(2) = rtol
iopt(1) = inr(5)
! Set floating point options:
call sumag ('math', ichap, iput, 1, iopt, sval)
! Integrate ODE/DAE. Use dummy external names for g(y, y')
! and partials.
ido = 1
Integration_Loop: do
call d2spg (n, t, tend, ido, y, ypr, dgspg, djspg, iwk, wk)
! Find where g(y, y') goes. (It only goes in one place here, but can
! vary where divided differences are used for partial derivatives.)
iopt(1) = in(27)

```
```

    call iumag ('math', ichap, iget, 1, iopt, ival)
    ! Direct user response:
select case(ido)
case (1,4)
! This should not occur.
write (*,*) ' Unexpected return with ido = ', ido
stop
case(3)
! Reset options to defaults. (This is good housekeeping but not
! required for this problem.)
in = -in
call iumag ('math', ichap, iput, 50, in, ival)
inr = -inr
call sumag ('math', ichap, iput, 20, inr, sval)
exit Integration_Loop
case(5)
! Evaluate partials of g(y,y').
t_y = y; t_ypr = ypr
t_g = r_diag*t_y + r_off*EOSHIFT(t_y,SHIFT=+1) \&
+ EO\overline{SHIFT(r_off*t_y,SHIFT=-1) \&}
- (a_diag*t_ypr + a_off*\overline{EOSHIFT}(t_ypr,SHIFT=+1) \&
+ E\overline{OSHIFT(a_off*t_ypr,SHIFT=-1))}
! Move data from the assumed size to assumed shape arrays.
do i=1, n
wk(ival(1)+i-1) = t_g(i)
end do
cycle Integration_Loop
case(6)
! Evaluate partials of g(y,y').
! Get value of c_j for partials.
iopt(1) = inr(9)
call sumag ('math', ichap, iget, 1, iopt, sval)
! Subtract c_j from diagonals to compute (partials for y')*c_j.
! The linear system is tridiagonal.
t_diag(1:n,1) = r_diag - sval(1)*a_diag
t_upper(1:n,1) = \overline{r_off - sval(1)*a_off}
t_lower = EOSHIFT(t_upper,SHIFT=+1,
cycle Integration_Loop
case (7)
! Compute the factorization.
iopti(1) = s_options(s_lin_sol_tri_factor_only,zero)
call lin_sol_tri (t_upper, t_diag, t_lower, \&
\overline{t_sol}, iopt=iopti)
cycle Integration_Loop
case (8)
! Solve the system.
iopti(1) = s_options(s_lin_sol_tri_solve_only,zero)

```
```

! Move data from the assumed size to assumed shape arrays.
t_sol(1:n,1)=wk(ival(1):ival(1) +n-1)
call lin_sol_tri (t_upper, t_diag, t_lower, \&
t_sol, iopt=iopti)
! Move data from the assumed shape to assumed size arrays.
wk(ival(1):ival(1)+n-1)=t_sol(1:n,1)
cycle Integration_Loop
case(2)
! Correct initial value to reach u 1 at t=tend.
u_0 = u_0 - (u_0*y(n/2) = (u_1-u_0)) / (y(n/2) + 1)
! Finish up internally in the integrator.
ido = 3
cycle Integration_Loop
end select
end do Integration_Loop
write (*,*) 'The equation u_t = u_xx, with u(0,t) = ', u_0
write (*,*) 'reaches the value ',\overline{u_1, ' at time = ', ten\overline{d}, '.'}
write (*,*) 'Example 4 for LIN_SOL_TRI is correct.'
end

```

\section*{Output}

Example 4 for LIN_SOL_TRI is correct.

\section*{LIN_SVD}

Computes the singular value decomposition (SVD) of a rectangular matrix, \(A\). This gives the decomposition
\[
A=U S V^{T}
\]
where \(V\) is an \(n \times n\) orthogonal matrix, \(U\) is an \(m \times m\) orthogonal matrix, and \(S\) is a real, rectangular diagonal matrix.

\section*{Required Arguments}

A- Array of size \(m \times n\) containing the matrix. (Input [/Output])
If the packaged option lin_svd_overwrite_input is used, this array is not saved on output.
\(\boldsymbol{S}\) - Array of size \(\min (m, n)\) containing the real singular values. These nonnegative values are in non-increasing order. (Output)
\(\boldsymbol{U}\) - Array of size \(m \times m\) containing the singular vectors, \(U\). (Output)
\(\boldsymbol{V}\) - Array of size \(n \times n\) containing the singular vectors, \(V\). (Output)

\section*{Optional Arguments}

MROWS \(=m\) (Input)
Uses array \(A(1: m, 1: n)\) for the input matrix.
Default: \(m=\operatorname{size}(A, 1)\)

NCOLS \(=\mathrm{n}\) (Input)
Uses array \(A(1: m, 1: n)\) for the input matrix.
Default: \(\mathrm{n}=\operatorname{size}(\mathrm{A}, 2)\)
RANK \(=k\) (Output)
Number of singular values that exceed the value Small. RANK will satisfy
\(k<=\min (m, n)\).
iopt = iopt(:) (Input)
Derived type array with the same precision as the input matrix. Used for passing optional data to the routine. The options are as follows:
\begin{tabular}{|c|c|c|}
\hline \multicolumn{3}{|c|}{Packaged Options for LIN_SVD} \\
\hline Option Prefix = ? & Option Name & Option Value \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_set_small & 1 \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_overwrite_input & 2 \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_scan_for_NaN & 3 \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_use_qr & 4 \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_skip_orth & 5 \\
\hline \(S_{-}, d_{-}, c_{-}, z_{-}\) & lin_svd_use_gauss_elim & 6 \\
\hline \(S_{-}, d_{-}, C_{-}, z_{-}\) & lin_svd_set_perf_ratio & 7 \\
\hline
\end{tabular}

\section*{iopt(IO) = ?_options(?_lin_svd_set_small, Small)}

If a singular value is smaller than \(\overline{\text { Small }} \overline{,}\) it is defined as zero for the purpose of computing the rank of \(A\).
Default: the smallest number that can be reciprocated safely
iopt(IO) = ?_options(?_lin_svd_overwrite_input, ?_dummy)
Does not save the input array \(A(:,:)\).
iopt (IO) = ?_options(?_lin_svd_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
\[
\text { isNaN }(a(i, j))==\text {.true. }
\]

See the isNaN () function, Chapter 10.
Default: The array is not scanned for NaNs.
iopt(IO) = ?_options(?_lin_svd_use_qr, ?_dummy)
Uses a rational \(Q R\) algorithm to compute eigenvalues. Accumulate the singular vectors using this algorithm.
Default: singular vectors computed using inverse iteration
```

iopt(IO) = ?_options(?_lin_svd_skip_Orth, ?_dummy)

```

If the eigenvalues are computed using inverse iteration, skips the final
orthogonalization of the vectors. This method results in a more efficient computation.
However, the singular vectors, while a complete set, may not be orthogonal.
Default: singular vectors are orthogonalized if obtained using inverse iteration
iopt(IO) = ?_options(?_lin_svd_use_gauss_elim, ?_dummy)
If the eigenvalues are computed using inverse iteration, uses standard elimination with partial pivoting to solve the inverse iteration problems.
Default: singular vectors computed using cyclic reduction
```

iopt(IO) = ?_options(?_lin_svd_set_perf_ratio,perf_ratio)

```

Uses residuals for approximate normalized singular vectors if they have a performance index no larger than perf_ratio. Otherwise an alternate approach is taken and the singular vectors are computed again: Standard elimination is used instead of cyclic reduction, or the standard \(Q R\) algorithm is used as a backup procedure to inverse iteration. Larger values of perf_ratio are less likely to cause these exceptions.
Default: perf_ratio \(=4\)

\section*{FORTRAN 90 Interface}

Generic: CALL LIN_SVD (A, S, U, V [, ...])
Specific: The specific interface names are S_LIN_SVD, D_LIN_SVD, C_LIN_SVD, and Z_LIN_SVD.

\section*{Description}

Routine lin_svd is an implementation of the \(Q R\) algorithm for computing the SVD of rectangular matrices. An orthogonal reduction of the input matrix to upper bidiagonal form is performed. Then, the SVD of a real bidiagonal matrix is calculated. The orthogonal decomposition \(A V=U S\) results from products of intermediate matrix factors. See Golub and Van Loan (1989, Chapter 8) for details.

\section*{Fatal, Terminal, and Warning Error Messages}

See the messages.gls file for error messages for LIN_SVD. These error messages are numbered 1001-1010; 1021-1030; 1041-1050; 1061-1070.

\section*{Example 1: Computing the SVD}

The SVD of a square, random matrix \(A\) is computed. The residuals \(R=A V-U S\) are small with respect to working precision. Also, see operator_ex21, supplied with the product examples.
```

        use lin_svd_int
        use ran\overline{d}_ge\overline{n}_int
        implicit none
    ! This is Example 1 for LIN_SVD.
integer, parameter :: n=32
real(kind(1d0)), parameter :: one=1d0
real(kind(1d0)) err
real(kind(1d0)), dimension(n,n) :: A, U, V, S(n), y(n* n)
! Generate a random n by n matrix.
call rand_gen(y)
A = reshape(y,(/n,n/))
! Compute the singular value decomposition.
call lin_svd(A, S, U, V)
! Check for small residuals of the expression A*V - U*S.
err = sum(abs(matmul(A,V) - U*spread(S,dim=1,ncopies=n))) \&
/ sum(abs(S))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 1 for LIN_SVD is correct.'
end if
end

```

\section*{Output}

Example 1 for LIN_SVD is correct.

\section*{Additional Examples}

\section*{Example 2: Linear Least Squares with a Quadratic Constraint}

An \(m \times n\) matrix equation \(A x \cong b, m>n\), is approximated in a least-squares sense. The matrix \(b\) is size \(m \times k\). Each of the \(k\) solution vectors of the matrix \(x\) is constrained to have Euclidean length of value \(\alpha_{j}>0\). The value of \(\alpha_{i}\) is chosen so that the constrained solution is 0.25 the length of the nonregularized or standard least-squares equation. See Golub and Van Loan (1989, Chapter 12) for more details. In the Example 2 code, Newton's method is used to solve for each regularizing parameter of the \(k\) systems. The solution is then computed and its length is checked. Also, see operator_ex22, supplied with the product examples.
```

use lin svd int
use ran\overline{d}_ge\overline{n}_int
implicit none

```
```

! This is Example 2 for LIN_SVD.
integer, parameter :: m=64, n=32, k=4
real(kind(1d0)), parameter :: one=1d0, zero=0d0
real(kind(1d0)) a (m,n), s(n), u(m,m), v(n,n), y(m*max(n,k)), \&
b(m,k), x(n,k), g(m,k), alpha(k), lamda(k), \&
delta_lamda(k), t_g(n,k), s_sq(n), phi(n,k), \&
phi_dot(n,k), rand(k), err
! Generate a random matrix for both A and B.
call rand gen(y)
a = reshape(y,(/m,n/))
call rand_gen(y)
b = reshape(y,(/m,k/))
! Compute the singular value decomposition.
call lin_svd(a, s, u, v)
! Choose alpha so that the lengths of the regularized solutions
! are 0.25 times lengths of the non-regularized solutions.
g = matmul(transpose(u),b)
x = matmul(v,spread(one/s,dim=2,ncopies=k)*g(1:n,1:k))
alpha = 0.25*sqrt (sum (x**2,dim=1))
t_g = g(1:n,1:k)*spread(s,dim=2,ncopies=k)
s_sq = s**2; lamda = zero
solve for lamda: do
x=ōne/(spread(s_sq,dim=2,ncopies=k)+ \&
spread(lamda,dim=1,ncopies=n))
phi = (t_g*x)**2; phi_dot = -2*phi*x
delta_lam}da=(sum(ph\overline{i},dim=1)-alpha**2)/sum(phi_dot,dim=1
! Make Newton method correction to solve the secular equations for
! lamda.
lamda = lamda - delta_lamda
if (sum(abs(delta_lamda)) <= \&
sqrt(epsilon(one))*sum(lamda)) \&
exit solve_for_lamda
! This is intended to fix up negative solution approximations.
call rand_gen(rand)
where (lamda < 0) lamda = s(1) * rand
end do solve_for_lamda
! Compute solutions and check lengths.
x = matmul(v,t_g/(spread(s_sq,dim=2,ncopies=k)+ \&
spread(lam}da,dim=1,ncopies=n)))
err = sum(abs(sum(x**2,dim=1) - alpha**2))/sum(abs(alpha**2))
if (err <= sqrt(epsilon(one))) then

```
```

    write (*,*) 'Example 2 for LIN_SVD is correct.'
    end if
end

```

\section*{Output}

Example 2 for LIN_SVD is correct.

\section*{Example 3: Generalized Singular Value Decomposition}

The \(n \times n\) matrices \(A\) and \(B\) are expanded in a Generalized Singular Value Decomposition (GSVD). Two \(n \times n\) orthogonal matrices, \(U\) and \(V\), and a nonsingular matrix \(X\) are computed such that
\[
A X=\operatorname{Udiag}\left(c_{1}, \ldots, c_{n}\right)
\]
and
\[
B X=\operatorname{Vdiag}\left(s_{1}, \ldots, s_{n}\right)
\]

The values \(s_{i}\) and \(c_{i i}\) are normalized so that
\[
s_{i}^{2}+c_{i}^{2}=1
\]

The \(c_{i}\) are nonincreasing, and the \(s_{i}\) are nondecreasing. See Golub and Van Loan (1989, Chapter 8) for more details. Our method is based on computing three SVDs as opposed to the \(Q R\) decomposition and two SVDs outlined in Golub and Van Loan. As a bonus, an SVD of the matrix \(X\) is obtained, and you can use this information to answer further questions about its conditioning. This form of the decomposition assumes that the matrix
\[
D=\left[\begin{array}{l}
A \\
B
\end{array}\right]
\]
has all its singular values strictly positive. For alternate problems, where some singular values of \(D\) are zero, the GSVD becomes
\[
U^{T} A=\operatorname{diag}\left(c_{1}, \ldots, c_{n}\right) W
\]
and
\[
V^{T} B=\operatorname{diag}\left(s_{1}, \ldots, s_{n}\right) W
\]

The matrix \(W\) has the same singular values as the matrix \(D\). Also, see operator_ex23, supplied with the product examples.
```

use lin_svd_int
use rand_gen_int
implicit none
! This is Example 3 for LIN_SVD.

```
```

        integer, parameter :: n=32
        integer i
        real(kind(1d0)), parameter :: one=1.0d0
        real(kind(1d0)) a(n,n), b(n,n), d(2*n,n), x(n,n), u_d(2*n,2*n), &
        v_d(n,n), v_c(n,n), u_c(n,n), v_s(n,n), u_s(n,n), &
        y(n*n), s_d\overline{(n), c(n),- s(n), sc_\overline{c}(n), sc_s(n), &}
        err1, err2
    ! Generate random square matrices for both A and B.
call rand_gen(y)
a = reshape(y,(/n,n/))
call rand_gen(y)
b = reshape(y,(/n,n/))
! Construct D; A is on the top; B is on the bottom.
d(1:n,1:n) = a
d(n+1:2*n,1:n) = b
! Compute the singular value decompositions used for the GSVD.
call lin_svd(d, s_d, u_d, v_d)
call lin_svd(u_d(\overline{1}:n,1:n), \overline{c, u_c, v_c)}
call lin_svd(u_d(n+1:,1:n), s, u__s, v
! Rearrange c(:) so it is non-increasing. Move singular
! vectors accordingly. (The use of temporary objects sc_c and
! x is required.)
sc_c = c(n:1:-1); c = sc_c
x = u_c(1:n,n:1:-1); u_c = x
x = v_c(1:n,n:1:-1); v_c = x
! The columns of v_c and v_s have the same span. They are
! equivalent by taking the signs of the largest magnitude values
! positive.
do i=1, n
sc_c(i) = sign(one,v_c(sum(maxloc(abs(v_c(1:n,i)))),i))
sc_s(i) = sign(one,v_s(sum(maxloc(abs(v_s(1:n,i)))),i))
end do
v_c = v_c*spread(sc_c,dim=1,ncopies=n)
u_c = u_c*spread(sc_c,dim=1,ncopies=n)
v_s = v_s*spread(sc_s,dim=1,ncopies=n)
u_s = u__s*spread(sc_s,dim=1,ncopies=n)
! In this form of the GSVD, the matrix X can be unstable if D
! is ill-conditioned.
x = matmul(v_d*spread(one/s_d,dim=1,ncopies=n),v_c)
! Check residuals for GSVD, A*X = u_c*diag(c_1, ..., c_n), and

```
```

! B*X = u_s*diag(s_1, ..., s_n).
err\overline{1}=\operatorname{sum(abss(matmul(\overline{a},x) - u_c*spread(c,dim=1,ncopies=n))) \&}
/ sum(s d)
err2 = sum(abs(matmul(b,x) - u_s*spread(s,dim=1,ncopies=n))) \&
/ sum(s_d)
if (err1 <= sqrE}(epsilon(one)) .and. \&
err2 <= sqrt(epsilon(one))) then
write (*,*) 'Example 3 for LIN_SVD is correct.'
end if
end

```

\section*{Example 4: Ridge Regression as Cross-Validation with Weighting}

This example illustrates a particular choice for the ridge regression problem: The least-squares problem \(A x \cong b\) is modified by the addition of a regularizing term to become
\[
\min _{x}\left(\|A x-b\|_{2}^{2}+\lambda^{2}\|x\|_{2}^{2}\right)
\]

The solution to this problem, with row \(k\) deleted, is denoted by \(x_{k}(\lambda)\). Using nonnegative weights ( \(w_{1}, \ldots, w_{m}\) ), the cross-validation squared error \(C(\lambda)\) is given by:
\[
m C(\lambda)=\sum_{k=1}^{m} w_{k}\left(a_{k}^{T} x_{k}(\lambda)-b_{k}\right)^{2}
\]

With the SVD \(A=U S V^{T}\) and product \(g=U^{T} b\), this quantity can be written as
\[
m C(\lambda)=\sum_{k=1}^{m} w_{k}\left(\frac{\left(b_{k}-\sum_{j=1}^{n} u_{k j} g_{j} \frac{s_{j}^{2}}{\left(s_{j}^{2}+\lambda^{2}\right)}\right)}{\left(1-\sum_{j=1}^{n} u_{k j}^{2} \frac{s_{j}^{2}}{\left(s_{j}^{2}+\lambda^{2}\right)}\right)}\right)^{2}
\]

This expression is minimized. See Golub and Van Loan (1989, Chapter 12) for more details. In the Example 4 code, \(m C(\lambda)\), at \(p=10\) grid points are evaluated using a log-scale with respect to \(\lambda\), \(0.1 s_{1} \leq \lambda \leq 10 s_{1}\). Array operations and intrinsics are used to evaluate the function and then to choose an approximate minimum. Following the computation of the optimum \(\lambda\), the regularized solutions are computed. Also, see operator_ex 24 , supplied with the product examples.
```

    use lin_svd_int
    use ran\overline{d_gen__int}
    implicit none
    ! This is Example 4 for LIN_SVD.
integer i
integer, parameter :: m=32, n=16, p=10, k=4

```
```

    real(kind(1d0)), parameter :: one=1d0
    real(kind(1d0)) log_lamda, log_lamda_t, delta_log_lamda
    real(kind(ld0)) a (m,n), b (m,k), w (m, \overline{k}), g(m,k), t\overline{(n), s(n), &}
        s_sq(n), u(m,m), v(n,n), y(m*max (n,k)), &
        c_lamda(p,k), lamda(k), x(n,k), res(n,k)
    ! Generate random rectangular matrices for A and right-hand
! sides, b.
call rand_gen(y)
a = reshape(y,(/m,n/))
call rand_gen(y)
b = reshape(y,(/m,k/))
! Generate random weights for each of the right-hand sides.
call rand_gen(y)
w = reshape(y,(/m,k/))
! Compute the singular value decomposition.
call lin_svd(a, s, u, v)
g = matmul(transpose(u),b)
s_sq = s**2
log_lamda = log(10.*s(1)); log_lamda_t=log_lamda
delta_log_lamda = (log_lamda - log(0.1*s(n))) / (p-1)
! Choose lamda to minimize the "cross-validation" weighted
! square error. First evaluate the error at a grid of points,
! uniform in log_scale.
cross_validation_error: do i=1, p
t = s_sq/(s_sq
c_lamda(i,:) = sum(w*((b-matmul(u(1:m,1:n),g(1:n,1:k)* \&
spread(t,DIM=2,NCOPIES=k)))/ \&
(one-matmul(u (1:m,1:n)**2, \&
spread(t,DIM=2,NCOPIES=k))))**2,DIM=1)
log_lamda = log_lamda - delta_log_lamda
end do cross_validàtion_error
! Compute the grid value and lamda corresponding to the minimum.
do i=1, k
lamda(i) = exp(log_lamda_t - delta_log_lamda* \&
(sum}(minloc(c_\overline{l}\operatorname{amda}(1:p,i)))-1)
end do
! Compute the solution using the optimum "cross-validation"
! parameter.
x = matmul(v,g(1:n,1:k)*spread(s,DIM=2,NCOPIES=k) / \&
(spread(s_sq,DIM=2,NCOPIES=k) + \&
spread(lamda,DIM=1,NCOPIES=n)))
! Check the residuals, using normal equations.
res = matmul(transpose(a),b-matmul(a,x)) - \&
spread(lamda,DIM=1,NCOPIES=n)*x
if (sum(abs(res))/sum(s_sq) <= \&

```
```

        sqrt(epsilon(one))) then
    write (*,*) 'Example 4 for LIN_SVD is correct.'
    end if
end

```

\section*{Output}

Example 4 for LIN_SVD is correct.

\section*{Parallel Constrained Least-Squares Solvers}

\section*{Solving Constrained Least-Squares Systems}

The routine PARALLEL_NONNEGATIVE_LSQ is used to solve dense least-squares systems. These are represented by \(A x \cong b\) where \(A\) is an \(m \times n\) coefficient data matrix, \(b\) is a given right-hand side \(m\)-vector, and \(x\) is the solution \(n\)-vector being computed. Further, there is a constraint requirement, \(x \geq 0\). The routine PARALLEL_BOUNDED_LSQ is used when the problem has lower and upper bounds for the solution, \(\alpha \leq x \leq \beta\). By making the bounds large, individual constraints can be eliminated. There are no restrictions on the relative sizes of \(m\) and \(n\). When \(n\) is large, these codes can substantially reduce computer time and storage requirements, compared with using a routine for solving a constrained system and a single processor.

The user provides the matrix partitioned by blocks of columns: \(A=\left[A_{1}\left|A_{2}\right| \ldots \mid A_{k}\right]\). An individual block of the partitioned matrix, say \(A_{p}\), is located entirely on the processor with rank MP_RANK \(=p-1\), where MP_RANK is packaged in the module MPI_SETUP_INT. This module, and the function MP_SETUP (), define the Fortran Library MPI communicator, MP_LIBRARY_WORLD. See Chapter 10, Dense Matrix Parallelism Using MPI.

\section*{PARALLEL_NONNEGATIVE_LSQ \\ HMPI \\ REQUIRED}

For a detailed description of MPI Requirements see "Dense Matrix Parallelism Using MPI" in Chapter 10 of this manual.

Solves a linear, non-negative constrained least-squares system.

\section*{Usage Notes}

CALL PARALLEL_NONNEGATIVE_LSQ \&
(A, B, X, RNORM, W, INDEX, IPART, IOPT = IOPT)

\section*{Required Arguments}
\(\boldsymbol{A}(\mathbf{1 : M , : )}\) - (Input/Output) Columns of the matrix with limits given by entries in the array \(\operatorname{IPART}\left(1: 2,1: \max \left(1, \mathrm{MP} \_\right.\right.\)NPROCS \(\left.)\right)\). On output \(A_{k}\) is replaced by the product \(Q A_{k}\), where \(Q\) is an orthogonal matrix. The value SIZE \((A, 1)\) defines the value of M. Each processor starts and exits with its piece of the partitioned matrix.
\(\boldsymbol{B} \mathbf{1}: \mathbf{M})\) - (Input/Output) Assumed-size array of length M containing the right-hand side vector, \(b\). On output \(b\) is replaced by the product \(Q b\), where \(Q\) is the orthogonal matrix applied to \(A\). All processors in the communicator start and exit with the same vector.
\(\boldsymbol{X}(\mathbf{1}: N)\) - (Output) Assumed-size array of length N containing the solution, \(x \geq 0\). The value SIZE (X) defines the value of \(N\). All processors exit with the same vector.

RNORM - (Output) Scalar that contains the Euclidean or least-squares length of the residual vector, \(\|A x-b\|\). All processors exit with the same value.
\(W(1: N)\) - (Output) Assumed-size array of length N containing the dual vector, \(w=A^{T}(b-A x) \leq 0\). All processors exit with the same vector.
\(\operatorname{INDEX}(1: N)\) - (Output) Assumed-size array of length \(N\) containing the NSETP indices of columns in the positive solution, and the remainder that are at their constraint. The number of positive components in the solution \(x\) is give by the Fortran intrinsic function value, NSETP=COUNT ( \(\mathrm{X}>0\) ) . All processors exit with the same array.

IPART(1:2,1:max(1,MP_NPROCS)) - (Input) Assumed-size array containing the partitioning describing the matrix \(A\). The value MP_NPROCS is the number of processors in the communicator, except when MPI has been finalized with a call to the routine MP_SETUP ('Final'). This causes MP_NPROCS to be assigned 0 . Normally users will give the partitioning to processor of rank \(=\) MP_RANK by setting IPART \(\left(1, M P \_\right.\)RANK +1\()=\) first column index, and \(\operatorname{IPART}\left(2, \mathrm{MP} \_\right.\)RANK +1 ) = last column index. The number of columns per node is typically based on their relative computing power. To avoid a node with rank MP_RANK doing any work except communication, set IPART (1,MP_RANK+1) = 0 and \(\operatorname{IPART}\left(2, M P \_R A N K+1\right)=-1\). In this exceptional case there is no reference to the array \(\boldsymbol{A}(:,:)\) at that node.

\section*{Optional Argument}

IOPT(:)— (Input) Assumed-size array of derived type S_OPTIONS or D_OPTIONS. This argument is used to change internal parameters of the algorithm. Normally users will not be concerned about this argument, so they would not include it in the argument list for the routine.
\begin{tabular}{|l|l|}
\hline \multicolumn{2}{|c|}{ Packaged Options for PARALLEL_NONNEGATIVE_LSQ } \\
\hline Option Name & Option Value \\
\hline PNLSQ_SET_TOLERANCE & 1 \\
\hline PNLSQ_SET_MAX_ITERATIONS & 2 \\
\hline PNLSQ_SET_MIN_RESIDUAL & 3 \\
\hline
\end{tabular}

IOPT (IO) =?_OPTIONS (PNLSQ_SET_TOLERANCE, TOLERANCE) Replaces the default rank tolerance for using a column, from EPSILON(TOLERANCE) to TOLERANCE. Increasing the value of TOLERANCE will cause fewer columns to be moved from their constraints, and may cause the minimum residual RNORM to increase.
 for the minimum residual vector length from 0 to RESID. Increasing the value of RESID can result in fewer iterations and thus increased efficiency. The descent in the optimization will stop at the first point where the minimum residual RNORM is smaller than RESID. Using this option may result in the dual vector not satisfying its optimality conditions, as noted above.
\(\operatorname{IOPT}(I O)=\) PNLSQ_SET_MAX_ITERATIONS
\(\operatorname{IOPT}(I O+1)=\) NEW_MAX_ITERATIONS Replaces the default maximum number of iterations from \(3 * N\) to NEW_MAX_ITERATIONS. Note that this option requires two entries in the derived type array.

\section*{FORTRAN 90 Interface}

Generic: CALL PARALLEL_NONNEGATIVE_LSQ (A, B, X, RNORM, W, INDEX, IPART [,...])

Specific: The specific interface names are S_PARALLEL_NONNEGATIVE_LSQ and D_PARALLEL_NONNEGATIVE_LSQ.

\section*{Description}

Subroutine PARALLEL_NONNEGATIVE_LSQ solves the linear least-squares system \(A x \cong b, x \geq 0\), using the algorithm \(N N L S\) found in Lawson and Hanson, (1995), pages 160-161. The code now updates the dual vector \(w\) of Step 2, page 161. The remaining new steps involve exchange of required data, using MPI.

\section*{Example 1: Distributed Linear Inequality Constraint Solver}

The program PNLSQ_EX1 illustrates the computation of the minimum Euclidean length solution of an \(m^{\prime} \times n^{\prime}\) system of linear inequality constraints, \(G y \geq h\). The solution algorithm is based on Algorithm \(L D P\), page \(165-166\), loc. cit. The rows of \(E=[G: h]\) are partitioned and assigned
random values. When the minimum Euclidean length solution to the inequalities has been calculated, the residuals \(r=G y-h \geq 0\) are computed, with the dual variables to the \(N N L S\) problem indicating the entries of \(r\) that are precisely zero.
The fact that matrix products involving both \(E\) and \(E^{T}\) are needed to compute the constrained solution \(y\) and the residuals \(r\), implies that message passing is required. This occurs after the NNLS solution is computed.
```

    PROGRAM PNLSQ_EX1
    ! Use Parallel_nonnegative_LSQ to solve an inequality
! constraint problem, Gy >= h. This algorithm uses
! Algorithm LDP of Solving Least Squares Problems,
! page 165. The constraints are allocated to the
! processors, by rows, in columns of the array A(:,:).
USE PNLSQ_INT
USE MPI SETUP INT
USE RAND_INT
USE SHOW_INT
IMPLICIT NONE
INCLUDE "mpif.h"
INTEGER, PARAMETER :: MP=500, NP=400, M=NP+1, N=MP
REAL(KIND(1D0)), PARAMETER :: ZERO=0D0, ONE=1D0
REAL(KIND(1D0)), ALLOCATABLE :: \&
A(:,:), B(:), X(:), Y(:), W(:), ASAVE(:,:)
REAL(KIND(1D0)) RNORM
INTEGER, ALLOCATABLE :: INDEX(:), IPART(:,:)
INTEGER K, L, DN, J, JSHIFT, IERROR
LOGICAL :: PRINT=.false.
! Setup for MPI:
MP_NPROCS=MP_SETUP()
DN=N/max(1,max(1,MP NPROCS)) -1
ALLOCATE (IPART (2, ma\overline{x}(1,MP_NPROCS)))
! Spread constraint rows evenly to the processors.
IPART (1, 1)=1
DO L=2,MP NPROCS
IPART (\overline{2},L-1)=IPART (1,L-1)+DN
IPART (1, L) = IPART (2,L-1) +1
END DO
IPART (2,MP_NPROCS )=N
! Define the constraint data using random values.
K=max (0,IPART (2,MP_RANK+1)-IPART (1,MP_RANK+1)+1)
ALLOCATE (A (M,K), A\overline{SAVE (M,K), X(N), W(N), \&}
B (M), Y(M), INDEX(N))
! The use of ASAVE can be removed by regenerating
! the data for A(:,:) after the return from

```
```

! Parallel_nonnegative_LSQ.
A=rand(A); ASAVE=A
IF(MP RANK == 0 .and. PRINT) \&
CALI SHOW(IPART, \&
"Partition of the constraints to be solved")
! Set the right-hand side to be one in the last component, zero elsewhere.
B=ZERO;B (M) =ONE
! Solve the dual problem.
CALL Parallel nonnegative LSQ \&
(A, B, X, RNORM, W, INDEX, IPART)
! Each processor multiplies its block times the part of
! the dual corresponding to that part of the partition.
Y=ZERO
DO J=IPART(1,MP_RANK+1),IPART (2,MP_RANK+1)
JSHIFT=J-IPART (1,MP_RANK+1) +1
Y=Y+ASAVE (:,JSHIFT)}\mp@subsup{}{~}{\prime}X(J
END DO
! Accumulate the pieces from all the processors. Put sum into B(:)
! on rank O processor.
B=Y
IF(MP NPROCS > 1) \&
CALI MPI_REDUCE(Y, B, M, MPI_DOUBLE_PRECISION,\&
MPI_SUM,
IF(MP_\overline{RANK == 0) THEN}
! Compute constrained solution at the root.
! The constraints will have no solution if B(M) = ONE.
! All of these example problems have solutions.
B (M) = B (M) -ONE; B=-B/B (M)
END IF
! Send the inequality constraint solution to all nodes.
IF(MP_NPROCS > 1) \&
CALI MPI BCAST(B, M, MPI DOUBLE PRECISION, \&
0, MP_L\overline{I}BRARY_WORLD, IE\overline{R}ROR)
! For large problems this printing needs to be removed.
IF(MP_RANK == 0 .and. PRINT) \&
CALL SHOW(B(1:NP), \&
"Minimal length solution of the constraints")
! Compute residuals of the individual constraints.
! If only the solution is desired, the program ends here.
X=ZERO
DO J=IPART(1,MP RANK+1),IPART(2,MP RANK+1)
JSHIFT=J-IPART (1,MP_RANK+1) +1
X(J)=dot_product(B,ASAVE (:,JSHIFT))
END DO
! This cleans up residuals that are about rounding
! error unit (times) the size of the constraint

```
```

! equation and right-hand side. They are replaced
! by exact zero.
WHERE (W == ZERO) X=ZERO; W=X
! Each group of residuals is disjoint, per processor.
! We add all the pieces together for the total set of
! constraints.
IF(MP NPROCS > 1) \&
CAL\overline{L MPI_REDUCE (X, W, N, MPI_DOUBLE_PRECISION, \&}
MPI_SUM, 0, MP_LIBRARY_WORLD, IER\overline{ROR)}
IF (MP R\overline{A}NK == 0 .a\overline{n}d. PRINT}) \&
CAL\overline{L}}\mathrm{ SHOW(W, "Residuals for the constraints")
! See to any errors and shut down MPI.
MP_NPROCS=MP_SETUP('Final')
IF\overline{(MP RANK == 0) THEN}
IF(\overline{COUNT (W < ZERO) == 0) WRITE(*,*) \&}
" Example 1 for PARALLEL_NONNEGATIVE_LSQ is correct."
END IF
END

```

\section*{Output}

Example 1 for PARALLEL_NONNEGATIVE_LSQ is correct.

\section*{Additional Examples}

\section*{Example 2: Distributed Non-negative Least-Squares}

The program PNLSQ_EX2 illustrates the computation of the solution to a system of linear leastsquares equations with simple constraints: \(a_{i}^{T} x \cong b_{i}, i=1, \ldots, m\), subject to \(x \geq 0\). In this example we write the row vectors \(\left[a_{i}{ }^{T}: b_{i}\right]\) on a file. This illustrates reading the data by rows and arranging the data by columns, as required by PARALLEL_NONNEGATIVE_LSQ. After reading the data, the right-hand side vector is broadcast to the group before computing a solution, \(x\). The block-size is chosen so that each participating processor receives the same number of columns, except any remaining columns sent to the processor with largest rank. This processor contains the right-hand side before the broadcast.
This example illustrates connecting a \(B L A C S\) 'context' handle and the Fortran Library MPI communicator, MP_LIBRARY_WORLD, described in Chapter 10.
```

    PROGRAM PNLSQ_EX2
    ! Use Parallel_Nonnegative_LSQ to solve a least-squares
! problem, A x = b, with x >= 0. This algorithm uses a
! distributed version of NNLS, found in the book
! Solving Least Squares Problems, page 165. The data is
! read from a file, by rows, and sent to the processors,
! as array columns.
USE PNLSQ INT
USE SCALAPACK_IO_INT
USE BLACS_INT

```
```

    USE MPI_SETUP_INT
    USE RAND INT
    USE ERROR_OPTION_PACKET
    IMPLICIT NONE
    INCLUDE "mpif.h"
    INTEGER, PARAMETER :: M=128, N=32, NP=N+1, NIN=10
    real(kind(1dO)), ALLOCATABLE, DIMENSION(:) :: &
        d_A(:,:), A(:,:), B, C, W, X, Y
    real(kind(1dO)) RNORM, ERROR
    INTEGER, ALLOCATABLE :: INDEX(:), IPART(:,:)
    INTEGER I, J, K, L, DN, JSHIFT, IERROR, &
        CONTXT, NPROW, MYROW, MYCOL, DESC A(9)
    TYPE(d_OPTIONS) IOPT(1)
    ! Routines with the "BLACS_" prefix are from the
! BLACS library.
CALL BLACS_PINFO(MP_RANK, MP_NPROCS)
! Make initialization for BLACS.
CALL BLACS_GET (0,0, CONTXT)
! Define processor grid to be 1 by MP NPROCS.
NPROW=1
CALL BLACS_GRIDINIT(CONTXT, 'N/A', NPROW, MP_NPROCS)
! Get this processor's role in the process grid.
CALL BLACS_GRIDINFO(CONTXT, NPROW, MP_NPROCS, \&
MYROW, MȲCOL)
! Connect BLACS context with communicator MP_LIBRARY_WORLD.
CALL BLACS_GET (CONTXT, 10, MP_LIBRARY_WOR\overline{LD})
! Setup for MPI:
MP_NPROCS=MP _SETUP()
DN=max(1,NP/MP NPROCS)
ALLOCATE(IPART\overline{(2,MP_NPROCS))}
! Spread columns evenly to the processors. Any odd
! number of columns are in the processor with highest
! rank.
IPART (1,:)=1; IPART (2,:)=0
DO L=2,MP_NPROCS
IPART (2, L-1) = IPART (1,L-1) +DN
IPART (1, L) = IPART (2,L-1) +1
END DO
IPART (2,MP NPROCS) =NP
IPART (2,:) =min (NP, IPART (2,:))
! Note which processor (L-1) receives the right-hand side.

```
```

    DO L=1,MP_NPROCS
        IF(IPART (1,L) <= NP .and. NP <= IPART (2,L)) EXIT
    END DO
    K=max (0,IPART (2,MP_RANK+1) -IPART (1,MP_RANK+1) +1)
    ALLOCATE (d_A (M, K), W(N), X(N), Y(N),&
        B(M), C(M), INDEX(N))
    IF(MP RANK == 0 ) THEN
        ALLOCATE (A (M,N))
    ! Define the matrix data using random values.
A=rand(A); B=rand(B)
! Write the rows of data to an external file.
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='UNKNOWN')
DO I=1,M
WRITE(NIN,*) (A(I,J),J=1,N), B(I)
END DO
CLOSE (NIN)
ELSE
! No resources are used where this array is not saved.
ALLOCATE (A (M,0))
END IF
! Define the matrix descriptor. This includes the
! right-hand side as an additional column. The row
! block size, on each processor, is arbitrary, but is
! chosen here to match the column block size.
DESC_A=(/1, CONTXT, M, NP, DN+1, DN+1, 0, 0, M/)
! Read the data by rows.
IOPT (1)=ScaLAPACK READ BY ROWS
CALL ScaLAPACK_READ ("Atest.dat", DESC_A, \&
d_A, IOPT=IOP\overline{T}
! Broadcast the right-hand side to all processors.
JSHIFT=NP-IPART (1,L) +1
IF(K > 0) B=d_A(:,JSHIFT)
IF(MP NPROCS > 1) \&
CAL\overline{L MPI_BCAST(B, M, MPI_DOUBLE_PRECISION , L-1, \&}
MP_LIB\overline{R}ARY_WORLD, IERR\overline{O}R)
! Adjust the partition of columns to ignore the
! last column, which is the right-hand side. It is
! now moved to B(:).
IPART (2,:)=min(N,IPART (2,:))
! Solve the constrained distributed problem.
C=B
CALL Parallel_Nonnegative_LSQ \&
(d_A, B, X, R\overline{N}ORM, W, IND\overline{EX, IPART)}
! Solve the problem on one processor, with data saved
! for a cross-check.

```
```

    IPART (2,:)=0; IPART (2,1)=N; MP_NPROCS=1
    ! Since all processors execute this code, all arrays
! must be allocated in the main program.
CALL Parallel_Nonnegative_LSQ \&
(A, C, Y, RNO\overline{RM, W, INDEX, IPART)}
! See to any errors.
CALL e1pop("Mp_Setup")
! Check the differences in the two solutions. Unique solutions
! may differ in the last bits, due to rounding.
IF(MP_RANK == 0) THEN
ERRO
IF(ERROR <= sqrt(EPSILON(ERROR))) write(*,*) \&
' Example 2 for PARALLEL NONNEGATIVE LSQ is correct.'
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='OLD')
CLOSE(NIN, STATUS='Delete')
END IF
! Exit from using this process grid.
CALL BLACS_GRIDEXIT( CONTXT )
CALL BLACS_EXIT(0)
END

```

\section*{Output}

Example 2 for PARALLEL_NONNEGATIVE_LSQ is correct.'

\section*{PARALLEL_BOUNDED_LSQ}

For a detailed description of MPI Requirements see "Dense Matrix Parallelism Using MPI" in Chapter 10 of this manual.

Solves a linear least-squares system with bounds on the unknowns.

\section*{Usage Notes}
```

CALL PARALLEL_BOUNDED_LSQ \& (A, B, BND, X, RNORM, W, INDEX, IPART, \&
NSETP, NSETZ, IOPT=IOPT)

```

\section*{Required Arguments}
\(\boldsymbol{A}(\mathbf{1}: \mathbf{M},:\) ) - (Input/Output) Columns of the matrix with limits given by entries in the array \(\operatorname{IPART}\left(1: 2,1: \max \left(1, \mathrm{MP} \_\right.\right.\)NPROCS \(\left.)\right)\). On output \(A_{k}\) is replaced by the product \(Q A_{k}\),
where \(Q\) is an orthogonal matrix. The value \(\operatorname{SIZE}(A, 1)\) defines the value of M. Each processor starts and exits with its piece of the partitioned matrix.
\(\boldsymbol{B} \mathbf{( 1 : M ) ~ - ~ ( I n p u t / O u t p u t ) ~ A s s u m e d - s i z e ~ a r r a y ~ o f ~ l e n g t h ~} M\) containing the right-hand side vector, \(b\). On output \(b\) is replaced by the product \(Q(b-A g)\), where \(Q\) is the orthogonal matrix applied to \(A\) and \(g\) is a set of active bounds for the solution. All processors in the communicator start and exit with the same vector.

BND(1:2,1:N) - (Input) Assumed-size array containing the bounds for \(x\). The lower bound \(\alpha_{j}\) is in \(\operatorname{BND}(1, J)\), and the upper bound \(\beta_{j}\) is in \(\operatorname{BND}(2, J)\).
\(\boldsymbol{X}(\mathbf{1}: N)\) - (Output) Assumed-size array of length N containing the solution, \(\alpha \leq x \leq \beta\). The value SIZE (X) defines the value of \(N\). All processors exit with the same vector.

RNORM - (Output) Scalar that contains the Euclidean or least-squares length of the residual vector, \(\|A x-b\|\). All processors exit with the same value.
\(\mathbf{W}(\mathbf{1}: N)\) - (Output) Assumed-size array of length N containing the dual vector, \(w=A^{T}(b-A x)\). At a solution exactly one of the following is true for each \(j, 1 \leq j \leq n\),
\[
\begin{aligned}
& \text { - } \alpha_{j}=x_{j}=\beta_{j}, \text { and } w_{j} \text { arbitrary } \\
& \text { - } \alpha_{j}=x_{j}, \text { and } w_{j} \leq 0 \\
& \text { - } x_{j}=\beta_{j}, \text { and } w_{j} \geq 0 \\
& \text { - } \alpha_{j}<x_{j}<\beta_{j}, \text { and } w_{j}=0
\end{aligned}
\]

All processors exit with the same vector.
INDEX(1:N) - (Output) Assumed-size array of length n containing the NSETP indices of columns in the solution interior to bounds, and the remainder that are at a constraint. All processors exit with the same array.

IPART(1:2,1:max(1,MP_NPROCS)) - (Input) Assumed-size array containing the partitioning describing the matrix \(A\). The value MP_NPROCS is the number of processors in the communicator, except when MPI has been finalized with a call to the routine MP_SETUP ('Final'). This causes MP_NPROCS to be assigned 0. Normally users will give the partitioning to processor of rank \(=\mathrm{MP}\) _RANK by setting \(\operatorname{IPART}\left(1, M P \_R A N K+1\right)=\) first column index, and \(\operatorname{IPART}\left(2, M P \_R A N K+1\right)=\) last column index. The number of columns per node is typically based on their relative computing power. To avoid a node with rank MP_RANK doing any work except communication, set IPART \(\left(1, \mathrm{MP} \_\right.\)RANK +1\()=\overline{0}^{-}\)and IPART \(\left(2, \mathrm{MP} \_\right.\)RANK +1\()=-1\). In this exceptional case there is no reference to the array \(\boldsymbol{A}(:,:)\) at that node.

NSETP- (Output) An INTEGER indicating the number of solution components not at constraints. The column indices are output in the array INDEX (: ).

NSETZ- (Output) An INTEGER indicating the solution components held at fixed values. The column indices are output in the array INDEX (: ).

\section*{Optional Argument}

IOPT(:)— (Input) Assumed-size array of derived type S_OPTIONS or D_OPTIONS. This argument is used to change internal parameters of the algorithm. Normally users will not be concerned about this argument, so they would not include it in the argument list for the routine.
\begin{tabular}{|l|l|}
\hline \multicolumn{2}{|c|}{ Packaged Options for PARALLEL_BOUNDED_LSQ } \\
\hline Option Name & Option Value \\
\hline PBLSQ_SET_TOLERANCE & 1 \\
\hline PBLSQ_SET_MAX_ITERATIONS & 2 \\
\hline PBLSQ_SET_MIN_RESIDUAL & 3 \\
\hline
\end{tabular}
\(\operatorname{IOPT}(I O)=\) ? OPTIONS (PBLSQ_SET_TOLERANCE, TOLERANCE) Replaces the default rank tolerance for using a column, from EPSILON (TOLERANCE) to TOLERANCE. Increasing the value of TOLERANCE will cause fewer columns to be increased from their constraints, and may cause the minimum residual RNORM to increase.

IOPT (IO) =? OPTIONS (PBLSQ_SET MIN_RESIDUAL, RESID) Replaces the default target for the minimum residual vector length from 0 to RESID. Increasing the value of RESID can result in fewer iterations and thus increased efficiency. The descent in the optimization will stop at the first point where the minimum residual RNORM is smaller than RESID. Using this option may result in the dual vector not satisfying its optimality conditions, as noted above.
\(\operatorname{IOPT}(I O)=\) PBLSQ_SET_MAX_ITERATIONS
\(\operatorname{IOPT}(I O+1)=\) NEW_MAX_ITERATIONS Replaces the default maximum number of iterations from \(3 * N\) to NEW_MAX_ITERATIONS. Note that this option requires two entries in the derived type array.

\section*{FORTRAN 90 Interface}

Generic: CALL PARALLEL_BOUNDED_LSQ (A, B, X [, ...])
Specific: The specific interface names are S_PARALLEL_BOUNDED_LSQ and D_PARALLEL_BOUNDED_LSQ.

\section*{Description}

Subroutine PARALLEL_BOUNDED_LSQ solves the least-squares linear system \(A x \cong b, \alpha \leq x \leq \beta\), using the algorithm \(B V L S\) found in Lawson and Hanson, (1995), pages 279-283. The new steps involve updating the dual vector and exchange of required data, using MPI. The optional changes to default tolerances, minimum residual, and the number of iterations are new features.

\section*{Example 1: Distributed Equality and Inequality Constraint Solver}

The program PBLSQ_EX1 illustrates the computation of the minimum Euclidean length solution of an \(m^{\prime} \times n^{\prime}\) system of linear inequality constraints, \(G y \geq h\). Additionally the first \(f>0\) of the constraints are equalities. The solution algorithm is based on Algorithm \(L D P\), page 165-166, loc. cit. By allowing the dual variables to be free, the constraints become equalities. The rows of \(E=[G: h]\) are partitioned and assigned random values. When the minimum Euclidean length solution to the inequalities has been calculated, the residuals \(r=G y-h \geq 0\) are computed, with the dual variables to the \(B V L S\) problem indicating the entries of \(r\) that are exactly zero.
```

    PROGRAM PBLSQ_EX1
    ! Use Parallel_boun\overline{ded_LSQ to solve an inequality}
! constraint problem, Gy >= h. Force F of the constraints
! to be equalities. This algorithm uses LDP of
! Solving Least Squares Problems, page 165.
! Forcing equality constraints by freeing the dual is
! new here. The constraints are allocated to the
! processors, by rows, in columns of the array A(:,:).
USE PBLSQ_INT
USE MPI_SETUP INT
USE RAND_INT
USE SHOW_INT
IMPLICIT NONE
INCLUDE "mpif.h"
INTEGER, PARAMETER :: MP=500, NP=400, M=NP+1, \&
N=MP, F=NP/10
REAL(KIND(1D0)), PARAMETER :: ZERO=0D0, ONE=1D0
REAL(KIND(1D0)), ALLOCATABLE :: \&
A(:,:), B(:), BND(:,:), X(:), Y(:), \&
W(:), ASAVE(:,:)
REAL(KIND(1DO)) RNORM
INTEGER, ALLOCATABLE :: INDEX(:), IPART(:,:)
INTEGER K, L, DN, J, JSHIFT, IERROR, NSETP, NSETZ
LOGICAL :: PRINT=.false.
! Setup for MPI:
MP_NPROCS=MP_SETUP()
DN=N/max (1, max (1,MP_NPROCS) ) - 1
ALLOCATE (IPART (2,max (1,MP_NPROCS)))
! Spread constraint rows evenly to the processors.

```
```

    IPART (1,1)=1
    DO L=2,MP NPROCS
        IPART ( }\overline{2},\textrm{L}-1)=\operatorname{IPART}(1,L-1)+D
        IPART (1,L) =IPART (2,L-1) +1
    END DO
    IPART (2,MP_NPROCS ) =N
    ! Define the constraints using random data.
K=max (0,IPART (2,MP_RANK+1) -IPART (1,MP_RANK+1) +1)
ALLOCATE (A (M, K), ASAVE (M,K), BND (2,N), \&
X(N), W(N), B(M), Y(M), INDEX(N))
! The use of ASAVE can be replaced by regenerating the
! data for A(:,:) after the return from
! Parallel_bounded_LSQ
A=rand(A); ASAVE=A
IF(MP_RANK == 0 .and. PRINT) \&
call show(IPART,\&
"Partition of the constraints to be solved")
! Set the right-hand side to be one in the last
! component, zero elsewhere.
B=ZERO;B (M) =ONE
! Solve the dual problem. Letting the dual variable
! have no constraint forces an equality constraint
! for the primal problem.
BND (1,1:F)=-HUGE (ONE); BND (1,F+1:)=ZERO
BND (2,:)=HUGE (ONE)
CALL Parallel bounded LSQ \&
(A, B, BND, X, RNORM, W, INDEX, IPART, \&
NSETP, NSETZ)
! Each processor multiplies its block times the part
! of the dual corresponding to that partition.
Y=ZERO
DO J=IPART(1,MP_RANK+1),IPART (2,MP_RANK+1)
JSHIFT=J-IPART (1,MP_RANK+1) +1
Y=Y+ASAVE (:,JSHIFT)}\overline{\star}X(J
END DO
! Accumulate the pieces from all the processors.
! Put sum into B(:) on rank O processor.
B=Y
IF(MP_NPROCS > 1) \&
CALI MPI_REDUCE(Y, B, M, MPI DOUBLE PRECISION,\&
MPI SUM, 0, MP LIBRARY WORLD}, IERROR)
IF(MP_\overline{RANK == 0) THEN}
! Compute constraint solution at the root.
! The constraints will have no solution if B(M) = ONE.
! All of these example problems have solutions.
B (M) =B (M) -ONE; B=-B/B (M)
END IF

```
```

! Send the inequality constraint or primal solution to all nodes.
IF(MP_NPROCS > 1) \&
CAL\overline{L MPI_BCAST(B, M, MPI_DOUBLE_PRECISION, 0, \&}
MP_LIB\overline{R}ARY_WORLD, IERRO\overline{O})
! For large problems this printing may need to be removed.
IF(MP_RANK == 0 .and. PRINT) \&
cal\ show(B(1:NP), \&
"Minimal length solution of the constraints")
! Compute residuals of the individual constraints.
X=ZERO
DO J=IPART(1,MP_RANK+1),IPART (2,MP_RANK+1)
JSHIFT=J-IPART
X(J)=dot_product(B, A
END DO
! This cleans up residuals that are about rounding error
! unit (times) the size of the constraint equation and
! right-hand side. They are replaced by exact zero.
WHERE (W == ZERO) X=ZERO
W=X
! Each group of residuals is disjoint, per processor.
! We add all the pieces together for the total set of
! constraints.
IF(MP NPROCS > 1) \&
CAL\overline{L MPI_REDUCE (X, W, N, MPI_DOUBLE_PRECISION, \&}
MPI_SUM, 0, MP_LIBRARY_WORLDD, IERR\overline{OR)}
IF (MP - RANK == 0 .and. PRINT) \&
cal\overline{l show(W, "Residuals for the constraints")}
! See to any errors and shut down MPI.
MP_NPROCS=MP_SETUP('Final')
IF(MP RANK == 0) THEN
IF(\overline{COUNT (W < ZERO) == 0 .and.\&}
COUNT(W == ZERO) >= F) WRITE(*,*) \&
" Example 1 for PARALLEL_BOUNDED_LSQ is correct."
END IF
END

```

\section*{Output}

Example 1 for PARALLEL_BOUNDED_LSQ is correct.

\section*{Additional Examples}

\section*{Example 2: Distributed Newton-Raphson Method with Step Control}

The program PBLSQ_EX2 illustrates the computation of the solution of a non-linear system of equations. We use a constrained Newton-Raphson method.
This algorithm works with the problem chosen for illustration. The step-size control used here, employing only simple bounds, may not work on other non-linear systems of equations. Therefore we do not recommend the simple non-linear solving technique illustrated here for an arbitrary
problem. The test case is Brown's Almost Linear Problem, Moré, et al. (1982). The components are given by:
\[
\begin{aligned}
& \text { - } f_{i}(x)=x_{i}+\sum_{j=1}^{n} x_{j}-(n+1), i=1, \ldots, n-1 \\
& \text { - } f_{n}(x)=x_{1} \ldots x_{n}-1
\end{aligned}
\]

The functions are zero at the point \(x=\left(\delta, \ldots, \delta, \delta^{1-n}\right)^{T}\), where \(\delta>1\) is a particular root of the polynomial equation \(n \delta^{n}-(n+1) \delta^{n-1}+1=0\). To avoid convergence to the local minimum \(x=(0, \ldots, 0, n+1)^{T}\), we start at the standard point \(x=(1 / 2, \ldots, 1 / 2,1 / 2)^{T}\) and develop the Newton method using the linear terms \(f(x-y) \approx f(x)-J(x) y \cong 0\), where \(J(x)\) is the Jacobian matrix. The update is constrained so that the first \(n-1\) components satisfy \(x_{j}-y_{j} \geq 1 / 2\), or \(y_{j} \leq x_{j}-1 / 2\). The last component is bounded from both sides, \(0<x_{n}-y_{n} \leq 1 / 2\), or \(x_{n}>y_{n} \geq\left(x_{n}-1 / 2\right)\). These bounds avoid the local minimum and allow us to replace the last equation by \(\sum_{j=1}^{n} \ln \left(x_{j}\right)=0\), which is better scaled than the original. The positive lower bound for \(x_{n}-y_{n}\) is replaced by the strict bound, EPSILON (1D0), the arithmetic precision, which restricts the relative accuracy of \(x_{n}\). The input for routine PARALLEL_BOUNDED_LSQ expects each processor to obtain that part of \(J(x)\) it owns. Those columns of the Jacobian matrix correspond to the partition given in the array \(\operatorname{IPART}(:,:)\). Here the columns of the matrix are evaluated, in parallel, on the nodes where they are required.

PROGRAM PBLSQ_EX2
! Use Parallel_bounded_LSQ to solve a non-linear system
! of equations. The example is an ACM-TOMS test problem, ! except for the larger size. It is "Brown's Almost Linear ! Function."

USE ERROR_OPTION_PACKET
USE PBLSQ \({ }^{-}\)INT
USE MPI_SETUP_INT
USE SHOW_INT
USE Numer̄ical_Libraries, ONLY : N1RTY

IMPLICIT NONE

INTEGER, PARAMETER : : N=200, MAXIT=5
REAL (KIND (1D0)), PARAMETER : : ZERO=0D0, ONE=1D0, \& HALF=5D-1, TWO=2D0 REAL (KIND (1D0)), ALLOCATABLE : : \& \(\mathrm{A}(:,:), \mathrm{B}(:), \operatorname{BND}(:,:), \mathrm{X}(:), \mathrm{Y}(:), \mathrm{W}(:)\) REAL (KIND (1D0)) RNORM INTEGER, ALLOCATABLE : : INDEX(:), IPART (: ,:)

INTEGER K, L, DN, J, JSHIFT, IERROR, NSETP, \& NSETZ, ITER LOGICAL : : PRINT=.false.
```

    TYPE(D_OPTIONS) IOPT(3)
    ! Setup for MPI:
MP_NPROCS=MP_SETUP()
DN=N/max (1, max (1,MP NPROCS) ) -1
ALLOCATE (IPART (2,max}(1,MP_NPROCS)))
! Spread Jacobian matrix columns evenly to the processors.
IPART (1,1)=1
DO L=2,MP NPROCS
IPART (2, L-1) =IPART (1,L-1) +DN
IPART (1,L) =IPART (2,L-1) +1
END DO
IPART (2,MP_NPROCS)=N
K=max (0,IPART (2,MP_RANK+1) -IPART (1,MP_RANK+1)+1)
ALLOCATE (A (N,K), BNDD (2,N), \&
X(N), W(N), B(N), Y(N), INDEX(N))
! This is Newton's method on "Brown's almost
! linear function."
X=HALF
ITER=0
! Turn off messages and stopping for FATAL class errors.
CALL ERSET (4, 0, 0)
NEWTON_METHOD: DO
! Set bounds for the values after the step is taken.
! All variables are positive and bounded below by HALF,
! except for variable N, which has an upper bound of HALF.
BND (1,1:N-1)=-HUGE (ONE)
BND (2,1:N-1)=X (1:N-1) -HALF
BND (1,N) =X (N) - HALF
BND (2,N)=X (N) -EPSILON (ONE)
! Compute the residual function.
B(1:N-1)=SUM(X) +X (1:N-1) - (N+1)
B (N)=LOG (PRODUCT (X))
if(mp_rank == 0 .and. PRINT) THEN
CALL SHOW(B, \&
"Developing non-linear function residual")
END IF
IF (MAXVAL (ABS (B (1:N-1))) <= SQRT(EPSILON(ONE)))\&
EXIT NEWTON_METHOD
! Compute the derivatives local to each processor.
A(1:N-1,:)=ONE
DO J=1,N-1
IF(J < IPART (1,MP_RANK+1)) CYCLE
IF(J > IPART (2,MP_RANK+1)) CYCLE
JSHIFT=J-IPART(1,MP_RANK+1) +1
A(J,JSHIFT) =TWO

```
```

    END DO
    A(N,:)=ONE/X(IPART(1,MP_RANK+1):IPART (2,MP_RANK+1))
    ! Reset the linear independence tolerance.
IOPT(1)=D_OPTIONS(PBLSQ_SET_TOLERANCE,\&
sqrt(EP\overline{SILON(ONE)))}
IOPT(2)=PBLSQ_SET_MAX_ITERATIONS
! If N iterations was not enough on a previous iteration, reset to 2*N.
IF(N1RTY(1) == 0) THEN
IOPT (3) =N
ELSE
IOPT (3) =2*N
CALL E1POP('MP SETUP')
CALL E1PSH('MP_SETUP')
END IF
CALL parallel_bounded_LSQ \&
(A, B, BND, Y, RNORM, W, INDEX, IPART, NSETP, \&
NSETZ,IOPT=IOPT)
! The array Y(:) contains the constrained Newton step.
! Update the variables.
X=X-Y
IF(mp_rank == 0 .and. PRINT) THEN
CALL show(BND, "Bounds for the moves")
CALL SHOW(X, "Developing Solution")
CALL SHOW((/RNORM/), \&
"Linear problem residual norm")
END IF
! This is a safety measure for not taking too many steps.
ITER=ITER+1
IF(ITER > MAXIT) EXIT NEWTON_METHOD
END DO NEWTON_METHOD
IF(MP RANK == 0) THEN
IF(\overline{ITER <= MAXIT) WRITE(*,*) \&}
" Example 2 for PARALLEL_BOUNDED_LSQ is correct."
END IF
! See to any errors and shut down MPI.
MP_NPROCS=MP_SETUP('Final')
END

```

\section*{LSARG}

Solves a real general system of linear equations with iterative refinement.

\section*{Required Arguments}
\(A-\quad \mathrm{N}\) by N matrix containing the coefficients of the linear system. (Input)
\(\boldsymbol{B}-\quad\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input) Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSARG (A, B, X [, ...])

Specific: The specific interface names are S_LSARG and D_LSARG.

\section*{FORTRAN 77 Interface}

Single: CALL LSARG (N, A, LDA, B, IPATH, X)
Double: The double precision name is DLSARG

\section*{ScaLAPACK Interface}

Generic: CALL LSARG (A0, B0, X0 [, ...])
Specific: The specific interface names are S_LSARG and D_LSARG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSARG solves a system of linear algebraic equations having a real general coefficient matrix. It first uses the LFCRG to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using the iterative refinement routine LFIRG. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

LSARG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSARG solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ARG/DL2ARG. The reference is:

CALL L2ARG (N, A, LDA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
FACT - Work vector of length \(\mathrm{N}^{2}\) containing the \(L U\) factorization of A on output.
IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficients of the linear system. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X O}\) - Local vector of length MXLDA containing the local portions of the distributed vector X . \(X\) contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has real general form and the right-hand-side vector \(b\) has three elements.
```

USE LSARG INT
USE WRRRN - INT
IMPLICIT NONE

```

!
Output
X
\(\begin{array}{rrr}1 & 2 & 3 \\ 1.000 & 1.500 & 1.000\end{array}\)

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has real general form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSARG INT
USE WRRRN_-INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE

```
```

    INCLUDE 'mpif.h'
    N, DESCA(9), DESCX(9)
    INTEGER INFO, MXLDA, MXCOL
    REAL, ALLOCATABLE :: A(:,:), B(:), X(:)
    REAL, ALLOCATABLE :: AO (:,:), BO(:), X0 (:)
    PARAMETER (N=3)
    MP_NPROCS = MP_SETUP()
    IF (MP_RANK .EQ. O) THEN
        ALLOCATE (A (N,N), B(N), X(N))
                        Set values for A and B
        A(1,:) = (/ 33.0, 16.0, 72.0/)
        A(2,:) = (/-24.0, -10.0, -57.0/)
        A(3,:) = (/ 18.0, -11.0, 7.0/)
        B = (/129.0, -96.0, 8.5/)
    ENDIF
        Set up a 1D processor grid and define
        its context id, MP ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
        Get the array descriptor entities MXLDA,
        AND MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
    Sèt up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT (DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICT\overline{XT, MXLDA, INFO)}
    Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA), X0 (MXLDA))
            Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, A0)
    CALL SCALAPACK_MAP(B, DESCX, BO)
    Solve the system of equations
    CALL LSARG (A0, B0, X0)
    Unmap the results from the distributed
    arrays back to a non-distributed array.
    After the unmap, only Rank=0 has the full
            array.
    CALL SCALAPACK_UNMAP(X0, DESCX, X)
            Print results.
            Only Rank=0 has the solution, X.
        IF (MP_RANK .EQ. 0) CALL WRRRN ('X', X, 1, N, 1)
        IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)
        DEALLOC
            Exit ScaLAPACK usage
        CALL SCALAPACK_EXIT(MP_ICTXT)
            Shut down MPI
    MP_NPROCS = MP_SETUP('FINAL')
    EN\overline{D}
    ```

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{3}{c}{X} & \\
1 & 2 & 3 \\
\(1.000^{2}\) & \(1.500^{2}\) & \(1.000^{2}\)
\end{tabular}

\section*{LSLRG}

Solves a real general system of linear equations without iterative refinement.

\section*{Required Arguments}
\(A-\mathrm{N}\) by N matrix containing the coefficients of the linear system. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLRG (A, B, X \([, \ldots]\) )
Specific: The specific interface names are S_LSLRG and D_LSLRG.

\section*{FORTRAN 77 Interface}

Single: CALL LSLRG (N, A, LDA, B, IPATH, X)
Double: The double precision name is DLSLRG.

\section*{ScaLAPACK Interface}

Generic: CALL LSLRG (A0, B0, X0 [, ...])

Specific: The specific interface names are _ LSLRG \(^{\text {and D_LSLRG. }}\)
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLRG solves a system of linear algebraic equations having a real general coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSLRG first uses the routine LFCRG to compute an \(L U\) factorization of the coefficient matrix based on Gauss elimination with partial pivoting. Experiments were analyzed to determine efficient implementations on several different computers. For some supercomputers, particularly those with efficient vendor-supplied BLAS, versions that call Level 1, 2 and 3 BLAS are used. The remaining computers use a factorization method provided to us by Dr. Leonard J. Harding of the University of Michigan. Harding's work involves "loop unrolling and jamming" techniques that achieve excellent performance on many computers. Using an option, LSLRG will estimate the condition number of the matrix. The solution of the linear system is then found using LFSRG.

The routine LSLRG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This occurs only if \(A\) is close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that small changes in \(A\) can cause large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that either LIN_SOL_SVD or LSARG be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LRG/DL2LRG. The reference is:

CALL L2LRG (N, A, LDA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
FACT \(-\mathrm{N} \times \mathrm{N}\) work array containing the \(L U\) factorization of A on output. If A is not needed, A and FACT can share the same storage locations. See Item 3 below to avoid memory bank conflicts.

IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(\mathbf{W K}\) - Work vector of length N .
2. Informational errors
Type Code

31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LRG the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2); respectively, in LSLRG. Additional memory allocation for FACT and option value restoration are done automatically in LSLRG. Users directly calling L2 LRG can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLRG or L2LRG. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLRG temporarily replaces IVAL(2) by IVAL(1). The routine L2CRG computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CRG skips this computation. LSLRG restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficients of the linear system. (Input)

BO - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)

XO - Local vector of length MXLDA containing the local portions of the distributed vector \(X\).
X contains the solution to the linear system. (Output)
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example 1}

A system of three linear equations is solved. The coefficient matrix has real general form and the right-hand-side vector \(b\) has three elements.
```

USE LSLRG INT
USE WRRRN_INT

```
```

    IMPLICIT NONE
    !
INTEGER }r\mathrm{ LDA, N
REAL A (LDA,N), B (N), X (N)
Set values for A and B
A(1,:) = (/ 33.0, 16.0, 72.0/)
A (2,:) = (/-24.0, -10.0, -57.0/)
A(3,:) = (/ 18.0, -11.0, 7.0/)
!
B = (/129.0 -96.0 8.5/)
! Solve the system of equations
CALL LSLRG (A, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrr}
\multicolumn{3}{c}{X} \\
1 & 2 & 3 \\
1.000 & 1.500 & 1.000
\end{tabular}

\section*{Additional Example}

\section*{Example 2}

A system of \(\mathrm{N}=16\) linear equations is solved using the routine L2LRG. The option manager is used to eliminate memory bank conflict inefficiencies that may occur when the matrix dimension is a multiple of 16 . The leading dimension of \(\operatorname{FACT}=A\) is increased from \(N\) to \(N+\operatorname{IVAL}(3)=17\), since \(\mathrm{N}=16=\operatorname{IVAL}(4)\). The data used for the test is a nonsymmetric Hadamard matrix and a right-hand side generated by a known solution, \(x_{j}=j, j=1, \ldots, N\).
```

USE L2LRG INT
USE IUMAG }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
USE SGEMV - INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (LDA=17,N=16)
INTEGER ICHP, IPATH, IPUT, KBANK
REAL ONE, ZERO
PARAMETER (ICHP=1, IPATH=1, IPUT=2, KBANK=16, ONE=1.0E0, \&
ZERO=0.0E0)
INTEGER I, IPVT(N), J, K, NN
REAL A(LDA,N), B(N), WK (N), X(N)
SPECIFICATIONS FOR SAVE VARIABLES
INTEGER IOPT(1), IVAL(4)

```
!
```

    SAVE IVAL
    DATA IVAL/1, 16, 1, 16/
    A(1,1) = ONE
    NN = 1
    DO 20 K=1, 4
        DO 10 J=1, NN
            DO 10 I=1, NN
                A(NN+I,J) = -A(I,J)
                A(I,NN+J) = A(I,J)
                A(NN+I,NN+J) = A(I,J)
    1 0
        CONTINUE
        NN = NN + NN
    2 0 ~ C O N T I N U E ~
    ! Generate right-hand-side.
DO 30 J=1, N
X(J) = J
CONTINUE
Set B = A*X.
CALL SGEMV ('N', N, N, ONE, A, LDA, X, 1, ZERO, B, 1)
Clear solution array.
X = ZERO
IOPT(1) = KBANK
CALL IUMAG ('MATH', ICHP, IPUT, 1, IOPT, IVAL)
Solve A*X = B.
CALL L2LRG (N, A, LDA, B, IPATH, X, A, IPVT, WK)
Print results
CALL WRRRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline & & & & & X & & & & \\
\hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\
\hline 1.00 & 2.00 & 3.00 & 4.00 & 5.00 & 6.00 & 7.00 & 8.00 & 9.00 & 10.00 \\
\hline 11 & 12 & 13 & 14 & 15 & & & & & \\
\hline 11.00 & 12.00 & 13.00 & 14.00 & 15.00 & 16. & & & & \\
\hline
\end{tabular}

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has real general form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and (see Chapter 11, "Utilities") are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

    USE MPI_SETUP_INT
    USE LSLRG INT
    USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'     INTEGER N, DESCA(9), DESCX(9)     INTEGER INFO, MXCOL, MXLDA     REAL, ALLOCATABLE :: A(:,:), B(:), X(:)     REAL, ALLOCATABLE :: A0(:,:), B0(:), X0(:)     PARAMETER (N=3)     MP NPROCS = MP SETUP()     IF(MP_RANK .EQ. 0) THEN         ALLOCATE (A (N,N), B(N), X(N))                             Set values for A and B         A(1,:) = (/ 33.0, 16.0, 72.0/)         A(2,:) = (/-24.0, -10.0, -57.0/)         A(3,:) = (/ 18.0, -11.0, 7.0/)         B = (/129.0, -96.0, 8.5/)     ENDIF         Set up a 1D processor grid and define         its context id, MP_ICTXT     CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)         Get the array descriptor entities MXLDA,                     and MXCOL     CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)     Set up the array descriptors     CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)     CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICT\overline{XT, MXLDA, INFO)}     Allocate space for the local arrays     ALLOCATE (A0(MXLDA,MXCOL), B0(MXLDA), XO(MXLDA))                             Map input arrays to the processor grid     CALL SCALAPACK MAP(A, DESCA, A0)     CALL SCALAPACK_MAP(B, DESCX, BO)                             Solve the system of equations     CALL LSLRG (A0, B0, X0)             Unmap the results from the distributed                     arrays back to a non-distributed array.                     After the unmap, only Rank=0 has the full                     array.     CALL SCALAPACK_UNMAP(X0, DESCX, X)                             Print results                             Only Rank=0 has the solution, X.     IF(MP_RANK .EQ. 0) CALL WRRRN ('X', X, 1, N, 1)     IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)     DEALLOCATE(A0, B0, X0)                             Exit ScaLAPACK usage     CALL SCALAPACK_EXIT(MP_ICTXT)                                     Shut down MPI     MP_NPROCS = MP_SETUP(`FINAL')
EN\overline{D}

```

\section*{Output}
\begin{tabular}{rrr}
\multicolumn{3}{c}{\(X\)} \\
1 & 2 & 3 \\
1.000 & 1.500 & 1.000
\end{tabular}

\section*{LFCRG}

CAPABLE
Computes the \(L U\) factorization of a real general matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix to be factored. (Input)
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, A and FACT can share the same storage locations.
IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)

RCOND - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).

FORTRAN 90 Interface
Generic: CALL LFCRG (A, FACT, IPVT, RCOND, [, ...])

Specific: The specific interface names are S_LFCRG and D_LFCRG.

\section*{FORTRAN 77 Interface}

Single: CALL LFCRG (N, A, LDA, FACT, LDFACT, IPVT, RCOND)
Double: The double precision name is DLFCRG.

\section*{ScaLAPACK Interface}

Generic: CALL LFCRG (A0, FACTO, IPVT0, RCOND [, ...])

Specific: The specific interface names are S_LFCRG and D_LFCRG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCRG performs an \(L U\) factorization of a real general coefficient matrix. It also estimates the condition number of the matrix. The underlying code is based on either LINPACK , LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm. Otherwise, partial pivoting is used.
The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described in a paper by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCRG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) either is singular or is very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with routines LFIRG, LFSRG and LFDRG. To solve systems of equations with multiple right-hand-side vectors, use LFCRG followed by either LFIRG or LFSRG called once for each right-hand side. The routine LFDRG can be called to compute the determinant of the coefficient matrix after LFCRG has performed the factorization.
Let \(F\) be the matrix FACT and let \(p\) be the vector IPVT. The triangular matrix \(U\) is stored in the upper triangle of \(F\). The strict lower triangle of \(F\) contains the information needed to reconstruct \(L\) using
\[
L^{-1}=L_{N-1} P_{N-1} \ldots L_{1} P_{1}
\]
where \(P_{k}\) is the identity matrix with rows \(k\) and \(p_{k}\) interchanged and \(L_{k}\) is the identity with \(F_{i k}\) for \(i=k+1, \ldots, N\) inserted below the diagonal. The strict lower half of \(F\) can also be thought of as containing the negative of the multipliers. LFCRG is based on the LINPACK routine SGECO; see Dongarra et al. (1979). SGECO uses unscaled partial pivoting.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CRG/DL2CRG. The reference is:

CALL L2CRG (N, A, LDA, FACT, LDFACT, IPVT, RCOND, WK)
The additional argument is

WK - Work vector of length N
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
42 The input matrix is singular

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix to be factored. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(L U\) factorization of the matrix A. (Output)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFCRG is called to factor the matrix and to check for singularity or ill-conditioning. LFIRG is called to determine the columns of the inverse.
```

USE LFCRG INT
USE UMACH INT
USE LFIRG INT
USE WRRRN INT
IMPLICIT NONE
Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT(N), J, NOUT
REAL A(LDA,N), AINV (LDA,N), FACT(LDFACT,N), RCOND, \&
RES(N), RJ(N)
Set values for A

```
!
```

        A(1,:) = (/ 1.0, 3.0, 3.0/)
        A(2,:) = (/ 1.0, 3.0, 4.0/)
        A(3,:) = (/ 1.0, 4.0, 3.0/)!
        CALL LFCRG (A, FACT, IPVT, RCOND)
        Print the reciprocal condition number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998) RCOND, 1.0E0/RCOND
        Set up the columns of the identity
        matrix one at a time in RJ
    RJ = 0.0E0
    DO 10 J=1, N
        RJ(J) = 1.0
        RJ is the J-th column of the identity
        matrix so the following LFIRG
        reference places the J-th column of
        the inverse of A in the J-th column
        of AINV
        CALL LFIRG (A, FACT, IPVT, RJ, AINV(:,J), RES)
        RJ(J) = 0.0
    1 0 ~ C O N T I N U E
    ! Print results
CALL WRRRN ('AINV', AINV)
!
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < . 02
L1 Condition number < 100.0

AINV
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 7.000 & -3.000 & -3.000 \\
2 & -1.000 & 0.000 & 1.000 \\
3 & -1.000 & 1.000 & 0.000
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. LFCRG is called to factor the matrix and to check for singularity or ill-conditioning. LFIRG is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFCRG_INT
USE UMACH_INT
USE LFIRG }\mp@subsup{}{-}{INT
USE WRRRN_INT
USE SCALA\overline{PACK_SUPPORT}

```
```

    IMPLICIT NONE
    INCLUDE 'mpif.h'
                            Declare variables
    INTEGER J, LDA, N, DESCA(9), DESCL(9)
    INTEGER INFO, MXCOL, MXLDA, NOUT
    INTEGER, ALLOCATABLE :: IPVTO(:)
    REAL, ALLOCATABLE :: A(:,:), AINV(:,:), X0(:), RJ(:)
    REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), RESO(:), RJO(:)
    REAL RCOND
    PARAMETER (LDA=3, N=3)
    MP_NPROCS = MP_SETUP()
    IF(MP RANK .EQ. O) THEN
        A\overline{LLOCATE (A (LDA,N), AINV (LDA,N))}
                            Set values for A
        A(1,:) = (/ 1.0, 3.0, 3.0/)
        A(2,:) = (/ 1.0, 3.0, 4.0/)
        A(3,:) = (/ 1.0, 4.0, 3.0/)
    ENDIF
    Set up a 1D processor grid and define
    its context id, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
        Get the array descriptor entities MXLDA,
        and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
                Set up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
                Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL) , X0 (MXLDA), FACTO (MXLDA,MXCOL), RJ (N) , &
        RJO (MXLDA) , RESO (MXLDA), IPVTO (MXLDA))
            Map input arrays to the processor grid
    CALL SCALAPACK MAP(A, DESCA, A0)
                            Call the factorization routine
    CALL LFCRG (A0, FACTO, IPVTO, RCOND)
                                    Print the reciprocal condition number
                                    and the L1 condition number
    IF(MP RANK .EQ. O) THEN
        CALL UMACH (2, NOUT)
        WRITE (NOUT,99998) RCOND, 1.0EO/RCOND
    ENDIF
    Set up the columns of the identity
    matrix one at a time in RJ
    RJ = 0.0E0
    DO 10 J=1, N
        RJ(J) = 1.0
        CALL SCALAPACK_MAP(RJ, DESCL, RJO)
                        RJ is the J-th column of the identity
                        matrix so the following LFIRG
                        reference computes the J-th column of
                    the inverse of A
    CALL LFIRG (A0, FACT0, IPVT0, RJ0, X0, RESO)
    RJ (J) = 0.0
    CALL SCALAPACK_UNMAP(X0, DESCL, AINV(:,J))
    1 0 ~ C O N T I N U E
    ```
```

! Print results
Only Rank=O has the solution, X.
IF(MP RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (MP_RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE(A0, IPVTO, FACTO, RESO, RJ, RJO, XO)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
MP_NPROCS = MP_SETUP(`FINAL')
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND < . }0
L1 Condition number < 100.0
AINV

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 7.000 | -3.000 | -3.000 |
| 2 | -1.000 | 0.000 | 1.000 |
| 3 | -1.000 | 1.000 | 0.000 |

```

\section*{LFTRG}

\section*{包MPI}

Computes the \(L U\) factorization of a real general matrix.

\section*{Required Arguments}
\(A-\mathrm{N}\) by N matrix to be factored. (Input)
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, A and FACT can share the same storage locations.
IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTRG (A, FACT, IPVT [, ...])

Specific: The specific interface names are S_LFTRG and D_LFTRG.

\section*{FORTRAN 77 Interface}

Single: CALL LfTRG (N, A, LDA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTRG.

\section*{ScaLAPACK Interface}

Generic: CALL LFTRG (A0, FACTO, IPVTO [, ...])

Specific: The specific interface names are S_LFTRG and D_LFTRG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFTRG performs an \(L U\) factorization of a real general coefficient matrix. The underlying code is based on either LINPACK , LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same norm. Otherwise, partial pivoting is used.

The routine LFTRG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) is singular or very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with routines LFIRG, LFSRG and LFDRG.
To solve systems of equations with multiple right-hand-side vectors, use LFTRG followed by either LFIRG or LFSRG called once for each right-hand side. The routine LFDRG can be called to compute the determinant of the coefficient matrix after LFTRG has performed the factorization. Let \(F\) be the matrix FACT and let \(p\) be the vector IPVT. The triangular matrix \(U\) is stored in the upper triangle of \(F\). The strict lower triangle of \(F\) contains the information needed to reconstruct \(L^{-1}\) using
\[
L^{-1}=L_{N-1} P_{N-1} \ldots L_{1} P_{1}
\]
where \(P_{k}\) is the identity matrix with rows \(k\) and \(p_{k}\) interchanged and \(L_{k}\) is the identity with \(F_{i k}\) for \(i=k+1, \ldots, N\) inserted below the diagonal. The strict lower half of \(F\) can also be thought of as containing the negative of the multipliers.

Routine Lftrg is based on the LINPACK routine SGEFA. See Dongarra et al. (1979). The routine SGEFA uses partial pivoting.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2TRG/ DL2TRG. The reference is:
```

CALL L2TRG (N, A, LDA, FACT, LDFACT, IPVT, WK)

```

The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N used for scaling.
2. Informational error

Type Code
42 The input matrix is singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
AO- MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix to be factored. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(L U\) factorization of the matrix A. (Output)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A linear system with multiple right-hand sides is solved. Routine LFTRG is called to factor the coefficient matrix. The routine LFSRG is called to compute the two solutions for the two righthand sides. In this case, the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCRG to perform the factorization, and LFIRG to compute the solutions.
```

        USE LFTRG INT
        USE LFSRG_INT
        USE WRRRN_INT
    ! Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT (N), J
REAL A(LDA,LDA), B (N, 2), FACT (LDFACT,LDFACT), X (N, 2)
Set values for A and B
A = ($$
\begin{array}{lll}{1.0}&{3.0}&{3.0)}\end{array}
$$)
( 1.0 3.0 4.0)
(1.0 4.0 3.0)
B =( 1.0 10.0)
(4.0 14.0)
( -1.0 9.0)
DATA A/1.0, 1.0, 1.0, 3.0, 3.0, 4.0, 3.0, 4.0, 3.0/
DATA B/1.0, 4.0, -1.0, 10.0, 14.0, 9.0/
CALL LFTRG (A, FACT, IPVT)
Solve for the two right-hand sides
DO 10 J=1, 2
CALL LFSRG (FACT, IPVT, B(:,J), X(:,J))
CONTINUE
CALL WRRRN ('X', X)
END
Output
X
1 2
-2.000 1.000
-2.000 -1.000
3.000 4.000

```

\section*{ScaLAPACK Example}

A linear system with multiple right-hand sides is solved. Routine LFTRG is called to factor the coefficient matrix. The routine LFSRG is called to compute the two solutions for the two righthand sides. In this case, the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCRG to perform the factorization, and LFIRG to compute the solutions. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFTRG_INT
USE LFSRG }\mp@subsup{}{-}{INT
USE WRRRN_INT
USE SCALAPACK_SUPPORT

```
```

    IMPLICIT NONE
    INCLUDE `mpif.h'
        Declare variables
    INTEGER J, LDA, N, DESCA(9), DESCL(9)
    INTEGER INFO, MXCOL, MXLDA
    INTEGER, ALLOCATABLE :: IPVTO(:)
    REAL, ALLOCATABLE :: A(:,:), B(:,:), X(:,:), X0(:)
    REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), BO(:)
    PARAMETER (LDA=3, N=3)
    MP NPROCS = MP SETUP()
    IF(MP_RANK .EQ. O) THEN
        ALLOCATE (A (LDA,N), B(N,2), X(N,2))
                            Set values for A and B
        A(1,:) = (/ 1.0, 3.0, 3.0/)
        A(2,:) = (/ 1.0, 3.0, 4.0/)
        A(3,:) = (/ 1.0, 4.0, 3.0/)
        B(1,:) = (/ 1.0, 10.0/)
        B(2,:) = (/ 4.0, 14.0/)
        B (3,:) = (/-1.0, 9.0/)
    ENDIF
    Set up a 1D processor grid and define
    its context id, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
                    Get the array descriptor entities MXLDA,
                    and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
            Sēt up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP_MB, 1,-0, 0, MP_ICT\overline{XT, MXLDA, INFO)}
                    Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), XO (MXLDA), FACTO (MXLDA,MXCOL), BO (MXLDA), &
        IPVT0 (MXLDA))
                            Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, AO)
        Call the factorization routine
    CALL LFTRG (A0, FACT0, IPVT0)
        Set up the columns of the B
        matrix one at a time in XO
    DO 10 J=1, 2
        CALL SCALAPACK_MAP(B(:,j), DESCL, BO)
                            Solve for the J-th column of X
        CALL LFSRG (FACTO, IPVTO, B0, X0)
        CALL SCALAPACK_UNMAP (X0, DESCL, X(:,J))
    1 0 ~ C O N T I N U E
                            Print results.
                            Only Rank=0 has the solution, X.
    IF(MP RANK.EQ.O) CALL WRRRN ('X', X)
    IF (MP _RANK .EQ. O) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, IPVT0, FACTO, X0)
                            Exit ScaLAPACK usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
            Shut down MPI
    MP_NPROCS = MP_SETUP('FINAL')
    ```

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{2}{c}{X} \\
& \multicolumn{1}{c}{1} & 2 \\
1 & -2.000 & 1.000 \\
2 & -2.000 & -1.000 \\
3 & 3.000 & 4.000
\end{tabular}

\section*{LFSRG}

\section*{勻MPI \\ CAPABLE}

Solves a real general system of linear equations given the \(L U\) factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - N by N matrix containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCRG. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from subroutine LFCRG or LFTRG. (Input).
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
\(\boldsymbol{L D F A C T}\) - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: ІРАтН \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSRG (FACT, IPVT, B, X [, ...])

Specific: The specific interface names are S_LFSRG and D_LFSRG.

\section*{FORTRAN 77 Interface}

Single: CALL LFSRG (N, FACT, LDFACT, IPVT, B, IPATH, X)
Double: The double precision name is DLFSRG.

\section*{ScaLAPACK Interface}

Generic: CALL LFSRG (FACT0, IPVT0, B0, X0 [, ...])
Specific: The specific interface names are S_LFSRG and D_LFSRG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFSRG computes the solution of a system of linear algebraic equations having a real general coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCRG or LFTRG. The solution to \(A x=b\) is found by solving the triangular systems \(L y=b\) and \(U x=y\). The forward elimination step consists of solving the system \(L y=b\) by applying the same permutations and elimination operations to \(b\) that were applied to the columns of \(A\) in the factorization routine. The backward substitution step consists of solving the triangular system \(U x=y\) for \(x\).

LFSRG and LFIRG both solve a linear system given its \(L U\) factorization. LFIRG generally takes more time and produces a more accurate answer than LFSRG. Each iteration of the iterative refinement algorithm used by LFIRG calls LFSRG. The underlying code is based on either LINPACK , LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT as output from routine LFCRG. FACT contains the \(L U\) factorization of the matrix A. (Input)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization as output from subroutine LFCRG or LFTRG/DLFTRG. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Local vector of length MXLDA containing the local portions of the distributed vector X . \(X\) contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse is computed for a real general \(3 \times 3\) matrix. The input matrix is assumed to be wellconditioned, hence, LFTRG is used rather than LFCRG.
```

    USE LFSRG INT
    USE LFTRG }\mp@subsup{}{}{-}\mathrm{ INT
    USE WRRRN_INT
    ! Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER I, IPVT(N), J
REAL A(LDA,LDA), AINV (LDA,LDA), FACT (LDFACT,LDFACT), RJ (N)
!
!
CALL LFTRG (A, FACT, IPVT)
RJ = 0.0E0
DO 10 J=1, N
RJ (J) = 1.0
RJ is the J-th column of the identity
matrix so the following LFSRG
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFSRG (FACT, IPVT, RJ, AINV(:,J))
RJ(J) = 0.0
10 CONTINUE
! Print results
CALL WRRRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{3}{c}{ AINV } \\
& 1 & 2 & 3 \\
1 & 7.000 & -3.000 & -3.000 \\
2 & -1.000 & 0.000 & 1.000 \\
3 & -1.000 & 1.000 & 0.000
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. The input matrix is assumed to be well-conditioned, hence, LFTRG is used rather than LFCRG. LFSRG is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFTRG_INT
USE UMACH - INT
USE LFSRG_INT
USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
INTEGER, ALLOCATABLE : : IPVTO(:)
REAL, ALLOCATABLE :: A(:,:), AINV(:,:), XO(:), RJ(:)
REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), RJO(:)
PARAMETER (LDA=3,N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ- O) THEN
ALLLOCATE (A (LDA,N), AINV(LDA,N))
Set values for A
A(1,:) = (/ 1.0, 3.0, 3.0/)
A(2,:) = (/ 1.0, 3.0, 4.0/)
A(3,:) = (/ 1.0, 4.0, 3.0/)
ENDIF
Set up a 1D processor grid and define
its context id, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Se}t\mathrm{ up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT (DESCL, N, 1, MP MB, 1, 0, 0, MP ICT\overline{XT, MXLDA, INFO)}
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL) , X0 (MXLDA) ,FACTO (MXLDA,MXCOL), RJ(N) , \&
RJ0 (MXLDA), IPVTO (MXLDA))

```
```

! Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, AO)
Call the factorization routine
CALL LFTRG (A0, FACT0, IPVT0)
Set up the columns of the identity
matrix one at a time in RJ
RJ = 0.0EO
DO 10 J=1, N
RJ(J) = 1.0
CALL SCALAPACK_MAP(RJ, DESCL, RJO)
RJ is the J-th column of the identity
matrix so the following LFIRG
reference computes the J-th column of
the inverse of A
CALL LFSRG (FACTO, IPVTO, RJO, X0)
RJ(J) = 0.0
CALL SCALAPACK_UNMAP(X0, DESCL, AINV(:,J))
1 0 ~ C O N T I N U E ~
Print results
Only Rank=0 has the solution, AINV.
IF(MP_RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (M\overline{P}}\mathrm{ RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE(A0, IPVTO, FACTO, RJ, RJO, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

\section*{AINV}
```

|  | 0 | 1 | 2 |
| ---: | ---: | ---: | ---: |
| 1 | 7.000 | -3.000 | -3.000 |
| 2 | -1.000 | 0.000 | 1.000 |
| 3 | -1.000 | 1.000 | 0.000 |

```

\section*{LFIRG}

Uses iterative refinement to improve the solution of a real general system of linear equations.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the coefficient matrix of the linear system. (Input)
\(\boldsymbol{F A C T}\) - N by N matrix containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCRG/DLFCRG or LFTRG/DLFTRG. (Input).
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCRG/DLFCRG or LFTRG/DLFTRG. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input).
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
\(\boldsymbol{R E S}\) - Vector of length \(N\) containing the final correction at the improved solution. (Output)

\section*{Optional Arguments}
\(N-\) Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A * X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LfiRG (A, FACT, IPVT, B, X, RES [,...])
Specific: The specific interface names are \(S_{\text {_LFIRG }}\) and D_LFIRG.

\section*{FORTRAN 77 Interface}

Single: CALL LFIRG (N, A, LDA, FACT, LDFACT, IPVT, B, IPATH, X, RES)
Double: The double precision name is DLFIRG.

\section*{ScaLAPACK Interface}

Generic: CALL LFIRG (A0, FACT0, IPVT0, B0, X0, RESO [, ...])
Specific: The specific interface names are S_LFIRG and D_LFIRG.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFIRG computes the solution of a system of linear algebraic equations having a real general coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCRG or LFTRG.
Iterative refinement fails only if the matrix is very ill-conditioned.
Routines LFIRG and LFSRG both solve a linear system given its \(L U\) factorization. LFIRG generally takes more time and produces a more accurate answer than LFSRG. Each iteration of the iterative refinement algorithm used by LFIRG calls LFSRG.

\section*{Comments}

Informational error
\begin{tabular}{ll} 
Type & Code \\
2 & 2 \begin{tabular}{l} 
The input matrix is too ill-conditioned for iterative refinement to be \\
effective.
\end{tabular}
\end{tabular}

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
\(\boldsymbol{A O}\) - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the linear system. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT as output from routine LFCRG or LFTRG. FACT contains the \(L U\) factorization of the matrix A. (Input)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization as output from subroutine LFCRG or LFTRG. (Input)

BO - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)

X0 - Local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

RESO - Local vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the final correction at the improved solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding 0.5 to the second element.
```

    USE LFIRG INT
    USE LFCRG -INT
    USE UMACH_INT
    USE WRRRN_INT
    ```
! Declare variables
    PARAMETER (LDA=3, LDFACT=3, \(N=3\) )
    INTEGER IPVT(N), NOUT
    REAL A(LDA,LDA), B(N), FACT (LDFACT,LDFACT), RCOND, RES(N), X(N)
    Set values for A and B
            \(A=\left(\begin{array}{lll}1.0 & 3.0 & 3.0\end{array}\right)\)
            \(\left(\begin{array}{lll}1.0 & 3.0 & 4.0\end{array}\right)\)
            \(\left(\begin{array}{lll}1.0 & 4.0 & 3.0\end{array}\right)\)
                            \(B=\left(\begin{array}{lll}-0.5 & -1.0 & 1.5\end{array}\right)\)
    DATA A/1.0, 1.0, 1.0, 3.0, 3.0, 4.0, 3.0, 4.0, 3.0/
    DATA B/-0.5, -1.0, 1.5/
\(!\)
    CALL LFCRG (A, FACT, IPVT, RCOND)
    CALL UMACH (2, NOUT)
    WRITE (NOUT, 99999) RCOND, 1.0EO/RCOND
    DO \(10 \mathrm{~J}=1,3\)
        CALL LFIRG (A, FACT, IPVT, B, X, RES)
        CALL WRRRN ('X', X, 1, N, 1)
        \(B(2)=B(2)+0.5\)
    10 CONTINUE
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ', F6.3)
    END

\section*{Output}
```

RCOND < 0.02
L1 Condition number < 100.0
X

| 1 | 2 | 3 |
| ---: | ---: | ---: |
| -5.000 | 2.000 | -0.500 |

            X
    | 1 | 2 | 3 |
| ---: | ---: | ---: |
| -6.500 | 2.000 | 0.000 |

            X
    | 1 | 2 | 3 |
| ---: | ---: | ---: |
| -8.000 | 2.000 | 0.500 |

```

\section*{ScaLAPACK Example}

The same set of linear systems is solved successively as a distributed example. The right-hand side vector is perturbed after solving the system each of the first two times by adding 0.5 to the second element. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFIRG_INT
USE UMACH - INT
USE LFCRG_INT
USE WRRRN INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA, NOUT
INTEGER, ALLOCATABLE :: IPVTO(:)
REAL, ALLOCATABLE :: A(:,:), B(:), X(:), XO(:), AINV(:,:)
REAL, ALLOCATABLE :: A0(:,:), FACTO(:,:), RESO(:), BO(:)
REAL RCOND
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP RANK .EQ. 0) THEN
ALLOCATE (A(LDA,N), AINV (LDA,N), B(N), X(N))
Set values for A and B
A(1,:) = (/ 1.0, 3.0, 3.0/)
A(2,:) = (/ 1.0, 3.0, 4.0/)
A(3,:) = (/ 1.0, 4.0, 3.0/)
B(:) = (/-0.5, -1.0, 1.5/)
ENDIF
Set up a 1D processor grid and define
its context id, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL

```
\(!\)
```

    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
                    Set up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXXT, MXLDA, INFO)
                            Allocate space for the local arrays
    ALLOCATE (AO (MXLDA,MXCOL), X0 (MXLDA) ,FACTO (MXLDA,MXCOL), &
        B0 (MXLDA), RESO(MXLDA), IPVTO (MXLDA))
        Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, AO)
                                    Call the factorization routine
    CALL LFCRG (AO, FACTO, IPVTO, RCOND)
        Print the reciprocal condition number
        and the L1 condition number
    IF(MP_RANK .EQ. O) THEN
        CALL UMACH (2, NOUT)
        WRITE (NOUT,99998) RCOND, 1.0E0/RCOND
    ENDIF
    ! Solve the three systems
! one at a time in X
DO 10 J=1, 3
CALL SCALAPACK_MAP(B, DESCL, BO)
CALL LFIRG (A0-, FACT0, IPVTO, B0, X0, RESO)
CALL SCALAPACK_UNMAP (X0, DESCL, X)
Print results
Only Rank=0 has the solution, X.
IF(MP_RANK.EQ.O) CALL WRRRN ('X', X, 1, N, 1)
IF(MP RANK.EQ.0) B(2) = B(2) + 0.5
1 0 ~ C O N T I N U E ~
IF (MP_RANK .EQ. O) DEALLOCATE (A, AINV, B)
DEALLOCATE (A0, B0, IPVT0, FACTO, RES0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.02
L1 Condition number < 100.0
\begin{tabular}{rrr} 
& \multicolumn{3}{c}{ X } & 3 \\
-5.000 & 2.000 & -0.500
\end{tabular}
\begin{tabular}{rrr} 
& X \\
1 & 2 & 3 \\
-6.500 & 2.000 & 0.000 \\
& \(X\) & \\
1 & 2 & 3 \\
-8.000 & 2.000 & 0.500
\end{tabular}

\section*{LFDRG}

Computes the determinant of a real general matrix given the \(L U\) factorization of the matrix.

\section*{Required Arguments}
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing the \(L U\) factorization of the matrix A as output from routine LFTRG/DLFTRG or LFCRG/DLFCRG. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization as output from routine LFTRG/DLFTRG or LFCRG/DLFCRG. (Input).

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\mathrm{DET} 2}\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDRG (FACT, IPVT, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDRG and D_LFDRG.

\section*{FORTRAN 77 Interface}

Single: CALL LFDRG (N, FACT, LDFACT, IPVT, DET1, DET2)
Double: The double precision name is DLFDRG.

\section*{Description}

Routine LFDRG computes the determinant of a real general coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCRG or LfTrg. The formula \(\operatorname{det} A=\operatorname{det} L \operatorname{det} U\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements
\[
\operatorname{det} U=\prod_{i=1}^{N} U_{i i}
\]
(The matrix \(U\) is stored in the upper triangle of FACT.) Since \(L\) is the product of triangular matrices with unit diagonals and of permutation matrices, \(\operatorname{det} L=(-1)^{k}\) where \(k\) is the number of pivoting interchanges.

Routine LFDRG is based on the LINPACK routine SGEDI; see Dongarra et al. (1979)

\section*{Example}

The determinant is computed for a real general \(3 \times 3\) matrix.
```

USE LFDRG INT
USE LFTRG_INT
USE UMACH_INT
! - Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT(N), NOUT
REAL A(LDA,LDA), DET1, DET2, FACT (LDFACT,LDFACT)
!
!
!
!
!
DATA A/33.0, -24.0, 18.0, 16.0, -10.0, -11.0, 72.0, -57.0, 7.0/
!
CALL LFTRG (A, FACT, IPVT)
CALL LFDRG (FACT, IPVT, DET1, DET2)
Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is ', F6.3, ' * 10**', F2.0)
END

```

\section*{Output}

The determinant of \(A\) is \(-4.761 * 10 * * 3\).

\section*{LINRG}

\section*{岛MPI}

Computes the inverse of a real general matrix.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the matrix to be inverted. (Input)

AINV - N by N matrix containing the inverse of A. (Output)
If A is not needed, A and AINV can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Order of the matrix A. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDAINV — Leading dimension of AINV exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAINV \(=\operatorname{size}(\) AINV, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LINRG (A, AINV [,...])

Specific: The specific interface names are S_LINRG and D_LINRG.

\section*{FORTRAN 77 Interface}

Single: CALL LINRG (N, A, LDA, AINV, LDAINV)
Double: The double precision name is DLINRG.

\section*{ScaLAPACK Interface}

Generic: CALL LINRG (A0, AINV0 [,...])
Specific: The specific interface names are S_LINRG and D_LINRG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LINRG computes the inverse of a real general matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LINRG first uses the routine LFCRG to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. Routine LFCRG computes \(U\) and the information needed to compute \(L^{-1}\). LINRT is then used to compute \(U^{-1}\). Finally, \(A^{-1}\) is computed using \(A^{-1}=U^{-1} L^{-1}\).

The routine LINRG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. This error occurs only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in \(A^{-1}\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2NRG/DL2NRG. The reference is:

CALL L2NRG (N, A, LDA, AINV, LDAINV, WK, IWK)
The additional arguments are as follows:
\(W K-\) Work vector of length \(N+N(N-1) / 2\).
IWK - Integer work vector of length N.
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The inverse might not be accurate.
42 The input matrix is singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix to be inverted. (Input)

AINVO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix AINV. AINV contains the inverse of the matrix A. (Output)
If A is not needed, A and AINV can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the \(\bar{S} c a L A P A C K ~ E x a m p l e ~ b e l o w . ~\)

\section*{Example}

The inverse is computed for a real general \(3 \times 3\) matrix.
```

USE LINRG INT
USE WRRRN_INT
PARAMETER (LDA=3, LDAINV=3)

```
```

        INTEGER I, J, NOUT
        REAL A(LDA,LDA), AINV (LDAINV,LDAINV)
        Set values for A
        A = ( 1.0 3.0 3.0)
        ( 1.0 3.0 4.0)
        ( 1.0 4.0 3.0)
        DATA A/1.0, 1.0, 1.0, 3.0, 3.0, 4.0, 3.0, 4.0, 3.0/
    CALL LINRG (A, AINV)
    CALL WRRRN ('AINV', AINV)
        END
    ```

\section*{Output}
```

                AINV
    ```
\begin{tabular}{rrrrr} 
& 1 & 2 & 3 \\
1 & 7.000 & -3.000 & -3.000 \\
2 & -1.000 & 0.000 & 1.000 \\
3 & -1.000 & 1.000 & 0.000
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LINRG_INT
USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER LDA, LDAINV, N, DESCA(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), AINV(:,:)
REAL, ALLOCATABLE :: AO(:,:), AINVO(:,:)
PARAMETER (LDA=3, LDAINV=3, N=3)
Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. O) THEN
ALLOCATE (A (LDA,N), AINV(LDAINV,N))
Set values for A
A(1,:) = (/ 1.0, 3.0, 3.0/)
A(2,:) = (/ 1.0, 3.0, 4.0/)
A(3,:) = (/ 1.0, 4.0, 3.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT

```
```

CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), AINV0 (MXLDA,MXCOL))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Get the inverse
CALL LINRG (A0, AINVO)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK UNMAP(AINVO, DESCA, AINV)
Print results
Only Rank=0 has the solution, AINV.
IF(MP_RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (MP_RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE (A0, AINVO)
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

\section*{AINV}
\begin{tabular}{rrrrr} 
& \multicolumn{3}{c}{ AINV } \\
& & 1 & 2 & 3 \\
1 & 7.000 & -3.000 & -3.000 \\
2 & -1.000 & 0.000 & 1.000 \\
3 & -1.000 & 1.000 & 0.000
\end{tabular}

\section*{LSACG}

Solves a complex general system of linear equations with iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficients of the linear system. (Input)
B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input) IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSACG (A, B, X [,...])

Specific: The specific interface names are S_LSACG and D_LSACG.

\section*{FORTRAN 77 Interface}

Single: CALL LSACG (N, A, LDA, B, IPATH, X)
Double: The double precision name is DLSACG.

\section*{ScaLAPACK Interface}

Generic: CALL LSACG (A0, B0, X0 [, ...])
Specific: The specific interface names are S_LSACG and D_LSACG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSACG solves a system of linear algebraic equations with a complex general coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSACG first uses the routine LFCCG to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using the iterative refinement routine LFICG.

LSACG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSACG solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ACG/DL2ACG. The reference is:

CALL L2ACG (N, A, LDA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
FACT - Complex work vector of length \(\mathrm{N}^{2}\) containing the \(L U\) factorization of A on output.

IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(\boldsymbol{W} \boldsymbol{K}\) - Complex work vector of length N.
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ACG the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2); respectively, in LSACG. Additional memory allocation for FACT and option value restoration are done automatically in LSACG. Users directly calling L2ACG can allocate additional space for \(\operatorname{FACT}\) and set \(\operatorname{IVAL}(3)\) and \(\operatorname{IVAL}(4)\) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSACG or L2ACG. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSACG temporarily replaces IVAL(2) by IVAL(1). The routine L2CCG computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CCG skips this computation. LSACG restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
\(\boldsymbol{A O}\) - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficients of the linear system. (Input)

B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)

X0 - Complex local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has complex general form and the right-hand-side vector \(b\) has three elements.
```

USE LSACG_INT
USE WRCRN_INT
! Declare variables
PARAMETER (LDA=3, N=3)
COMPLEX A(LDA,LDA), B (N), X(N)
Set values for A and B
A = ( 3.0-2.0i 2.0+4.0i 0.0-3.0i)
( 1.0+1.0i 2.0-6.0i 1.0+2.0i)
( 4.0+0.0i -5.0+1.0i 3.0-2.0i)
B = (10.0+5.0i 6.0-7.0i -1.0+2.0i)
DATA A/ (3.0,-2.0), (1.0,1.0), (4.0,0.0), (2.0,4.0), (2.0,-6.0), \&
(-5.0,1.0),(0.0,-3.0), (1.0,2.0), (3.0,-2.0)/
DATA B/ (10.0,5.0), (6.0,-7.0), (-1.0,2.0)/
Solve AX = B (IPATH = 1)
CALL LSACG (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}
( \(1.000,-1.000)^{1}\)
\((2.000,1.000)^{2}\)\(\left(^{3}\right.\)

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has complex general form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSACG_INT
USE WRCRN - INT
USE SCALAP\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:)
COMPLEX, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. O) THEN
ALLOCATE (A(LDA,N), B(N), X(N))
Set values for A and B
A(1,:) = (/ (3.0, -2.0), (2.0, 4.0), (0.0, -3.0)/)
A(2,:) = (/ (1.0, 1.0), (2.0, -6.0), (1.0, 2.0)/)
A(3,:) = (/ (4.0, 0.0), (-5.0, 1.0), (3.0, -2.0)/)
B = (/ (10.0, 5.0), (6.0, -7.0), (-1.0, 2.0)/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXTT, MXLDA, INFO)
A
ALLOCATE (AO (MXLDA,MXCOL), B0 (MXLDA), X0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK_MAP(B, DESCX, B0)
Solve the system of equations
CALL LSACG (A0, B0, XO)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(X0, DESCX, X)
Print results
Only Rank=0 has the solution, X.
IF(MP_RANK .EQ. O)CALL WRCRN ('X', X, 1, N, 1)

```
```

    IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, X0)
    ! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP(`FINAL')
END

```

\section*{Output}

X
```

(1.000,-1.000) (2.000, 1.000) ( 0.000, 3.000)

```

\section*{LSLCG}

Solves a complex general system of linear equations without iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficients of the linear system. (Input)

B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1 ).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved
Default: IPATH \(=1\)

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCG (A, B, X [, ...])

Specific: The specific interface names are _ LSLCG \(^{\text {and D_LSLCG. }}\)

\section*{FORTRAN 77 Interface}

Single: CALL LSLCG (N, A, LDA, B, IPATH, X)
Double: The double precision name is DLSLCG.

\section*{ScaLAPACK Interface}

Generic: CALL LSLCG (A0, B0, XO [, ...])

Specific: The specific interface names are __LSLCG \(^{\text {and }}\) D_LSLCG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLCG solves a system of linear algebraic equations with a complex general coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSLCG first uses the routine LFCCG to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using LFSCG.

LSLCG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This occurs only if \(A\) either is a singular matrix or is very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSACG be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LCG/DL2LCG. The reference is:
```

CALL L2LCG (N, A, LDA, B, IPATH, X, FACT, IPVT, WK)

```

The additional arguments are as follows:
\(\boldsymbol{F A C T}-\mathrm{N} \times \mathrm{N}\) work array containing the \(L U\) factorization of A on output. If A is not needed, A and FACT can share the same storage locations.

IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(\boldsymbol{W} \boldsymbol{K}\) - Complex work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LCG the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2); respectively, in LSLCG. Additional memory allocation for FACT and option value restoration are done automatically in LSLCG. Users directly calling L2LCG can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLCG or L2LCG. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLCG temporarily replaces IVAL(2) by IVAL(1). The routine L2CCG computes the condition number if IVAL \((2)=2\). Otherwise L2CCG skips this computation. LSLCG restores the option. Default values for the option \(\operatorname{are} \operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficients of the linear system. (Input)

B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has complex general form and the right-hand-side vector \(b\) has three elements.
```

USE LSLCG_INT
USE WRCRN_INT
!
PARAMETER (LDA=3, N=3)
Declare variables
COMPLEX A(LDA,LDA), B(N), X(N)
Set values for A and B
A = ( 3.0-2.0i 2.0+4.0i 0.0-3.0i)
( 1.0+1.0i 2.0-6.0i 1.0+2.0i)
( 4.0+0.0i -5.0+1.0i 3.0-2.0i)
B=(10.0+5.0i 6.0-7.0i -1.0+2.0i)
DATA A/ (3.0,-2.0), (1.0,1.0), (4.0,0.0), (2.0,4.0), (2.0,-6.0),\&
(-5.0,1.0), (0.0,-3.0), (1.0,2.0), (3.0,-2.0)/
DATA B/ (10.0,5.0), (6.0,-7.0), (-1.0,2.0)/
Solve AX = B (IPATH = 1)
CALL LSLCG (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}

X
\((1.000,-1.00)^{1} \quad(2.000,1.000)^{2} \quad(0.000,3.000)^{3}\)

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has complex general form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LSLC\overline{GG INT}
USE WRCRN_INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:)
COMPLEX, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ- 0) THEN
A\overline{LLOCATE (A (LDA,N), B(N), X(N))}
Set values for A and B

```
!
```

        A(1,:) = (/ (3.0, -2.0), (2.0, 4.0), (0.0, -3.0)/)
        A(2,:) = (/ (1.0, 1.0), (2.0, -6.0), (1.0, 2.0)/)
        A(3,:) = (/ (4.0, 0.0), (-5.0, 1.0), (3.0, -2.0)/)
        B = (/ (10.0, 5.0), (6.0, -7.0), (-1.0, 2.0)/)
    ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Sēt up the array descriptors
CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), X0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK MAP(A, DESCA, AO)
CALL SCALAPACK_MAP(B, DESCX, BO)
Solve the system of equations
CALL LSLCG (A0, B0, X0)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(X0, DESCX, X)
Print results.
Only Rank=0 has the solution, X.
IF(MP RANK .EQ. O)CALL WRCRN ('X', X, 1, N, 1)
IF (MP_RANK .EQ. O) DEALLOCATE(A, B, X)
DEALLOC}ATE(A0, B0, X0
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}
x
\((1.000,-1.000)^{1} \quad(2.000,1.000)^{2} \quad(0.000,3.000)^{3}\)

\section*{LFCCG \\ 勻MPI}

Computes the \(L U\) factorization of a complex general matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-\) Complex N by N matrix to be factored. (Input)
\(\boldsymbol{F A C T}\) - Complex n by n matrix containing the \(L U\) factorization of the matrix A (Output) If \(A\) is not needed, \(A\) and FACT can share the same storage locations)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)

RCOND - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCCG (A, FACT, IPVT, RCOND [,...])

Specific: The specific interface names are \(S\) _LFCCG and D_LFCCG.

\section*{FORTRAN 77 Interface}

Single: CALL LFCCG (N, A, LDA, FACT, LDFACT, IPVT, RCOND)

Double: The double precision name is DLFCCG.

\section*{ScaLAPACK Interface}

Generic: CALL LFCCG (A0, FACT0, IPVT0, RCOND [, ...])
Specific: The specific interface names are \(S_{\text {_ }}\) LFCCG and D_LFCCG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCCG performs an \(L U\) factorization of a complex general coefficient matrix. It also estimates the condition number of the matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm.
The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.
LFCCG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) either is singular or is very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with routines LFICG, LFSCG and LFDCG. To solve systems of equations with multiple right-hand-side vectors, use LFCCG followed by either LFICG or LFSCG called once for each right-hand side. The routine LFDCG can be called to compute the determinant of the coefficient matrix after LFCCG has performed the factorization.
Let \(F\) be the matrix FACT and let \(p\) be the vector IPVT. The triangular matrix \(U\) is stored in the upper triangle of \(F\). The strict lower triangle of \(F\) contains the information needed to reconstruct \(L\) using
\[
L^{11}=L_{N-1} P_{N-1} \ldots L_{1} P_{1}
\]
where \(P_{k}\) is the identity matrix with rows \(k\) and \(p_{k}\) interchanged and \(L_{k}\) is the identity with \(F_{i k}\) for \(i=k+1, \ldots, \mathrm{~N}\) inserted below the diagonal. The strict lower half of \(F\) can also be thought of as containing the negative of the multipliers.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CCG/DL2CCG. The reference is:

CALL L2CCG (N, A, LDA, FACT, LDFACT, IPVT, RCOND, WK)
The additional argument is:
WK - Complex work vector of length N .
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
42 The input matrix is singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the matrix to be factored. (Input)

FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(L U\) factorization of the matrix A. (Output)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFCCG is called to factor the matrix and to check for singularity or ill-conditioning. LFICG is called to determine the columns of the inverse.
```

USE IMSL_LIBRARIES

```
```

! Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT (N), NOUT
REAL RCOND, THIRD
COMPLEX A(LDA,N), AINV (LDA,N), RJ(N), FACT (LDFACT,N), RES(N)
Declare functions
COMPLEX CMPLX
Set values for A
A = ( 1.0+1.0i 2.0+3.0i 3.0+3.0i)
( 2.0+1.0i 5.0+3.0i 7.0+4.0i)
(-2.0+1.0i -4.0+4.0i -5.0+3.0i)
DATA A/(1.0,1.0), (2.0,1.0), (-2.0,1.0), (2.0,3.0), (5.0,3.0),\&
(-4.0,4.0), (3.0,3.0), (7.0,4.0), (-5.0,3.0)/
Scale A by dividing by three
THIRD = 1.0/3.0
DO 10 I=1, N
CALL CSSCAL (N, THIRD, A(:, I), 1)
1 0 ~ C O N T I N U E ~
CALL LFCCG (A, FACT, IPVT, RCOND)
Print the L1 condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
Set up the columns of the identity

```
```

! matrix one at a time in RJ
CALL CSET (N, (0.0,0.0), RJ, 1)
DO 20 J=1, N
RJ(J) = CMPLX(1.0,0.0)
RJ is the J-th column of the identity
matrix so the following LFIRG
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFICG (A, FACT, IPVT, RJ, AINV(:,J), RES)
RJ(J) = CMPLX(0.0,0.0)
2 0 ~ C O N T I N U E ~
!
CALL WRCRN ('AINV', AINV)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND < . }0

```
L1 Condition number \(<100.0\)

AINV
1223
\((6.400,-2.800)(-3.800,2.600)(-2.600,1.200)\)
\((-1.600,-1.800)(0.200,0.600)(0.400,-0.800)\)
\((-0.600,2.200)(1.200,-1.400)(0.400,0.200)\)

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. LFCCG is called to factor the matrix and to check for singularity or ill-conditioning. LFICG is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFC\overline{CG INT}
USE UMACH_INT
USE LFICG_INT
USE WRCRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALAPAACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
! Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA, NOUT
INTEGER, ALLOCATABLE :: IPVTO(:)
COMPLEX, ALLOCATABLE :: A(:,:), AINV(:,:), X0(:), RJ(:)
COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), RESO(:), RJO(:)
REAL RCOND, THIRD
PARAMETER (LDA=3, N=3)

```
```

! Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N), AINV(LDA,N))
Set values for A
A(1,:) = (/ ( 1.0, 1.0), ( 2.0, 3.0), ( 3.0, 3.0)/)
A(2,:) = (/ ( 2.0, 1.0), ( 5.0, 3.0), ( 7.0, 4.0)/)
A(3,:) = (/ (-2.0, 1.0), (-4.0, 4.0), (-5.0, 3.0)/)
Scale A by dividing by three
THIRD = 1.0/3.0
A = A * THIRD
ENDIF
Set up a 1D processor grid and define
its context id, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Sēt up t\overline{he array descriptors}
CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCL, N, 1, MP_MB, 1, - 0, 0, MP_ICTXTT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), RJ(N), \&
RJO(MXLDA), RESO(MXLDA), IPVTO(MXLDA))
Map input array to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Factor A
CALL LFCCG (A0, FACT0, IPVT0, RCOND)
Print the reciprocal condition number
and the L1 condition number
IF(MP RANK .EQ. 0) THEN
CALL UMACH (2, NOUT)
WRITE (NOUT,99998) RCOND, 1.OEO/RCOND
ENDIF
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0, 0.0)
DO 10 J=1, N
RJ(J) = (1.0, 0.0)
CALL SCALAPACK_MAP(RJ, DESCL, RJO)
RJ is the J-th column of the identity
matrix so the following LFICG
reference computes the J-th column of
the inverse of A
CALL LFICG (A0, FACTO, IPVTO, RJO, X0, RESO)
RJ(J) = (0.0, 0.0)
CALL SCALAPACK_UNMAP(X0, DESCL, AINV(:,J))
10 CONTINUE
Print results
Only Rank=0 has the solution, AINV.
IF(MP RANK.EQ.O) CALL WRCRN ('AINV', AINV)
IF (M\overline{P}}\mathrm{ RANK .EQ. 0) DEALLOCATE (A, AINV)
DEALLOCATE(A0, FACT0, IPVT0, RJ, RJ0, RESO, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)

```
```

! Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
99998 FO\overline{RMAT (' RCON}D=',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < . 02
L1 Condition number < 100.0

AINV
\begin{tabular}{rrrrr}
\(1^{1}\) & \((6.400,-2.800)\) & \((-3.800,2.600)\) & \((-2.600,1.200)^{3}\) \\
2 & \((-1.600,-1.800)\) & \((0.200,0.600)\) & \((0.400,-0.800)\) \\
3 & \((-0.600,2.200)\) & \((1.200,-1.400)\) & \((0.400,0.200)\)
\end{tabular}

\section*{LFTCG}

Computes the \(L U\) factorization of a complex general matrix.

\section*{Required Arguments}

A-Complex n by n matrix to be factored. (Input)
FACT - Complex n by n matrix containing the \(L U\) factorization of the matrix A. (Output) If \(A\) is not needed, A and FACT can share the same storage locations.

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTCG (A, FACT, IPVT [,...])

Specific: The specific interface names are S_LFTCG and D_LFTCG.

\section*{FORTRAN 77 Interface}

Single: CALL LFTCG (N, A, LDA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTCG.

\section*{ScaLAPACK Interface}

Generic: CALL LfTCG (A0, FACT0, IPVTO [, ...])
Specific: The specific interface names are S_LFTCG and D_LFTCG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFTCG performs an \(L U\) factorization of a complex general coefficient matrix. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm.

LFTCG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) either is singular or is very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with routines LFICG, LFSCG and LFDCG. To solve systems of equations with multiple right-hand-side vectors, use LFTCG followed by either LFICG or LFSCG called once for each right-hand side. The routine LFDCG can be called to compute the determinant of the coefficient matrix after LFCCG has performed the factorization.

Let \(F\) be the matrix FACT and let \(p\) be the vector IPVT. The triangular matrix \(U\) is stored in the upper triangle of \(F\). The strict lower triangle of \(F\) contains the information needed to reconstruct L using
\[
L=L_{N-1} P_{N-1} \ldots L_{1} P_{1}
\]
where \(P_{k}\) is the identity matrix with rows \(k\) and \(P_{k}\) interchanged and \(L_{k}\) is the identity with \(F_{i k}\) for \(i=k+1, \ldots, N\) inserted below the diagonal. The strict lower half of \(F\) can also be thought of as containing the negative of the multipliers.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2TCG/DL2TCG. The reference is:

CALL L2TCG (N, A, LDA, FACT, LDFACT, IPVT, WK)
The additional argument is:

WK - Complex work vector of length \(N\).
2. Informational error

Type Code
42 The input matrix is singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:

A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the matrix to be factored. (Input)

FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(L U\) factorization of the matrix A. (Output) If \(A\) is not needed, A and FACT can share the same storage locations.

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the \(\bar{S} c a L A P A C K ~ E x a m p l e ~ b e l o w . ~\)

\section*{Example}

A linear system with multiple right-hand sides is solved. LFTCG is called to factor the coefficient matrix. LFSCG is called to compute the two solutions for the two right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCCG to perform the factorization, and LFICG to compute the solutions.
```

USE LFTCG INT
USE LFSCG_INT
USE WRCRN INT
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT(N)
COMPLEX A(LDA,LDA), B (N,2), X (N,2), FACT (LDFACT,LDFACT)
Set values for A

```
```

!
DATA A/ (1.0,1.0), (2.0,1.0), (-2.0,1.0), (2.0,3.0), (5.0,3.0),\&
(-4.0,4.0),(3.0,-3.0),(7.0,-5.0), (5.0,3.0)/
Set the right-hand sides, B
B = ( 3.0+ 5.0i 9.0+ 0.0i)
( 22.0+10.0i 13.0+ 9.0i)
(-10.0+4.0i 6.0+10.0i)
DATA B/(3.0,5.0), (22.0,10.0), (-10.0,4.0), (9.0,0.0),\&
(13.0,9.0), (6.0,10.0)/
CALL LFTCG (A, FACT, IPVT)
DO 10 J=1, 2
CALL LFSCG (FACT, IPVT, B(:,J), X(:,J))
1 0 CONTINUE
CALL WRCRN ('X', X)
END

```

\section*{Output}
\begin{tabular}{cc}
X & \\
\((1\) & 2 \\
\((1.000,-1.000)\) & \((0.000,2.000)\) \\
\((2.000,4.000)\) & \((-2.000,-1.000)\) \\
\((3.000,0.000)\) & \((1.000,3.000)\)
\end{tabular}

\section*{ScaLAPACK Example}

The same linear system with multiple right-hand sides is solved as a distributed example. LFTCG is called to factor the matrix. LFSCG is called to compute the two solutions for the two right-hand sides. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFTCG_INT
USE LFSCG_}\mp@subsup{}{}{-}\mathrm{ INT
USE WRCRN_INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
INTEGER, ALLOCATABLE :: IPVTO(:)

```
```

    COMPLEX, ALLOCATABLE :: A(:,:), B(:,:), X(:,:), X0(:)
    COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), BO(:)
    PARAMETER (LDA=3, N=3)
        Set up for MPI
    MP_NPROCS = MP_SETUP()
    IF(MP RANK .EQ. 0) THEN
        ALLOCATE (A (LDA,N), B(N,2), X(N,2))
                                    Set values for A and B
        A(1,:) = (/ ( 1.0, 1.0), ( 2.0, 3.0), ( 3.0,-3.0)/)
        A(2,:) = (/ ( 2.0, 1.0), ( 5.0, 3.0), ( 7.0,-5.0)/)
        A(3,:) = (/ (-2.0, 1.0), (-4.0, 4.0), ( 5.0, 3.0)/)
        B(1,:) = (/ ( 3.0, 5.0), ( 9.0, 0.0)/)
        B(2,:) = (/ ( 22.0, 10.0), (13.0, 9.0)/)
        B(3,:) = (/ (-10.0, 4.0), ( 6.0, 10.0)/)
    ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
    CALL SCALAPACK SETUP(N, N, .TRUE., .TRUE.)
    Get the array descriptor entities MXLDA,
    and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
                            Se}t up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP MB, 1, - 0, 0, MP ICTXTT, MXLDA, INFO)
                            Allocate space for the local arrays
    ALLOCATE (AO (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), &
        BO(MXLDA), IPVTO(MXLDA))
            Map input array to the processor grid
    CALL SCALAPACK MAP(A, DESCA, A0)
                            Factor A
    CALL LFTCG (A0, FACTO, IPVTO)
                                    Solve for the two right-hand sides
    DO 10 J=1, 2
        CALL SCALAPACK MAP(B(:,J), DESCL, BO)
        CALL LFSCG (FACTOO, IPVTO, B0, X0)
        CALL SCALAPACK_UNMAP(X0, DESCL, X(:,J))
    10 CONTINUE
    Print results.
    Only Rank=O has the solution, X.
    IF(MP RANK.EQ.O) CALL WRCRN ('X', X)
    IF (MP_RANK .EQ. 0) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, FACT0, IPVTO, XO)
                            Exit ScaLAPACK usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
        Shut down MPI
    MP_NPROCS = MP_SETUP('FINAL')
    END
    ```

\section*{Output}

X
\(1 \quad(1.000,-1.000) \quad(0.000,2.000)\)
```

2(2.000, 4.000) (-2.000,-1.000)
3(3.000, 0.000) ( 1.000, 3.000)

```

\section*{LFSCG}

Solves a complex general system of linear equations given the \(L U\) factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - Complex n by n matrix containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input) Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input) Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).

IPATH - Path indicator. (Input) IPATH \(=1\) means the system \(A X=B\) is solved. IPATH \(=2\) means the system \(A^{H} \mathrm{X}=\mathrm{B}\) is solved. Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSCG (FACT, IPVT, B, X [, ...])
Specific: The specific interface names are \(S_{\text {_L }}\) LFSCG and D_LFSCG.

\section*{FORTRAN 77 Interface}

Single: CALL LFSCG (N, FACT, LDFACT, IPVT, B, IPATH, X)

Double: The double precision name is DLFSCG.

\section*{ScaLAPACK Interface}

Generic: CALL LFSCG (FACTO, IPVTO, BO, XO [, ...])

Specific: The specific interface names are S_LFSCG and D_LFSCG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFSCG computes the solution of a system of linear algebraic equations having a complex general coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCG or LFTCG. The solution to \(A x=b\) is found by solving the triangular systems \(L y=b\) and \(U x=y\). The forward elimination step consists of solving the system \(L y=b\) by applying the same permutations and elimination operations to \(b\) that were applied to the columns of \(A\) in the factorization routine. The backward substitution step consists of solving the triangular system \(U x=y\) for \(x\).

Routines LFSCG and LFICG both solve a linear system given its \(L U\) factorization. LFICG generally takes more time and produces a more accurate answer than LFSCG. Each iteration of the iterative refinement algorithm used by LFICG calls LFSCG.

The underlying code is based on either LINPACK , LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. FACT contains the \(L U\) factorization of the matrix A. (Input)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization as output from subroutine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)

B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector \(\mathrm{X} . \mathrm{X}\) contains the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse is computed for a complex general \(3 \times 3\) matrix. The input matrix is assumed to be well-conditioned, hence LFTCG is used rather than LFCCG.

USE IMSL_LIBRARIES
```

! Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT (N)
REAL THIRD
COMPLEX A(LDA,LDA), AINV (LDA,LDA), RJ(N), FACT(LDFACT,LDFACT)
Declare functions
COMPLEX CMPLX
Set values for A
A = ( 1.0+1.0i 2.0+3.0i 3.0+3.0i)
( 2.0+1.0i 5.0+3.0i 7.0+4.0i)
( -2.0+1.0i -4.0+4.0i -5.0+3.0i)
DATA A/ (1.0.1.0), (2.0,1.0), (-2.0,1.0), (2.0,3.0), (5.0,3.0),\&
(-4.0,4.0),(3.0,3.0), (7.0,4.0), (-5.0,3.0)/
THIRD = 1.0/3.0
DO 10 I=1, N
CALL CSSCAL (N, THIRD, A(:,I), 1)
10 CONTINUE
CALL LFTCG (A, FACT, IPVT)
Set up the columns of the identity
matrix one at a time in RJ
CALL CSET (N, (0.0,0.0), RJ, 1)
DO 20 J=1, N
RJ(J) = CMPLX(1.0,0.0)
RJ is the J-th column of the identity
matrix so the following LFSCG
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFSCG (FACT, IPVT, RJ, AINV(:,J))
RJ(J) = CMPLX(0.0,0.0)
CONTINUE
CALL WRCRN ('AINV', AINV)
END

```

\section*{Output}

AINV
\begin{tabular}{rrrr}
\((6.400,-2.800)^{1}\) & \((-3.800,2.600)^{3}\) & \((-2.600,1.200)^{3}\) \\
\((-1.600,-1.800)\) & \((0.200,0.600)\) & \((0.400,-0.800)\) \\
\((-0.600,2.200)\) & \((1.200,-1.400)\) & \((0.400,0.200)\)
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. The input matrix is assumed to be well-conditioned, hence LFTCG is used rather than LFCCG. LFSCG is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFTCG_INT
USE LFSCG_INT
USE WRCRN -}\mp@subsup{}{}{-
USE SCALAPAACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
INTEGER, ALLOCATABLE :: IPVTO(:)
COMPLEX, ALLOCATABLE :: A(:,:), AINV(:,:), X0(:)
COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), RJ(:), RJO(:)
REAL THIRD
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N), AINV (LDA,N))
Set values for A
A(1,:) = (/ ( 1.0, 1.0), ( 2.0, 3.0), ( 3.0, 3.0)/)
A(2,:)=(/ ( 2.0, 1.0), ( 5.0, 3.0), ( 7.0, 4.0)/)
A(3,:)=(/ (-2.0, 1.0), (-4.0, 4.0), (-5.0, 3.0)/)
Scale A by dividing by three
THIRD = 1.0/3.0
A = A * THIRD
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Sēt up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXTT, MXLDA, INFO)
\overline{Allocate space for the local arrays}
ALLOCATE (A0 (MXLDA, MXCOL), X0 (MXLDA), FACTO (MXLDA,MXCOL), RJ(N), \&
RJ0 (MXLDA), IPVTO (MXLDA))

```
```

! Map input array to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Factor A
CALL LFTCG (A0, FACTO, IPVTO)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0, 0.0)
DO 10 J=1, N
RJ(J) = (1.0, 0.0)
CALL SCALAPACK_MAP(RJ, DESCL, RJO)
RJ is the J-th column of the identity
matrix so the following LFICG
reference computes the J-th column of
the inverse of A
CALL LFSCG (FACTO, IPVTO, RJO, XO)
RJ(J) = (0.0, 0.0)
CALL SCALAPACK_UNMAP(X0, DESCL, AINV(:,J))
10 CONTINUE
Print results.
Only Rank=0 has the solution, AINV.
IF(MP RANK.EQ.O) CALL WRCRN ('AINV', AINV)
IF (M\overline{P}}\mathrm{ RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE(A0, FACTO, IPVT0, RJ, RJ0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP(`FINAL')
END

```

\section*{Output}

AINV
123
\((6.400,-2.800)(-3.800,2.600)(-2.600,1.200)\)
\((-1.600,-1.800)(0.200,0.600)(0.400,-0.800)\)
\((-0.600,2.200)(1.200,-1.400)(0.400,0.200)\)

\section*{LFICG}

\section*{- MPI \\ CAPABLE}

Uses iterative refinement to improve the solution of a complex general system of linear equations.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the linear system. (Input)
FACT - Complex n by n matrix containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length n containing the solution to the linear system. (Output)
\(\boldsymbol{R E S}\) - Complex vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\) size (A, 1 ).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFICG (A, FACT, IPVT, B, X, RES [, ...])
Specific: The specific interface names are \(S_{\text {_LFICG }}\) and \(D_{-}\)LFICG.

\section*{FORTRAN 77 Interface}

Single: CALL LFICG (N, A, LDA, FACT, LDFACT, IPVT, B, IPATH, X, RES)
Double: The double precision name is DLFICG.

\section*{ScaLAPACK Interface}

Generic: CALL LFICG (A0, FACT0, IPVT0, B0, X0, RESO [,...])
Specific: The specific interface names are S_LFICG and D_LFICG.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFICG computes the solution of a system of linear algebraic equations having a complex general coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCG, or LFTCG.

Iterative refinement fails only if the matrix is very ill-conditioned. Routines LFICG and LFSCG both solve a linear system given its \(L U\) factorization. LFICG generally takes more time and produces a more accurate answer than LFSCG. Each iteration of the iterative refinement algorithm used by LFICG calls LfSCG.

\section*{Comments}

Informational error
Type Code
32 The input matrix is too ill-conditioned for iterative refinement to be effective

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the linear system. (Input)

FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT as output from routine LFCCG or LFTCG. FACT contains the \(L U\) factorization of the matrix A. (Input)

IPVT0 - Local vector of length MXLDA containing the local portions of the distributed vector IPVT. IPVT contains the pivoting information for the \(L U\) factorization as output from subroutine LFCCG or LFTCG. (Input)

B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector \(\mathrm{X} . \mathrm{X}\) contains the solution to the linear system. (Output)

RESO - Complex local vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the final correction at the improved solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding \(0.5+0.5 i\) to the second element.
```

    USE LFICG_INT
    USE LFCCG INT
    USE WRCRN INT
    USE UMACH_INT
    PARAMETER (LDA=3, LDFACT=3, N=3)
    INTEGER IPVT (N), NOUT
    REAL RCOND
    COMPLEX A(LDA,LDA), B(N), X(N), FACT(LDFACT,LDFACT), RES (N)
        Declare functions
    COMPLEX CMPLX
        Set values for A
        A = ( 1.0+1.0i 2.0+3.0i 3.0-3.0i)
        ( 2.0+1.0i 5.0+3.0i 7.0-5.0i)
        ( -2.0+1.0i -4.0+4.0i 5.0+3.0i)
    DATA A/ (1.0,1.0), (2.0,1.0), (-2.0,1.0), (2.0,3.0), (5.0,3.0), &
        (-4.0,4.0), (3.0,-3.0), (7.0,-5.0), (5.0,3.0)/
        Set values for B
        B = ( 3.0+5.0i 22.0+10.0i -10.0+4.0i)
    DATA B/(3.0,5.0), (22.0,10.0), (-10.0,4.0)/
        Factor A
    CALL LFCCG (A, FACT, IPVT, RCOND)
        Print the L1 condition number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) RCOND, 1.0EO/RCOND
        Solve the three systems
    DO 10 J=1, 3
    CALL LFICG (A, FACT, IPVT, B, X, RES)
            Print results
        CALL WRCRN ('X', X, 1, N, 1)
                            Perturb B by adding 0.5+0.5i to B(2)
        B(2) = B(2) + CMPLX (0.5,0.5)
    10 CONTINUE
    99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```
!

\section*{Output}
```

RCOND < 0.025
L1 Condition number < 75.0
X
1 2 3
( 1.000,-1.000) ( 2.000, 4.000) ( 3.000, 0.000)
X
1 2 3
(0.910,-1.061) (1.986, 4.175) ( 3.123, 0.071)
X
1 2 3
( 0.821,-1.123) ( 1.972, 4.349) ( 3.245, 0.142)

```

\section*{ScaLAPACK Example}

The same set of linear systems is solved successively as a distributed example. The right-handside vector is perturbed after solving the system each of the first two times by adding \(0.5+0.5 i\) to the second element. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFICG_INT
USE LFCCG - INT
USE WRCRN-INT
USE UMACH_INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA, NOUT
INTEGER, ALLOCATABLE :: IPVTO(:)
COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:), XO(:), RES(:)
COMPLEX, ALLOCATABLE :: A0(:,:), FACTO(:,:), BO(:), RESO(:)
REAL RCOND
PARAMETER (LDA=3, N=3)
Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N), B(N), X(N), RES(N))
Set values for A and B
A(1,:) = (/ ( 1.0, 1.0), ( 2.0, 3.0), ( 3.0, 3.0)/)
A(2,:) = (/ ( 2.0, 1.0), ( 5.0, 3.0), ( 7.0, 4.0)/)
A(3,:) = (/ (-2.0, 1.0), (-4.0, 4.0), (-5.0, 3.0)/)
B = (/ (3.0, 5.0), (22.0, 10.0), (-10.0, 4.0)/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL

```
```

    CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
                    Set up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
                            A
    ALLOCATE (AO (MXLDA,MXCOL), XO (MXLDA), FACTO (MXLDA,MXCOL), &
        B0 (MXLDA), IPVTO (MXLDA), RESO (MXLDA))
                            Map input array to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, A0)
            Factor A
    CALL LFCCG (A0, FACTO, IPVTO, RCOND)
                            Print the L1 condition number
    IF (MP_RANK .EQ. O) THEN
        CAL\overline{L UMACH (2, NOUT)}
        WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
    ENDIF
    ! Solve the three systems
DO 10 J=1, 3
CALL SCALAPACK MAP(B, DESCL, BO)
CALL LFICG (AO, FACTO, IPVTO, BO, X0, RESO)
CALL SCALAPACK_UNMAP (X0, DESCL, X)
Print results
Only Rank=0 has the solution, X.
IF (MP_RANK .EQ. O) CALL WRCRN ('X', X, 1, N, 1)
Perturb B by adding 0.5+0.5i to B(2)
IF(MP_RANK .EQ. 0) B (2) = B (2) + (0.5,0.5)
10 CONTINUE
IF (MP RANK .EQ. O) DEALLOCATE (A, B, X, RES)
DEALLOCATE (A0, B0, FACT0, IPVT0, X0, RES0)
Exit Scalapack usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
99999 FORMAT (' RCONDD = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.025
L1 Condition number \(<75.0\)
X
\((1.000,-1.000)^{1} \quad(2.000,4.000)^{2} \quad(3.000,0.000)^{3}\)

X
\((0.910,-1.061)^{1} \quad(1.986,4.175)^{2} \quad(3.123,0.071)^{3}\)

X
\((0.821,-1.123)^{1} \quad(1.972,4.349)^{2}(3.245,0.142)^{3}\)

\section*{LFDCG}

Computes the determinant of a complex general matrix given the \(L U\) factorization of the matrix.

\section*{Required Arguments}

FACT - Complex \(n\) by n matrix containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCCG/DLFCCG or LFTCG/DLFTCG. (Input)

DET1 - Complex scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq|\operatorname{DET} 1|<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET} 1 * 10^{\mathrm{DET}}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: \(\operatorname{LDFACT}=\operatorname{size}(\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDCG (FACT, IPVT, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDCG and D_LFDCG.

\section*{FORTRAN 77 Interface}

Single: CALL LFDCG (N, FACT, LDFACT, IPVT, DET1, DET2)
Double: The double precision name is DLFDCG.

\section*{Description}

Routine LFDCG computes the determinant of a complex general coefficient matrix. To compute the determinant the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCG or LFTCG. The formula \(\operatorname{det} A=\operatorname{det} L \operatorname{det} U\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} U=\prod_{i=1}^{N} U_{i i}
\]
(The matrix \(U\) is stored in the upper triangle of FACT.) Since \(L\) is the product of triangular matrices with unit diagonals and of permutation matrices, \(\operatorname{det} L=(-1)^{k}\) where \(k\) is the number of pivoting interchanges.

LFDCG is based on the LINPACK routine CGEDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a complex general \(3 \times 3\) matrix.
```

        USE LFDCG INT
        USE LFTCG_INT
        USE UMACH_INT
    ! - Declare variables
PARAMETER (LDA=3, LDFACT=3, N=3)
INTEGER IPVT (N), NOUT
REAL DET2
COMPLEX A(LDA,LDA), FACT(LDFACT,LDFACT), DET1
Set values for A
A = ( 3.0-2.0i 2.0+4.0i 0.0-3.0i)
( 1.0+1.0i 2.0-6.0i 1.0+2.0i)
( 4.0+0.0i -5.0+1.0i 3.0-2.0i)
DATA A/ (3.0,-2.0), (1.0,1.0), (4.0,0.0), (2.0,4.0), (2.0,-6.0),\&
(-5.0,1.0), (0.0,-3.0), (1.0,2.0), (3.0,-2.0)/
CALL LFTCG (A, FACT, IPVT)
Compute the determinant for the
factored matrix
CALL LFDCG (FACT, IPVT, DET1, DET2)
Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is',3X,' (',F6.3,',',F6.3,\&
r) * 10**',F2.0)
END

```

\section*{Output}

The determinant of \(A\) is \((0.700,1.100) * 10 * * 1\).

\section*{LINCG}


Computes the inverse of a complex general matrix.

\section*{Required Arguments}
\(\boldsymbol{A}-\) Complex n by N matrix containing the matrix to be inverted. (Input)
AINV - Complex n by n matrix containing the inverse of A. (Output)
If \(A\) is not needed, A and AINV can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\) size (A,2).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDAINV - Leading dimension of AINV exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAINV = size (AINV,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LINCG (A, AINV [,...])

Specific: The specific interface names are \(S_{\text {_LINCG }}\) and \(D_{-}\)LINCG.

\section*{FORTRAN 77 Interface}

Single: CALL LINCG (N, A, LDA, AINV, LDAINV)
Double: The double precision name is DLINCG.

\section*{ScaLAPACK Interface}

Generic: CALL LINCG (A0, AINV0 [,...])

Specific: The specific interface names are S_LINCG and D_LINCG.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LINCG computes the inverse of a complex general matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

LINCG first uses the routine LFCCG to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. LFCCG computes \(U\) and the information needed to compute \(L\). LINCT is then used to compute U. Finally A is computed using A=UL.

LINCG fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. This errors occurs only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in \(\mathrm{A}^{-1}\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2NCG/DL2NCG. The reference is:

CALL L2NCG (N, A, LDA, AINV, LDAINV, WK, IWK)
The additional arguments are as follows:
\(\boldsymbol{W} \boldsymbol{K}-\) Complex work vector of length \(\mathrm{N}+\mathrm{N}(\mathrm{N}-1) / 2\).
IWK - Integer work vector of length N.
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The inverse might not be accurate.

422 The input matrix is singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:

A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the matrix to be inverted. (Input)

AINVO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix AINV. AINV contains the inverse of the matrix A. (Output) If \(A\) is not needed, A and AINV can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse is computed for a complex general \(3 \times 3\) matrix.
```

    USE LINCG_INT
    USE WRCRN - INT
    USE CSSCAL_INT
    !
PARAMETER (LDA=3, LDAINV=3, N=3)
REAL THIRD
COMPLEX A(LDA,LDA), AINV(LDAINV,LDAINV)
Set values for A
A = ( 1.0+1.0i 2.0+3.0i 3.0+3.0i)
( 2.0+1.0i 5.0+3.0i 7.0+4.0i)
( -2.0+1.0i -4.0+4.0i -5.0+3.0i)
DATA A/ (1.0,1.0), (2.0,1.0), (-2.0,1.0), (2.0,3.0), (5.0,3.0),\&
(-4.0,4.0), (3.0,3.0), (7.0,4.0), (-5.0,3.0)/
Scale A by dividing by three
THIRD = 1.0/3.0
DO 10 I=1, N
CALL CSSCAL (N, THIRD, A(:,I), 1)
1 0 ~ C O N T I N U E ~
CALL LINCG (A, AINV)
CALL WRCRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{rcr} 
& AINV \\
\((6.400,-2.800)\) & \((-3.800,2.600)\) & \((-2.600,1.200)\) \\
\((-1.600,-1.800)\) & \((0.200,0.600)\) & \((0.400,-0.800)\) \\
\((-0.600,2.200)\) & \((1.200,-1.400)\) & \((0.400,0.200)\)
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LINCG_INT
USE WRCRN INT
USE SCALAP\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE `mpif.h'
INTEGER J, LDA, N, DESCA(9)
INTEGER INFO, MXCOL, MXLDA, NPROW, NPCOL
COMPLEX, ALLOCATABLE :: A(:,:), AINV(:,:)

```
```

COMPLEX, ALLOCATABLE :: AO(:,:), AINVO(:,:)
REAL THIRD
PARAMETER (LDA=3, N=3)
Set up for MPI
MP_NPROCS = MP_SETUP()
IF\overline{(MP RANK .EQ-. O) THEN}
ALLOCATE (A (LDA,N), AINV (LDA,N))
Set values for A
A(1,:) = (/ ( 1.0, 1.0), ( 2.0, 3.0), ( 3.0, 3.0)/)
A(2,:) = (/ ( 2.0, 1.0), ( 5.0, 3.0), ( 7.0, 4.0)/)
A(3,:) = (/ (-2.0, 1.0), (-4.0, 4.0), (-5.0, 3.0)/)
Scale A by dividing by three
THIRD = 1.0/3.0
A = A * THIRD
ENDIF
Set up a 1D processor grid and define
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Se}t up th\overline{e} array descriptor
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
A
ALLOCATE (A0 (MXLDA,MXCOL), AINVO (MXLDA,MXCOL))
Map input array to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Factor A
CALL LINCG (A0, AINVO)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(AINVO, DESCA, AINV)
Print results.
Only Rank=0 has the solution, X.
IF(MP_RANK.EQ.O) CALL WRCRN ('AINV', AINV)
IF (MP RANK .EQ. O) DEALLOCATE(A, AINV)
DEALLO\overline{C}ATE(A0, AINV0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

AINV
\begin{tabular}{lrrr} 
& & \\
1 & \((6.400,-2.800)\) & \((-3.800,2.600)\) & \((-2.600,1.200)\) \\
2 & \((-1.600,-1.800)\) & \((0.200,0.600)\) & \((0.400,-0.800)\) \\
3 & \((-0.600,2.200)\) & \((1.200,-1.400)\) & \((0.400,0.200)\)
\end{tabular}

\section*{LSLRT}

\section*{匀 MPI \\ CAPABLE}

Solves a real triangular system of linear equations.

\section*{Required Arguments}

A -N by N matrix containing the coefficient matrix for the triangular linear system. (Input) For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means solve \(\mathrm{AX}=\mathrm{B}, \mathrm{A}\) lower triangular.
\(I P A T H=2\) means solve \(A X=B\), A upper triangular.
IPATH \(=3\) means solve \(A^{T} \mathrm{X}=\mathrm{B}, \mathrm{A}\) lower triangular.
IPATH \(=4\) means solve \(A^{T} \mathrm{X}=\mathrm{B}\), A upper triangular.
Default: ІРАТН \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLRT (A, B, X [, ...])
Specific: The specific interface names are S_LSLRT and D_LSLRT.

\section*{FORTRAN 77 Interface}

Single: CALL LSLRT (N, A, LDA, B, IPATH, X)
Double: The double precision name is DLSLRT.

\section*{ScaLAPACK Interface}

Generic: CALL LSLRT (A0, B0, X0 [,...])

Specific: The specific interface names are S_LSLRT and D_LSLRT.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLRT solves a system of linear algebraic equations with a real triangular coefficient matrix. LSLRT fails if the matrix A has a zero diagonal element, in which case A is singular. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:

A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficients of the linear system. (Input)

For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)

X0 - Local vector of length MXLDA containing the local portions of the distributed vector X. X contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the \(\bar{S} c a L A P A C K ~ E x a m p l e ~ b e l o w . ~\)

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has lower triangular form and the right-hand-side vector, \(b\), has three elements.
```

USE LSLRT_INT
USE WRRRN_INT
PARAMETER (LDA=3)
REAL A(LDA,LDA), B(LDA), X(LDA)
Set values for A and B

```
!
```

!
!
!
!
!
!
DATA A/2.0, 2.0, -4.0, 0.0, -1.0, 2.0, 0.0, 0.0, 5.0/
DATA B/2.0, 5.0, 0.0/
! Solve AX = B (IPATH = 1)
CALL LSLRT (A, B, X)
CALL WRRRN ('X', X, 1, 3, 1)
END

```

\section*{Output}
\begin{tabular}{rrr}
\multicolumn{3}{c}{X} \\
1 & 2 & 3 \\
1.000 & -3.000 & 2.000
\end{tabular}

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has lower triangular form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSLRT_INT
USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), B(:), X(:)
REAL, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=3)
MP NPROCS = MP SETUP()
IF(MP_RANK .EQ. 0) THEN
A\overline{LLOCATE (A (LDA,N), B(N), X(N))}
Set values for A and B
A(1,:) = (/ 2.0, 0.0, 0.0/)
A(2,:) = (/ 2.0, -1.0, 0.0/)
A(3,:) = (/-4.0, 2.0, 5.0/)
B = (/ 2.0, 5.0, 0.0/)
ENDIF

```
!
!
```

! its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Se}t up the array descriptor
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP MB, 1, 0, 0, MP ICTXT, MXLDA, INFO)
A
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), XO (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK_MAP(B, DESCX, BO)
Solve AX = B (IPATH = 1)
CALL LSLRT (A0, B0, X0)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP (X0, DESCX, X)
Print results.
Only Rank=0 has the solution, X.
IF(MP_RANK .EQ. O)CALL WRRRN ('X', X, 1, N, 1)
IF (MP_RANK .EQ. O) DEALLOCATE(A, B, X)
DEALLOCATE (A0, B0, X0)
! Exit Scalapack usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{3}{c}{\(x\)} \\
1 & 2 & 3 \\
1.000 & -3.000 & 2.000
\end{tabular}

\section*{LFCRT}

\section*{MPI \\ CAPABLE}

Estimates the condition number of a real triangular matrix.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the coefficient matrix for the triangular linear system. (Input) For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).
IPATH - Path indicator. (Input)
IPATH \(=1\) means A is lower triangular.
IPATH \(=2\) means A is upper triangular.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCRT (A, RCOND [, ...])

Specific: The specific interface names are S_LFCRT and D_LFCRT.

\section*{FORTRAN 77 Interface}

Single: CALL LFCRT (N, A, LDA, IPATH, RCOND)
Double: The double precision name is DLFCRT.

\section*{ScaLAPACK Interface}

Generic: CALL LFCRT (A0, RCOND [,...])
Specific: The specific interface names are S_LFCRT and D_LFCRT.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCRT estimates the condition number of a real triangular matrix. The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\).

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CRT/ DL2CRT. The reference is:

CALL L2CRT (N, A, LDA, IPATH, RCOND, WK)
The additional argument is:
\(W K\) - Work vector of length N .
2. Informational error

Type Code
31 The input triangular matrix is algorithmically singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
AO- MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix for the triangular linear system. (Input) For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the \(\bar{S} c a L A P A C K ~ E x a m p l e ~ b e l o w . ~\)

\section*{Example}

An estimate of the reciprocal condition number is computed for a \(3 \times 3\) lower triangular coefficient matrix.
```

    USE LFCRT_INT
    USE UMACH_INT
    ! Declare variables
PARAMETER (LDA=3)
REAL A (LDA,LDA), RCOND
INTEGER NOUT
Set values for A and B

```
```

!
!
!
DATA A/2.0, 2.0, -4.0, 0.0, -1.0, 2.0, 0.0, 0.0, 5.0/
Compute the reciprocal condition
number (IPATH=1)
CALL LFCRT (A, RCOND)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0EO/RCOND
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.1
L1 Condition number < 15.0

\section*{ScaLAPACK Example}

The same lower triangular matrix as in the example above is used in this distributed computing example. An estimate of the reciprocal condition number is computed for the \(3 \times 3\) lower triangular coefficient matrix. SCALAPACK_MAP is an IMSL utility routine (see Chapter 11, "Utilities") used to map an array to the processor grid. It is used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFCRT INT
USE SCALAP\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE `mpif.h'
! Declare variables
INTEGER LDA, N, NOUT, DESCA(9)
INTEGER INFO, MXCOL, MXLDA
REAL RCOND
REAL, ALLOCATABLE :: A(:,:)
REAL, ALLOCATABLE :: A0(:,:)
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N))
! Set values for A
A(1,:) = (/ 2.0, 0.0, 0.0/)
A(2,:) = (/ 2.0, -1.0, 0.0/)
A(3,:) = (/-4.0, 2.0, 5.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL

```
!
```

    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
    Set up the array descriptor
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL))
    Map input array to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, A0)
        Compute the reciprocal condition
        number (IPATH=1)
        Print results.
    Only Rank=0 has the solution, RCOND.
    IF(MP_RANK .EQ. O) THEN
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
    ENDIF
    IF (MP_RANK .EQ. 0) DEALLOCATE(A)
    DEALLOCATE (A0)
    ! Exit Scalapack usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.1
L1 Condition number < 15.0

\section*{LFDRT}

Computes the determinant of a real triangular matrix.

\section*{Required Arguments}

A-N by N matrix containing the triangular matrix. (Input) The matrix can be either upper or lower triangular.

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq|\operatorname{DET1}|<10.0\) or DET1 \(=0.0\).

DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(\mathrm{A})=\operatorname{DET} 1 * 10^{\mathrm{DET}} 2\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDRT (A, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDRT and D_LFDRT.

\section*{FORTRAN 77 Interface}

Single: CALL LFDRT (N, A, LDA, DET1, DET2)
Double: The double precision name is DLFDRT.

\section*{Description}

Routine LFDRT computes the determinant of a real triangular coefficient matrix. The determinant of a triangular matrix is the product of the diagonal elements
\[
\operatorname{det} A=\prod_{i=1}^{N} A_{i i}
\]

LFDRT is based on the LINPACK routine STRDI; see Dongarra et al. (1979).

\section*{Comments}

Informational error
Type Code
31 The input triangular matrix is singular.

\section*{Example}

The determinant is computed for a \(3 \times 3\) lower triangular matrix.
```

USE LFDRT INT
USE UMACH INT
PARAMETER (LDA=3)
REAL A(LDA,LDA), DET1, DET2
INTEGER NOUT
Set values for A
A = (rrrar
DATA A/2.0, 2.0, -4.0, 0.0, -1.0, 2.0, 0.0, 0.0, 5.0/
CALL LFDRT (A, DET1, DET2)

```
!
```

! Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
99999 FORMAT (' The determinant of A is ', F6.3, ' * 10**', F2.0)
END

```

\section*{Output}

The determinant of \(A\) is -1.000 * \(10 * * 1\).

\section*{LINRT}

Computes the inverse of a real triangular matrix.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the triangular matrix to be inverted. (Input)
For a lower triangular matrix, only the lower triangular part and diagonal of A are referenced. For an upper triangular matrix, only the upper triangular part and diagonal of A are referenced.

AINV - N by N matrix containing the inverse of A. (Output)
If A is lower triangular, AINV is also lower triangular. If A is upper triangular, AINV is also upper triangular. If A is not needed, A and AINV can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means A is lower triangular.
IPATH \(=2\) means A is upper triangular.
Default: IPATH \(=1\).

LDAINV - Leading dimension of AINV exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAINV \(=\operatorname{size}(\) AINV, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LINRT (A, AINV [,...])

Specific: The specific interface names are S_LINRT and D_LINRT.

\section*{FORTRAN 77 Interface}

Single: CALL LINRT (N, A, LDA, IPATH, AINV, LDAINV)
Double: The double precision name is DLINRT.

\section*{Description}

Routine LINRT computes the inverse of a real triangular matrix. It fails if A has a zero diagonal element.

\section*{Example}

The inverse is computed for a \(3 \times 3\) lower triangular matrix.
```

USE LINRT_INT
USE WRRRN INT
PARAMETER (LDA=3)
REAL A(LDA,LDA), AINV(LDA,LDA)
Set values for A
A = ( 2.0 ( )
(rrr
DATA A/2.0, 2.0, -4.0, 0.0, -1.0, 2.0, 0.0, 0.0, 5.0/
CALL LINRT (A, AINV)
CALL WRRRN ('AINV', AINV)
END

```
!
\(!\)

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{3}{c}{ AINV } \\
& 1 & 2 & 3 \\
1 & 0.500 & 0.000 & 0.000 \\
2 & 1.000 & -1.000 & 0.000 \\
3 & 0.000 & 0.400 & 0.200
\end{tabular}

\section*{LSLCT}

Solves a complex triangular system of linear equations.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the triangular linear system. (Input)
For a lower triangular system, only the lower triangle of A is referenced. For an upper triangular system, only the upper triangle of \(A\) is referenced.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

IPATH - Path indicator. (Input)
IPATH \(=1\) means solve \(A X=B\), A lower triangular
IPATH \(=2\) means solve \(A X=B\), A upper triangular
IPATH \(=3\) means solve \(A^{H} X=B\), A lower triangular
IPATH \(=4\) means solve \(A^{H} X=B\), A upper triangular Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCT (A, B, X [, ...])
Specific: The specific interface names are S_LSLCT and D_LSLCT.

\section*{FORTRAN 77 Interface}

Single: CALL LSLCT (N, A, LDA, B, IPATH, X)
Double: \(\quad\) The double precision name is DLSLCT.

\section*{ScaLAPACK Interface}

Generic: CALL LSLCT (A0, B0, X0 [,...])
Specific: The specific interface names are S_LSLCT and D_LSLCT.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLCT solves a system of linear algebraic equations with a complex triangular coefficient matrix. LSLCT fails if the matrix A has a zero diagonal element, in which case A is singular. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational error
\begin{tabular}{cll} 
Type & Code \\
4 & 1 & The input triangular matrix is singular. Some of its diagonal elements are \\
& near zero.
\end{tabular}

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
\(\boldsymbol{A O}\) - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the triangular linear system. (Input) For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.

B0 - Local complex vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Local complex vector of length MXLDA containing the local portions of the distributed vector \(X . X\) contains the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has lower triangular form and the right-hand-side vector, \(b\), has three elements.
```

USE LSLCT INT
USE WRCRN_INT
INTEGER LDA

```
```

PARAMETER (LDA=3)
COMPLEX A(LDA,LDA), B(LDA), X(LDA)
Set values for A and B
A = ( -3.0+2.0i )
( -2.0-1.0i 0.0+6.0i )
( -1.0+3.0i 1.0-5.0i -4.0+0.0i )
B = (-13.0+0.0i -10.0-1.0i -11.0+3.0i)
DATA A/ (-3.0,2.0), (-2.0,-1.0), (-1.0, 3.0), (0.0,0.0), (0.0,6.0),\&
(1.0,-5.0), (0.0,0.0), (0.0,0.0), (-4.0,0.0)/
DATA B/(-13.0,0.0), (-10.0,-1.0), (-11.0,3.0)/
CALL LSLCT (A, B, X)
Solve AX = B
Print results
CALL WRCRN ('X', X, 1, 3, 1)
END

```

\section*{Output}
\((3.000,2.000)^{1}\)
\((1.000,1.000)^{2}\)\((2.000,0.000)^{3}\)

\section*{ScaLAPACK Example}

The same lower triangular matrix as in the example above is used in this distributed computing example. The system of three linear equations is solved. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCInIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSLCT_INT
USE WRCRN_INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:)
COMPLEX, ALLOCATABLE :: AO(:,:), BO(:), XO(:)
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. O) THEN
A\overline{LLOCATE (A (LDA,N), B(N), X(N))}
Set values for A
A(1,:) = (/ (-3.0, 2.0), (0.0, 0.0), ( 0.0, 0.0)/)
A(2,:) = (/ (-2.0, -1.0), (0.0, 6.0), ( 0.0, 0.0)/)
A(3,:) = (/ (-1.0, 3.0), (1.0, -5.0), (-4.0, 0.0)/)

```
```

!

```
```

        B
    ```
        B
    ENDIF
    ENDIF
    =(/ (-13.0, 0.0), (-10.0, -1.0), (-11.0, 3.0)
    =(/ (-13.0, 0.0), (-10.0, -1.0), (-11.0, 3.0)
    Set up a 1D processor grid and define
    Set up a 1D processor grid and define
    its context ID, MP ICTXT
    its context ID, MP ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
    Get the array descriptor entities MXLDA,
    Get the array descriptor entities MXLDA,
    and MXCOL
    and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
    Set up the array descriptor
    Set up the array descriptor
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
    Allocate space for the local arrays
    Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), XO (MXLDA))
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), XO (MXLDA))
                            Map input arrays to the processor grid
                            Map input arrays to the processor grid
CALL SCALAPACK MAP(A, DESCA, AO)
CALL SCALAPACK MAP(A, DESCA, AO)
CALL SCALAPACK_MAP(B, DESCX, BO)
CALL SCALAPACK_MAP(B, DESCX, BO)
                            Solve AX = B
                            Solve AX = B
CALL LSLCT (A0, B0, X0)
CALL LSLCT (A0, B0, X0)
    Unmap the results from the distributed
    Unmap the results from the distributed
    arrays back to a non-distributed array.
    arrays back to a non-distributed array.
    After the unmap, only Rank=0 has the full
    After the unmap, only Rank=0 has the full
    array.
    array.
    CALL SCALAPACK_UNMAP(X0, DESCX, X)
    CALL SCALAPACK_UNMAP(X0, DESCX, X)
                            Print results.
                            Print results.
                            Only Rank=0 has the solution, X.
                            Only Rank=0 has the solution, X.
    IF(MP RANK .EQ. O) CALL WRCRN ('X', X, 1, 3, 1)
    IF(MP RANK .EQ. O) CALL WRCRN ('X', X, 1, 3, 1)
    IF (MP RANK .EQ. 0) DEALLOCATE (A, B, X)
    IF (MP RANK .EQ. 0) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, X0)
    DEALLOCATE (A0, B0, X0)
! Exit ScaLAPACK usage
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
MP NPROCS = MP SETUP('FINAL')
END
```

END

```

\section*{Output}

\section*{X}
\((3.000,2.00)^{1} \quad(1.000,1.000)^{2}(2.000,0.000)^{3}\)

\section*{LFCCT}

Estimates the condition number of a complex triangular matrix.

\section*{Required Arguments}

A-Complex n by n matrix containing the triangular matrix. (Input)
For a lower triangular system, only the lower triangle of A is referenced. For an upper triangular system, only the upper triangle of A is referenced.
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means A is lower triangular.
IPATH \(=2\) means A is upper triangular.
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCCT (A, RCOND [,...])

Specific: The specific interface names are S_LFCCT and D_LFCCT.

\section*{FORTRAN 77 Interface}

Single: CALL LFCCT (N, A, LDA, IPATH, RCOND)
Double: The double precision name is DLFCCT.

\section*{ScaLAPACK Interface}

Generic: CALL LFCCT (A0, RCOND [,...])

Specific: The specific interface names are S_LFCCT and D_LFCCT.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCCT estimates the condition number of a complex triangular matrix. The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979). If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CCT/DL2CCT. The reference is:
```

CALL L2CCT (N, A, LDA, IPATH, RCOND, CWK)

```

The additional argument is:
CWK - Complex work vector of length n.
2. Informational error
Type Code

31 The input triangular matrix is algorithmically singular.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the triangular linear system. (Input)
For a lower triangular system, only the lower triangular part and diagonal of A are referenced. For an upper triangular system, only the upper triangular part and diagonal of \(A\) are referenced.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

An estimate of the reciprocal condition number is computed for a \(3 \times 3\) lower triangular coefficient matrix.
```

USE LFCCT INT
USE UMACH_

```
```

! Declare variables
INTEGER LDA, N
PARAMETER (LDA=3)
INTEGER NOUT
REAL RCOND
COMPLEX A(LDA,LDA)
!
!
!
DATA A/ (-3.0,2.0), (-2.0,-1.0), (-1.0, 3.0), (0.0,0.0), (0.0,6.0),\&
(1.0,-5.0), (0.0,0.0), (0.0,0.0), (-4.0,0.0)/
Compute the reciprocal condition
number
CALL LFCCT (A, RCOND)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.2
L1 Condition number < 10.0

\section*{ScaLAPACK Example}

The same lower triangular matrix as in the example above is used in this distributed computing example. An estimate of the reciprocal condition number is computed for a \(3 \times 3\) lower triangular coefficient matrix. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFC\overline{CT INT}
USE UMACH_INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
! Declare variables
INTEGER LDA, N, NOUT, DESCA(9)
INTEGER INFO, MXCOL, MXLDA
REAL RCOND
COMPLEX, ALLOCATABLE :: A(:,:)
COMPLEX, ALLOCATABLE :: A0 (:,:)
PARAMETER (LDA=3, N=3)
! Set up for MPI

```
```

    MP NPROCS = MP SETUP()
    IF(MP_RANK .EQ. O) THEN
        ALLOCATE (A (LDA,N))
        A(1,:) = (/ (-3.0, 2.0), (0.0, 0.0), ( 0.0, 0.0)/)
        A(2,:) = (/ (-2.0, -1.0), (0.0, 6.0), ( 0.0, 0.0)/)
        A(3,:)=(/ (-1.0, 3.0), (1.0, -5.0), (-4.0, 0.0)/)
    ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
Set up the array descriptor
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
A
ALLOCATE (A0 (MXLDA,MXCOL))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Compute the reciprocal condition
number
Print results.
Only Rank=0 has the solution, RCOND.
IF (MP_RANK .EQ. O) THEN
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0EO/RCOND
ENDIF
IF (MP RANK .EQ. O) DEALLOCATE (A)
DEALLO\overline{CATE (AO)}
!
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
99999 FORMAT (' RCON}D=',F5.3,/,' L1 Condition number = ',F6.3
END

```

\section*{Output}

RCOND < 0.2
L1 Condition number < 10.0

\section*{LFDCT}

Computes the determinant of a complex triangular matrix.

\section*{Required Arguments}

A-Complex n by n matrix containing the triangular matrix.(Input)
DET1 - Complex scalar containing the mantissa of the determinant. (Output) The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).

DET2 - Scalar containing the exponent of the determinant. (Output) The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET} 1 * 10^{\text {DET2 }}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDCT (A, DET1, DET2 [, ...])
Specific: The specific interface names are S_LFDCT and D_LFDCT.

\section*{FORTRAN 77 Interface}

Single: CALL LFDCT (N, A, LDA, DET1, DET2)
Double: The double precision name is DLFDCT.

\section*{Description}

Routine LFDCT computes the determinant of a complex triangular coefficient matrix. The determinant of a triangular matrix is the product of the diagonal elements
\[
\operatorname{det} A=\prod_{i=1}^{N} A_{i i}
\]

LFDCT is based on the LINPACK routine CTRDI; see Dongarra et al. (1979).

\section*{Comments}

Informational error
Type Code
31 The input triangular matrix is singular.

\section*{Example}

The determinant is computed for a \(3 \times 3\) complex lower triangular matrix.
```

USE LFDCT_INT
USE UMACH_INT
Declare variables
INTEGER LDA, N
PARAMETER (LDA=3,N=3)
INTEGER NOUT

```
```

    REAL DET2
    COMPLEX A(LDA,LDA), DET1
                Set values for A
                A = ( -3.0+2.0i )
                    ( -2.0-1.0i 0.0+6.0i )
                            ( -1.0+3.0i 1.0-5.0i -4.0+0.0i )
    DATA A/ (-3.0,2.0), (-2.0,-1.0), (-1.0, 3.0), (0.0,0.0), (0.0,6.0),&
        (1.0,-5.0),(0.0,0.0),(0.0,0.0), (-4.0,0.0)/
    CALL LFDCT (A, DET1, DET2)
                            Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) DET1, DET2
    99999 FORMAT (' The determinant of A is (',F4.1,',',F4.1,') * 10**',\&
F2.0)
END

```

\section*{Output}

The determinant of \(A\) is ( \(0.5,0.7\) ) * \(10 * * 2\).

\section*{LINCT}

Computes the inverse of a complex triangular matrixs.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex n by n matrix containing the triangular matrix to be inverted. (Input) For a lower triangular matrix, only the lower triangle of \(A\) is referenced. For an upper triangular matrix, only the upper triangle of \(A\) is referenced.

AINV - Complex n by n matrix containing the inverse of A. (Output)
If A is lower triangular, AINV is also lower triangular. If A is upper triangular, AINV is also upper triangular. If A is not needed, A and AINV can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

IPATH - Path indicator. (Input)
IPATH \(=1\) means \(A\) is lower triangular.

IPATH \(=2\) means \(A\) is upper triangular.
Default: \(\operatorname{IPATH}=1\).

LDAINV - Leading dimension of AINV exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAINV \(=\operatorname{size}(\) AINV, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LINCT (A, AINV [,...])
Specific: The specific interface names are \(S_{\text {_LINCT }}\) and D_LINCT.

\section*{FORTRAN 77 Interface}

Single: CALL LINCT (N, A, LDA, IPATH, AINV, LDAINV)
Double: The double precision name is DLINCT.

\section*{Description}

Routine LINCT computes the inverse of a complex triangular matrix. It fails if A has a zero diagonal element.

\section*{Comments}

Informational error
Type Code
41 The input triangular matrix is singular. Some of its diagonal elements are close to zero.

\section*{Example}

The inverse is computed for a \(3 \times 3\) lower triangular matrix.
```

USE LINCT INT
USE WRCRN_INT
! Declare variables
INTEGER LDA
PARAMETER (LDA=3)
COMPLEX A(LDA,LDA), AINV (LDA,LDA)
Set values for A
A = (-3.0+2.0i )
( -2.0-1.0i 0.0+6.0i )
( -1.0+3.0i 1.0-5.0i -4.0+0.0i )
DATA A/ (-3.0,2.0), (-2.0,-1.0), (-1.0, 3.0), (0.0,0.0), (0.0,6.0),\&
(1.0,-5.0),(0.0,0.0),(0.0,0.0), (-4.0,0.0)/

```
```

! Compute the inverse of A
CALL LINCT (A, AINV)
CALL WRCRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{rrrrr} 
& \multicolumn{4}{c}{ AINV } \\
1 & \((-0.2308,-0.1538)\) & \((0.0000,0.0000)\) & \((0.0000,0.0000)\) \\
2 & \((-0.0897,-0.0513)\) & \((0.0000,-0.1667)\) & \((0.0000,0.0000)\) \\
3 & \((0.2147,-0.0096)\) & \((-0.2083,-0.0417)\) & \((-0.2500,0.0000)\)
\end{tabular}

\section*{LSADS}

Solves a real symmetric positive definite system of linear equations with iterative refinement.

\section*{Required Arguments}

A-n by n matrix containing the coefficient matrix of the symmetric positive definite linear system. (Input)
Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSADS (A, B, X [, ...])

Specific: The specific interface names are S_LSADS and D_LSADS.

\section*{FORTRAN 77 Interface}

Single: CALL LSADS (N, A, LDA, B, X)
Double: The double precision name is DLSADS.

\section*{ScaLAPACK Interface}

Generic: CALL LSADS (A0, B0, X0 [, ...])

Specific: The specific interface names are S_LSADS and D_LSADS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSADS solves a system of linear algebraic equations having a real symmetric positive definite coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSADS first uses the routine LFCDS to compute an \(R^{T} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. The matrix \(R\) is upper triangular. The solution of the linear system is then found using the iterative refinement routine LFIDS. LSADS fails if any submatrix of \(R\) is not positive definite, if \(R\) has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is either very close to a singular matrix or a matrix which is not positive definite. If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSADS solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ADS / DL2ADS. The reference is:

CALL L2ADS (N, A, LDA, B, X, FACT, WK)
The additional arguments are as follows:
FACT- Work vector of length \(\mathrm{N}^{2}\) containing the \(R^{T} R\) factorization of A on output.
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is not positive definite.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ADS the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values \(\operatorname{IVAL}(3)\) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSADS. Additional memory allocation for FACT and option value restoration are done automatically in LSADS. Users directly calling L2ADS can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSADS or L2ADS. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).
17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSADS temporarily replaces IVAL(2) by IVAL(1). The routine L2CDS computes the condition number if IVAL \((2)=2\). Otherwise L2CDS skips this computation. LSADS restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the symmetric positive definite linear system. (Input)

B0 - Local vector of length MXIDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Local vector of length MXLDA containing the local portions of the distributed vector X . \(X\) contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the \(\bar{S} c a L A P A C K ~ E x a m p l e ~ b e l o w . ~\)

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has real positive definite form and the right-hand-side vector \(b\) has three elements.
```

    USE LSADS_INT
    USE WRRRN_INT
    !
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
REAL A (LDA,LDA), B (N), X(N)
Set values for A and B
A = ($$
\begin{array}{lll}{1.0}&{-3.0}&{2.0}\end{array}
$$)
(-3.0 10.0 -5.0)
( 2.0 -5.0 6.0)
B =( 27.0 -78.0 64.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
DATA B/27.0, -78.0, 64.0/
CALL LSADS (A, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```
    Output
        X
\(\begin{array}{rrr}1 & 2 & 3 \\ 0 & -4.000 & 7.000\end{array}\)

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has real positive definite form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSADS_INT
USE WRRRN_INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), B(:), X(:)
REAL, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=3)
Set up for MPI

```
```

MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. O) THEN
A\overline{LLOCATE (A (LDA,N), B(N), X(N))}
Set values for A and B
A(1,:) = (/ 1.0, -3.0, 2.0/)
A (2,:) = (/ -3.0, 10.0, -5.0/)
A(3,:) = (/ 2.0, -5.0, 6.0/)
B = (/27.0, -78.0, 64.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA), X0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK_MAP(B, DESCX, BO)
Solve the system of equations
CALL LSADS (A0, B0, X0)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(X0, DESCX, X)
Print results.
Only Rank=0 has the solution, X.
IF(MP_RANK .EQ. 0) CALL WRRRN ('X', X, 1, N, 1)
IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)
DEALLOCATE (A0, B0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK EXIT(MP ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP(`FINAL')
EN\overline{D}

```

\section*{Output}
\begin{tabular}{ccc} 
& X \\
1 & 2 & 3 \\
1.000 & -4.000 & 7.000
\end{tabular}

\section*{LSLDS \\ CAPABLE}

Solves a real symmetric positive definite system of linear equations without iterative refinement .

\section*{Required Arguments}

A—n by n matrix containing the coefficient matrix of the symmetric positive definite linear system. (Input)
Only the upper triangle of A is referenced.
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\) size (A, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLDS (A, B, X [,...])
Specific: The specific interface names are S_LSLDS and D_LSLDS.

\section*{FORTRAN 77 Interface}

Single: CALL LSLDS (N, A, LDA, B, X)
Double: The double precision name is DLSLDS.

\section*{ScaLAPACK Interface}

Generic: CALL LSLDS (A0, B0, X0 [, ...])

Specific: The specific interface names are S_LSLDS and D_LSLDS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLDS solves a system of linear algebraic equations having a real symmetric positive definite coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a
detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSLDS first uses the routine LFCDS to compute an \(R^{T} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. The matrix \(R\) is upper triangular. The solution of the linear system is then found using the routine LFSDS. LSLDS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) either is very close to a singular matrix or to a matrix which is not positive definite. If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned, it is recommended that LSADS be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LDS / DL2LDS. The reference is:

CALL L2LDS (N, A, LDA, B, X, FACT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-\mathrm{N} \times \mathrm{N}\) work array containing the \(R^{T} R\) factorization of A on output. If A is not needed, A can share the same storage locations as FACT.
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is not positive definite.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LDS the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSLDS. Additional memory allocation for FACT and option value restoration are done automatically in LSLDS. Users directly calling L2 LDS can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLDS or L2LDS. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLDS temporarily replaces IVAL(2) by IVAL(1). The routine L2CDS computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CDS skips this computation. LSLDS restores the option. Default values for the option \(\operatorname{are} \operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
AO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the symmetric positive definite linear system. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)

XO - Local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has real positive definite form and the right-hand-side vector \(b\) has three elements.
```

    USE LSLDS_INT
    USE WRRRN INT
    INTEGER LDA, N
    PARAMETER (LDA=3, N=3)
    REAL A(LDA,LDA), B (N), X(N)
        Set values for A and B
        A = ( 1.0 -3.0 2.0)
        (-3.0 10.0 -5.0)
        ( 2.0 -5.0 6.0)
        B =( 27.0 -78.0 64.0)
    DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
    DATA B/27.0, -78.0, 64.0/
    CALL LSLDS (A, B, X)
    CALL WRRRN ('X', X, 1, N, 1)
    !
END

```

\section*{Output}
```

        X
    2 3
    1.000 -4.000 7.000

```

\section*{ScaLAPACK Example}

The same system of three linear equations is solved as a distributed computing example. The coefficient matrix has real positive definite form and the right-hand-side vector \(b\) has three elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSLDS INT
USE WRRRN_INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), B(:), X(:)
REAL, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=3)
Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP RANK .EQ. O) THEN
ALLOCATE (A (LDA,N), B(N), X(N))
Set values for A and B
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/ -3.0, 10.0, -5.0/)
A(3,:) = (/ 2.0, -5.0, 6.0/)
B = (/27.0, -78.0, 64.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
Se}t up t\overline{he array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP MB, 1, 0, 0, MP ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), XO(MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK MAP(B, DESCX, BO)
Solve the system of equations
CALL LSLDS (A0, B0, XO)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full

```
```

! array.
CALL SCALAPACK_UNMAP(X0, DESCX, X)
Print results.
Only Rank=O has the solution, X.
IF(MP RANK .EQ. O)CALL WRRRN ('X', X, 1, N, 1)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}
\begin{tabular}{ccc} 
& X \\
1 & 2 & 3 \\
1.000 & -4.000 & 7.000
\end{tabular}

\section*{LFCDS}

Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N symmetric positive definite matrix to be factored. (Input) Only the upper triangle of A is referenced.
\(\boldsymbol{F A C T}\) - N by N matrix containing the upper triangular matrix \(R\) of the factorization of \(A\) in the upper triangular part. (Output) Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCDS (A, FACT, RCOND [, ...])
Specific: The specific interface names are S_LFCDS and D_LFCDS.

\section*{FORTRAN 77 Interface}

Single: CALL LFCDS (N, A, LDA, FACT, LDFACT, RCOND)
Double: The double precision name is DLFCDS.

\section*{ScaLAPACK Interface}

Generic: CALL LFCDS (A0, FACTO, RCOND [,...])
Specific: The specific interface names are S_LFCDS and D_LFCDS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCDS computes an \(R^{T} R\) Cholesky factorization and estimates the condition number of a real symmetric positive definite coefficient matrix. The matrix \(R\) is upper triangular.
The \(L_{1}\) condition number of the matrix A is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCDS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

The \(R^{T} R\) factors are returned in a form that is compatible with routines LFIDS, LFSDS and LFDDS. To solve systems of equations with multiple right-hand-side vectors, use LFCDS followed by either LFIDS or LFSDS called once for each right-hand side. The routine LFDDS can be called to compute the determinant of the coefficient matrix after LFCDS has performed the factorization.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CDS/DL2CDS. The reference is:

CALL L2CDS (N, A, LDA, FACT, LDFACT, RCOND, WK)
The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
\(42 \quad\) The input matrix is not positive definite.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
\(\boldsymbol{A O}\) - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the symmetric positive definite matrix to be factored. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the upper triangular matrix \(R\) of the factorization of \(A\) in the upper triangular part. (Output)
Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFCDS is called to factor the matrix and to check for nonpositive definiteness or ill-conditioning. LFIDS is called to determine the columns of the inverse.
```

USE LFCDS_INT
USE UMACH_INT
USE WRRRN_INT
USE LFIDS_INT
Declare variables
INTEGER LDA, LDFACT, N, NOUT
PARAMETER (LDA=3, LDFACT=3, N=3)
REAL A(LDA,LDA), AINV (LDA,LDA), RCOND, FACT (LDFACT,LDFACT) , \&
RES(N), RJ (N)
Set values for A

```
```

! A = ( 1.0 1.0 -3.0
(-3.0 10.0 -5.0)
( 2.0 -5.0 6.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
Factor the matrix A
CALL LFCDS (A, FACT, RCOND)
Set up the columns of the identity
matrix one at a time in RJ
RJ = 0.0E0
DO 10 J=1, N
RJ (J) = 1.0EO
RJ is the J-th column of the identity
matrix so the following LFIDS
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFIDS (A, FACT, RJ, AINV(:,J), RES)
RJ(J) = 0.0EO
1 0 ~ C O N T I N U E ~
!
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
CALL WRRRN ('AINV', AINV)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F9.3)
END

```

\section*{Output}

RCOND < 0.005
L1 Condition number < 875.0

AINV
\begin{tabular}{rrr}
1 & 2 & 3 \\
35.00 & 8.00 & -5.00 \\
8.00 & 2.00 & -1.00 \\
-5.00 & -1.00 & 1.00
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. LFCDS is called to factor the matrix and to check for singularity or ill-conditioning. LFIDS is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFCDS_INT
USE UMACH_INT
USE LFIDS_INT
USE WRRRN_INT
USE SCALA\overline{PACK_SUPPORT}

```
```

IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER J, LDA, N, NOUT, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), AINV(:,:), X0(:), RJ(:)
REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), RESO(:), RJO(:)
REAL RCOND
PARAMETER (LDA=3,N=3)
MP NPROCS = MP SETUP()
IF(MP_RANK .EQ. O) THEN
ALLOCATE (A (LDA,N), AINV (LDA,N))
Set values for A
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/ -3.0, 10.0, -5.0/)
A(3,:) = (/ 2.0, -5.0, 6.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT(DESCA, N, N, MP MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL) , X0 (MXLDA) ,FACTO (MXLDA,MXCOL), RJ (N) , \&
RJ0 (MXLDA), RESO (MXLDA))
Map input array to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Call the factorization routine
CALL LFCDS (A0, FACTO, RCOND)
Print the reciprocal condition number
and the L1 condition number
IF(MP RANK .EQ. O) THEN
CALL UMACH (2, NOUT)
WRITE (NOUT,99998) RCOND, 1.0E0/RCOND
ENDIF
Set up the columns of the identity
matrix one at a time in RJ
RJ = 0.0E0
DO 10 J=1, N
RJ(J) = 1.0
Map input array to the processor grid
CALL SCALAPACK_MAP(RJ, DESCL, RJO)
RJ is the J-th column of the identity
matrix so the following LFIDS
reference computes the J-th column of
the inverse of A
CALL LFIDS (A0, FACTO, RJO, X0, RESO)
RJ(J) = 0.0
CALL SCALAPACK_UNMAP(X0, DESCL, AINV (:,J))

```
```

    1 0 \text { CONTINUE}
    !
! O
IF(MP_RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (MP RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLO\overline{CATE (A0, FACT0, RJ, RJ0, RESO, X0)}
Exit ScaLAPACK usage
CALL SCALAPACK EXIT(MP ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F9.3)
END

```

\section*{Output}

RCOND < 0.005
L1 Condition number \(<875.0\)
```

            AINV
    ```
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 35.00 & 8.00 & -5.00 \\
2 & 8.00 & 2.00 & -1.00 \\
3 & -5.00 & -1.00 & 1.00
\end{tabular}

\section*{LFTDS}

Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix.

\section*{Required Arguments}

A-n by n symmetric positive definite matrix to be factored. (Input)
Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing the upper triangular matrix \(R\) of the factorization of A in the upper triangle, and the lower triangular matrix \(R^{T}\) in the lower triangle. (Output) If A is not needed, A and FACT can share the same storage location.

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTDS (A, FACT [,...])
Specific: The specific interface names are S_LFTDS and D_LFTDS.

\section*{FORTRAN 77 Interface}

Single: CALL LFTDS (N, A, LDA, FACT, LDFACT)
Double: \(\quad\) The double precision name is DLFTDS.

\section*{ScaLAPACK Interface}

Generic: CALL LFTDS (A0, FACTO [,...])
Specific: The specific interface names are S_LFTDS and D_LFTDS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFTDS computes an \(R^{T} R\) Cholesky factorization of a real symmetric positive definite coefficient matrix. The matrix \(R\) is upper triangular.
LFTDS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

The \(R^{T} R\) factors are returned in a form that is compatible with routines LFIDS, LFSDS and LFDDS. To solve systems of equations with multiple right-hand-side vectors, use LFTDS followed by either LFIDS or LFSDS called once for each right-hand side. The routine LFDDS can be called to compute the determinant of the coefficient matrix after LFTDS has performed the factorization.
The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational error
\begin{tabular}{cl} 
Type & Code \\
4 & 2 The input matrix is not positive definite.
\end{tabular}

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the symmetric positive definite matrix to be factored. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the upper triangular matrix \(R\) of the factorization of \(A\) in the upper triangular part. (Output)
Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFTDS is called to factor the matrix and to check for nonpositive definiteness. LFSDS is called to determine the columns of the inverse.
```

USE LFTDS_INT
USE LFSDS_INT
USE WRRRN_INT
INTEGER LDA, LDFACT, N
PARAMETER (LDA=3, LDFACT=3, N=3)
REAL A(LDA,LDA), AINV (LDA,LDA), FACT (LDFACT,LDFACT), RJ (N)
Set values for A
A = ( 1.0 -3.0 2.0)
( -3.0 10.0 -5.0)
( 2.0 -5.0 6.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
CALL LFTDS (A, FACT)

```
\(R J=0.0 E 0\)
DO \(10 \mathrm{~J}=1\), N
    \(R J(J)=1.0 E 0\)
\(!\)
\(!\)
\(!\)
\(!\)
matrix so the following LFSDS
reference places the J-th column of
the inverse of \(A\) in the \(J\)-th column
of AINV
        CALL LFSDS (FACT, RJ, AINV (:, J))
        \(R J(J)=0.0 \mathrm{EO}\)
    10 CONTINUE
! Print the results

\section*{Output}
\begin{tabular}{rr} 
AINV & \\
2 & 3 \\
8.00 & -5.00 \\
2.00 & -1.00 \\
-1.00 & 1.00
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. LFTDS is called to factor the matrix and to check for nonpositive definiteness. LFSDS is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFTDS INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE LFSDS_INT
USE WRRRN INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), AINV(:,:), X0(:)
REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), RESO(:), RJO(:)
PARAMETER (LDA=3,N=3)
MP NPROCS = MP SETUP()
IF\overline{(MP_RANK .EQ. O) THEN}
ALLOCATE (A (LDA,N), AINV (LDA,N))
Set values for A
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/ -3.0, 10.0, -5.0/)
A (3,:) = (/ 2.0, -5.0, 6.0/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
Sèt up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCL, N, 1, MP_MB, 1, - 0, 0, MP_ICT\overline{XT, MXLDA, INFO)}

```
```

! Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), RJ (N), \&
RJ0 (MXLDA), RES0 (MXLDA), IPVT0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Call the factorization routine
CALL LFTDS (A0, FACTO)
RJ = 0.0EO
DO 10 J=1, N
RJ (J) = 1.0
CALL SCALAPACK_MAP(RJ, DESCL, RJO)
RJ is the J-th column of the identity
matrix so the following LFSDS
reference computes the J-th column of
the inverse of A
CALL LFSDS (FACTO, RJO, X0)
RJ(J) = 0.0
CALL SCALAPACK UNMAP(X0, DESCL, AINV (:,J))
1 0 ~ C O N T I N U E
Print results.
Only Rank=0 has the solution, AINV.
IF(MP RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (MP RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE (A0, FACT0, IPVT0, RJ, RJ0, RESO, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
END

```

\section*{Output}

RCOND < 0.005
L1 Condition number \(<875.0\)

AINV
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 35.00 & 8.00 & -5.00 \\
2 & 8.00 & 2.00 & -1.00 \\
3 & -5.00 & -1.00 & 1.00
\end{tabular}

\section*{LFSDS}

Solves a real symmetric positive definite system of linear equations given the \(R^{T} R\) Cholesky factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - N by N matrix containing the \(R^{T} R\) factorization of the coefficient matrix A as output from routine LFCDS/DLFCDS or LFTDS/DLFTDS. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSDS (FACT, B, X [,...])

Specific: The specific interface names are S_LFSDS and D_LFSDS.

\section*{FORTRAN 77 Interface}

Single: CALL LFSDS (N, FACT, LDFACT, B, X)

Double: \(\quad\) The double precision name is DLFSDS.

\section*{ScaLAPACK Interface}

Generic: CALL LFSDS (FACTO, B0, XO [,...])
Specific: The specific interface names are S_LFSDS and D_LFSDS.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

This routine computes the solution for a system of linear algebraic equations having a real symmetric positive definite coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either LFCDS or LFTDS. \(R\) is an upper triangular matrix.

The solution to \(A x=b\) is found by solving the triangular systems \(R^{T} y=b\) and \(R x=y\).

LFSDS and LFIDS both solve a linear system given its \(R^{T} R\) factorization. LFIDS generally takes more time and produces a more accurate answer than LFSDS. Each iteration of the iterative refinement algorithm used by LFIDS calls LFSDS.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational error
\begin{tabular}{cl} 
Type & Code \\
4 & 1
\end{tabular}

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(R^{T} R\) factorization of the coefficient matrix A as output from routine LFCDS/DLFCDS or LFTDS/DLFTDS. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. LFTDS is called to factor the coefficient matrix. LFSDS is called to compute the four solutions for the four right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCDS to perform the factorization, and LFIDS to compute the solutions.
```

USE LFSDS INT
USE LFTDS }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
INTEGER LDA, LDFACT, N
PARAMETER (LDA=3, LDFACT=3, N=3)
REAL A(LDA,LDA), B(N,4), FACT (LDFACT,LDFACT), X(N, 4)
Set values for A and B

```
```

!
!
!
!
!
!
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
DATA B/-1.0, -3.0, -3.0, 3.6, -4.2, -5.2, -8.0, 11.0, -6.0,\&
-9.4, 17.6, -23.4/
CALL LFTDS (A, FACT)
DO 10 I=1, 4
CALL LFSDS (FACT, B(:,I), X(:,I))
1 0
CONTINUE
CALL WRRRN ('The solution vectors are', X)
END

```
\(!\)

\section*{Output}
```

        The solution vectors are
            1 2 2 3 4
    -44.0 118.4 -162.0 -71.2
    -11.0 25.6 -36.0 -16.6
        5.0 -19.0 23.0 6.0
    ```

\section*{ScaLAPACK Example}

The same set of linear systems is solved successively as a distributed example. Routine LFTDS is called to factor the coefficient matrix. The routine LFSDS is called to compute the four solutions for the four right-hand sides. In this case, the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCDS to perform the factorization, and LFIDS to compute the solutions. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFSDS_INT
USE LFTDS_INT
USE WRRRN_INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER J, LDA, N, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), B(:,:), X(:,:), X0(:)
REAL, ALLOCATABLE :: AO(:,:), FACTO(:,:), BO(:)
PARAMETER (LDA=3,N=3)

```
```

! Set up for MPI
MP_NPROCS = MP_SETUP()
IF\overline{(MP RANK .EQ-. O) THEN}
ALLOCATE (A (LDA,N), B(N,4), X(N,4))
Set values for A and B
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/ -3.0, 10.0, -5.0/)
A(3,:) = (/ 2.0, -5.0, 6.0/)
B(1,:) = (/ -1.0, 3.6, -8.0, -9.4/)
B (2,:) = (/ -3.0, -4.2, 11.0, 17.6/)
B(3,:) = (/ -3.0, -5.2, -6.0, -23.4/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
S\overline{e}t up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT (DESCL, N, 1, MP MB, 1, 0, 0, MP ICTXTT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), B0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL LFTDS (A0, FACTO)
Call the factorization routine
Set up the columns of the B
matrix one at a time in XO
DO 10 J=1, 4
CALL SCALAPACK_MAP(B(:,j), DESCL, B0)
Solve for the J-th column of X
CALL LFSDS (FACTO, BO, X0)
CALL SCALAPACK_UNMAP (X0, DESCL, X(:,J))
1 0 CONTINUE
Print results.
Only Rank=0 has the solution, X.
IF(MP RANK.EQ.O) CALL WRRRN ('The solution vectors are', X)
IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)
DEALLOCATE (A0, FACT0, B0, X0)
Exit Scalapack usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}
```

    The solution vectors are
    1 rrrrer
    -11.0 25.6 -36.0 -16.6
5.0 -19.0 23.0 6.0

```

\section*{LFIDS}

\section*{勾MPI}

Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations.

\section*{Required Arguments}
\(A-\mathrm{N}\) by N matrix containing the symmetric positive definite coefficient matrix of the linear system. (Input) Only the upper triangle of A is referenced.

FACT - N by N matrix containing the \(R^{T} R\) factorization of the coefficient matrix A as output from routine LFCDS/DLFCDS or LFTDS/DLFTDS. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
\(\boldsymbol{R E S}\) - Vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimesion statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIDS (A, FACT, B, X, RES [,...])
Specific: The specific interface names are S_LFIDS and D_LFIDS.

\section*{FORTRAN 77 Interface}

Single: CALL LFIDS (N, A, LDA, FACT, LDFACT, B, X, RES)

Double: The double precision name is DLFIDS.

\section*{ScaLAPACK Interface}

Generic: CALL LFIDS (A0, FACT0, B0, X0, RESO [, ...])
Specific: The specific interface names are S_LFIDS and D_LFIDS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFIDS computes the solution of a system of linear algebraic equations having a real symmetric positive definite coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

To compute the solution, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either LFCDS or LFTDS.

Iterative refinement fails only if the matrix is very ill-conditioned.
LFIDS and LFSDS both solve a linear system given its \(R^{T} R\) factorization. LFIDS generally takes more time and produces a more accurate answer than LFSDS. Each iteration of the iterative refinement algorithm used by LFIDS calls LFSDS.

\section*{Comments}

Informational error
Type Code
32 The input matrix is too ill-conditioned for iterative refinement to be effective.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the symmetric positive definite coefficient matrix of the linear system. (Input)

FACTO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the \(R^{T} R\) factorization of the coefficient matrix A as output from routine LFCDS/DLFCDS or LFTDS/DLFTDS. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)
\(X 0\) - Local vector of length MXLDA containing the local portions of the distributed vector \(X\). X contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

RESO - Local vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the residual vector at the improved solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding 0.2 to the second element.
```

USE LFIDS_INT
USE LFCDS_INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
INTEGER LDA, LDFACT, N
PARAMETER (LDA=3, LDFACT=3, N=3)
REAL A(LDA,LDA), B(N), RCOND, FACT(LDFACT,LDFACT), RES (N,3) ,\&
X(N, 3)
Set values for A and B
A = ( 1.0 -3.0 2.0)
(-3.0 10.0 -5.0)
( 2.0 -5.0 6.0)
B = ( 1.0 -3.0 2.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
DATA B/1.0, -3.0, 2.0/
CALL LFCDS (A, FACT, RCOND)
Factor the matrix A
Print the estimated condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
! Compute the solutions
DO 10 I=1, 3
CALL LFIDS (A, FACT, B, X(:,I), RES(:,I))
B(2) = B(2) + . 2E0
1 0 CONTINUE
CALL WRRRN ('The solution vectors are', X)
CALL WRRRN ('The residual vectors are', RES)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F9.3)

```
!

\section*{Output}


\section*{ScaLAPACK Example}

The same set of linear systems is solved successively as a distributed example. The right-handside vector is perturbed after solving the system each of the first two times by adding 0.2 to the second element. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFIDS_INT
USE LFCDS - INT
USE UMACH_INT
USE WRRRN INT
USE SCALAP\overline{PACK SUPPORT}
IMPLICIT NONE
INCLUDE `mpif.h'
INTEGER J, LDA, N, NOUT, DESCA(9), DESCL(9)
INTEGER INFO, MXCOL, MXLDA
REAL RCOND
REAL, ALLOCATABLE :: A(:,:), B(:), X(:,:), RES(:,:), X0(:)
REAL, ALLOCATABLE :: A0(:,:), FACTO(:,:), B0(:), RESO(:)
PARAMETER (LDA=3, N=3)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A(LDA,N), B(N), X(N,3), RES(N,3))
Set values for A and B
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/-3.0, 10.0, -5.0/)
A(3,:) = (/ 2.0, -5.0, 6.0/)
B = (/ 1.0, -3.0, 2.0/)
ENDIF

```
```

                                    Set up a 1D processor grid and define
                                    its context ID, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
                            Get the array descriptor entities MXLDA,
                                    and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
                                    Set up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCL, N, 1, MP_MB, 1, 0, 0, MP_ICTXTT, MXLDA, INFO)
            Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA), FACTO (MXLDA,MXCOL), BO (MXLDA) , &
                RESO (MXLDA) )
                            Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, AO)
                            Call the factorization routine
    CALL LFCDS (A0, FACTO, RCOND)
                                    Print the estimated condition number
    CALL UMACH (2, NOUT)
    IF(MP RANK .EQ. O) WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
                                    Set up the columns of the B
                                    matrix one at a time in XO
    DO 10 J=1, 3
        CALL SCALAPACK_MAP(B, DESCL, B0)
                                    Solve for the J-th column of X
        CALL LFIDS (A0, FACTO, B0, X0, RESO)
        CALL SCALAPACK_UNMAP (X0, DESCL, X(:,J))
        CALL SCALAPACK UNMAP (RESO, DESCL, RES (:,J))
        IF(MP_RANK .EQ. 0) B (2) = B (2) + . 2EO
    1 0 ~ C O N T I N U E ~
                                    Print results.
                                    Only Rank=0 has the full arrays
    IF(MP_RANK.EQ.O) CALL WRRRN ('The solution vectors are', X)
    IF(MP RANK.EQ.O) CALL WRRRN ('The residual vectors are', RES)
    IF (MP_RANK .EQ. 0) DEALLOCATE (A, B, X, RES)
    DEALLOCATE (A0, B0, FACT0, RES0, X0)
                                    Exit ScaLAPACK usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
        Shut down MPI
    MP_NPROCS = MP_SETUP('FINAL')
    99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F9.3)
END

```

\section*{Output}

\begin{tabular}{llll}
1 & 0.0000 & 0.0000 & 0.0000 \\
2 & 0.0000 & 0.0000 & 0.0000 \\
3 & 0.0000 & 0.0000 & 0.0000
\end{tabular}

\section*{LFDDS}

Computes the determinant of a real symmetric positive definite matrix given the \(R^{T} R\) Cholesky factorization of the matrix .

\section*{Required Arguments}

FACT - N by N matrix containing the \(R^{T} R\) factorization of the coefficient matrix A as output from routine LFCDS/DLFCDS or LFTDS/DLFTDS. (Input)

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that, \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form, \(\operatorname{det}(\mathrm{A})=\operatorname{DET} 1 * 10^{\mathrm{DET} 2}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDDS (FACT, DET1, DET2 [, ...])
Specific: The specific interface names are S_LFDDS and D_LFDDS.

\section*{FORTRAN 77 Interface}

Single: CALL LFDDS (N, FACT, LDFACT, DET1, DET2)
Double: The double precision name is DLFDDS.

\section*{Description}

Routine LFDDS computes the determinant of a real symmetric positive definite coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either LFCDS or LFTDS. The formula \(\operatorname{det} A=\operatorname{det} R^{T} \operatorname{det} R=(\operatorname{det} R)^{2}\) is
used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} R=\prod_{i=1}^{N} R_{i i}
\]
(The matrix \(R\) is stored in the upper triangle of FACT.)
LFDDS is based on the LINPACK routine SPODI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a real positive definite \(3 \times 3\) matrix.
```

    USE LFDDS_INT
    USE LFTDS_INT
    USE UMACH INT
    ! - Declare variables
INTEGER LDA, LDFACT, NOUT
PARAMETER (LDA=3, LDFACT=3)
REAL A(LDA,LDA), DET1, DET2, FACT (LDFACT,LDFACT)
Set values for A
A = ( 1.0 -3.0 2.0)
(-3.0 20.0 -5.0)
( 2.0 -5.0 6.0)
DATA A/1.0, -3.0, 2.0, -3.0, 20.0, -5.0, 2.0, -5.0, 6.0/
Factor the matrix
CALL LFTDS (A, FACT)
CALL LFDDS (FACT, DET1, DET2)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is ',F6.3,' * 10**',F2.0)
END

```

\section*{Output}
```

The determinant of A is 2.100 * 10**1.

```

\section*{LINDS}

Computes the inverse of a real symmetric positive definite matrix.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the symmetric positive definite matrix to be inverted. (Input) Only the upper triangle of \(A\) is referenced.

AINV — N by n matrix containing the inverse of A. (Output)
If A is not needed, A and AINV can share the same storage locations.

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix A. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDAINV — Leading dimension of AINV exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAINV = size (AINV,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LINDS (A, AINV [,...])
Specific: The specific interface names are S_LINDS and D_LINDS.

\section*{FORTRAN 77 Interface}

Single: CALL LINDS (N, A, LDA, AINV, LDAINV)
Double: The double precision name is DLINDS.

\section*{ScaLAPACK Interface}

Generic: CALL LINDS (A0, AINV0 [,...])

Specific: The specific interface names are S_LINDS and D_LINDS.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LINDS computes the inverse of a real symmetric positive definite matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see
"Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LINDS first uses the routine LFCDS to compute an \(R^{T} R\) factorization of the coefficient matrix and to estimate the condition number of the matrix. LINRT is then used to compute \(R^{-1}\). Finally \(A^{-1}\) is computed using \(R^{-1}=R^{-1} R^{-T}\).

LINDS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.
If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in \(A\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2NDS/DL2NDS. The reference is:

CALL L2NDS (N, A, LDA, AINV, LDAINV, WK)
The additional argument is:
\(W K\) - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is not positive definite.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the symmetric positive definite matrix to be inverted. (Input)

AINVO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix AINV. AINV contains the inverse of the matrix A. (Output)
If \(A\) is not needed, A and AINV can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse is computed for a real positive definite \(3 \times 3\) matrix.
```

        USE LINDS_INT
        USE WRRRN_INT
    !
INTEGER LDA, LDAINV
PARAMETER (LDA=3, LDAINV=3)
REAL A (LDA,LDA), AINV (LDAINV,LDAINV)
Set values for A
A =( llll
( -3.0 10.0 -5.0)
( 2.0 -5.0 6.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
CALL LINDS (A, AINV)
CALL WRRRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{rrr} 
& AINV & \\
1 & 2 & 3 \\
35.00 & 8.00 & -5.00 \\
8.00 & 2.00 & -1.00 \\
-5.00 & -1.00 & 1.00
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(3 \times 3\) matrix is computed as a distributed example. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LINDS INT
USE WRRRN INT
USE SCALAP\overline{P}ACK SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER J, LDA, LDFACT, N, DESCA(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), AINV(:,:)
REAL, ALLOCATABLE :: AO(:,:), AINVO(:,:)
PARAMETER (LDA=3, N=3)
MP NPROCS = MP SETUP()
IF(MP RANK .EQ_ 0) THEN
ALLOCATE (A (LDA,N), AINV(LDA,N))
Set values for A
A(1,:) = (/ 1.0, -3.0, 2.0/)
A(2,:) = (/ -3.0, 10.0, -5.0/)

```
```

A(3,:) = (/ 2.0, -5.0, 6.0/)

```

ENDIF
```

                                    Set up a 1D processor grid and define
                                    its context ID, MP ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), AINVO (MXLDA,MXCOL))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Call the routine to get the inverse
CALL LINDS (A0, AINVO)
Unmap the results from the distributed
arrays back to a nondistributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(AINVO, DESCA, AINV)
Print results.
Only Rank=0 has the solution, AINV.
IF(MP_RANK.EQ.O) CALL WRRRN ('AINV', AINV)
IF (MP RANK .EQ. O) DEALLOCATE(A, AINV)
DEALLOC
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

AINV
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 35.00 & 8.00 & -5.00 \\
2 & 8.00 & 2.00 & -1.00 \\
3 & -5.00 & -1.00 & 1.00
\end{tabular}

\section*{LSASF}

Solves a real symmetric system of linear equations with iterative refinement.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the coefficient matrix of the symmetric linear system. (Input) Only the upper triangle of A is referenced.
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSASF (A, B, X [, ...])
Specific: The specific interface names are S_LSASF and D_LSASE.

\section*{FORTRAN 77 Interface}

Single: CALL LSASF (N, A, LDA, B, X)
Double: The double precision name is DLSASF.

\section*{Description}

Routine LSASF solves systems of linear algebraic equations having a real symmetric indefinite coefficient matrix. It first uses the routine LFCSF to compute a \(U D U^{T}\) factorization of the coefficient matrix and to estimate the condition number of the matrix. \(D\) is a block diagonal matrix with blocks of order 1 or 2 , and \(U\) is a matrix composed of the product of a permutation matrix and a unit upper triangular matrix. The solution of the linear system is then found using the iterative refinement routine LFISF.

LSASF fails if a block in \(D\) is singular or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSASF solves the problem that is represented in the computer, however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ASF/DL2ASF. The reference is

CALL L2ASF (N, A, LDA, B, X, FACT, IPVT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-\mathrm{N} \times \mathrm{N}\) work array containing information about the \(U D U^{T}\) factorization of A on output. If A is not needed, A and FACT can share the same storage location.

IPVT - Integer work vector of length N containing the pivoting information for the factorization of \(A\) on output.

WK - Work vector of length n.
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ASF the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and \(\operatorname{IVAL}(4)\) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSASF. Additional memory allocation for FACT and option value restoration are done automatically in LSASF. Users directly calling L2ASF can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSASF or L2ASF. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSASF temporarily replaces IVAL(2) by IVAL(1). The routine L2CSF computes the condition number if IVAL \((2)=2\). Otherwise L2CSF skips this computation. LSASF restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has real symmetric form and the right-hand-side vector \(b\) has three elements.
```

    USE LSASF_INT
    USE WRRRN_INT
    ! Declare variables
PARAMETER (LDA=3, N=3)
REAL A(LDA,LDA), B (N), X(N)
Set values for A and B
A =( 1.0 -2.0 1.0)
( -2.0 3.0 -2.0)
( 1.0 -2.0 3.0)

```
```

! B = ( }$$
\begin{array}{llll}{4.1}&{-4.7}&{6.5}\end{array}
$$
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
DATA B/4.1, -4.7, 6.5/
!
CALL LSASF (A, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{1}{c}{\(x\)} & 3 \\
-4.100 & -3.500 & 1.200
\end{tabular}

\section*{LSLSF}

Solves a real symmetric system of linear equations without iterative refinement .

\section*{Required Arguments}
\(A-\mathrm{N}\) by N matrix containing the coefficient matrix of the symmetric linear system. (Input) Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}-\) Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLSF (A, B, X [, ...])
Specific: The specific interface names are S_LSLSF and D_LSLSF.

\section*{FORTRAN 77 Interface}

Single: CALL LSLSF (N, A, LDA, B, X)
Double: The double precision name is DLSLSF.

\section*{Description}

Routine LSLSF solves systems of linear algebraic equations having a real symmetric indefinite coefficient matrix. It first uses the routine LFCSF to compute a \(U D U^{T}\) factorization of the coefficient matrix. \(D\) is a block diagonal matrix with blocks of order 1 or 2 , and \(U\) is a matrix composed of the product of a permutation matrix and a unit upper triangular matrix.

The solution of the linear system is then found using the routine LFSSF.
LSLSF fails if a block in \(D\) is singular. This occurs only if \(A\) either is singular or is very close to a singular matrix.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LSF/DL2LSF. The reference is:

CALL L2LSF (N, A, LDA, B, X, FACT, IPVT, WK)
The additional arguments are as follows:
FACT- \(\mathrm{N} \times \mathrm{N}\) work array containing information about the \(U D U^{T}\) factorization of A on output. If A is not needed, A and FACT can share the same storage locations.

IPVT - Integer work vector of length N containing the pivoting information for the factorization of A on output.

WK - Work vector of length \(N\).
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine LSLSF the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and IVAL(4) are temporarily replaced by \(\operatorname{IVAL}(1)\) and IVAL(2), respectively, in LSLSF. Additional memory allocation for FACT and option value restoration are done automatically in LSLSF. Users directly calling LSLSF can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLSF or LSLSF. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLSF temporarily replaces IVAL(2) by IVAL(1). The
routine L2CSF computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CSF skips this computation. LSLSF restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has real symmetric form and the right-hand-side vector \(b\) has three elements.
```

USE LSLSF_INT
USE WRRRN_INT
PARAMETER (LDA=3, N=3)
REAL A(LDA,LDA), B(N), X(N)
Set values for A and B
A = ( 1.0 -2.0 1.0)
( -2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
B = ( 4.1 -4.7 6.5)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
DATA B/4.1, -4.7, 6.5/
CALL LSLSF (A, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{3}{c}{X} \\
& 1 & 3 \\
-4.100 & \(-3.500^{3}\) & 1.200
\end{tabular}

\section*{LFCSF}

Computes the \(U D U^{T}\) factorization of a real symmetric matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N symmetric matrix to be factored. (Input)
Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing information about the factorization of the symmetric matrix A. (Output)
Only the upper triangle of FACT is used. If A is not needed, A and FACT can share the same storage locations.
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the factorization. (Output)
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCSF (A, FACT, IPVT, RCOND [,...])
Specific: The specific interface names are \(S_{\_}\)LFCSF and D_LFCSF.

\section*{FORTRAN 77 Interface}

Single: CALL LFCSF (N, A, LDA, FACT, LDFACT, IPVT, RCOND)
Double: \(\quad\) The double precision name is DLFCSF.

\section*{Description}

Routine LFCSF performs a \(U D U^{T}\) factorization of a real symmetric indefinite coefficient matrix. It also estimates the condition number of the matrix. The \(U D U^{T}\) factorization is called the diagonal pivoting factorization.

The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCSF fails if A is singular or very close to a singular matrix.
The \(U D U^{T}\) factors are returned in a form that is compatible with routines LFISF, LFSSF and LFDSF. To solve systems of equations with multiple right-hand-side vectors, use LFCSF followed
by either LFISF or LFSSF called once for each right-hand side. The routine LFDSF can be called to compute the determinant of the coefficient matrix after LFCSF has performed the factorization.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CSF/DL2CSF. The reference is:

CALL L2CSF (N, A, LDA, FACT, LDFACT, IPVT, RCOND, WK)
The additional argument is:
\(\boldsymbol{W K}\) - Work vector of length n.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
42 The input matrix is singular.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFCSF is called to factor the matrix and to check for singularity or ill-conditioning. LFISF is called to determine the columns of the inverse.
```

USE LFCSF_INT
USE UMACH INT
USE LFISF}\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
PARAMETER (LDA=3, N=3)
INTEGER IPVT (N), NOUT
REAL A(LDA,LDA), AINV (N,N), FACT (LDA,LDA), RJ (N), RES (N),\&
RCOND
Set values for A
A = ($$
\begin{array}{lll}{1.0}&{-2.0}&{1.0}\end{array}
$$)
( -2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
Factor A and return the reciprocal
condition number estimate
CALL LFCSF (A, FACT, IPVT, RCOND)
Print the estimate of the condition
number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND

```
!
```

! matrix one at a time in RJ
RJ = 0.EO
DO 10 J=1, N
RJ(J) = 1.0E0
RJ is the J-th column of the identity
matrix so the following LFISF
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFISF (A, FACT, IPVT, RJ, AINV(:,J), RES)
RJ(J) = 0.0E0
1 0 ~ C O N T I N U E ~
! Print the inverse
CALL WRRRN ('AINV', AINV)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND < 0.05
L1 Condition number < 40.0
AINV
1 rrrr
2 -2.000 -1.000 0.000
3-0.500 0.000 0.500

```

\section*{LFTSF}

Computes the \(U D U^{T}\) factorization of a real symmetric matrix.

\section*{Required Arguments}

A—N by N symmetric matrix to be factored. (Input) Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing information about the factorization of the symmetric matrix A. (Output)
Only the upper triangle of FACT is used. If A is not needed, A and FACT can share the same storage locations.

IPVT - Vector of length N containing the pivoting information for the factorization. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTSF (A, FACT, IPVT [,...])
Specific: The specific interface names are S_LFTSF and D_LFTSF.

\section*{FORTRAN 77 Interface}

Single: CALL LFTSF (N, A, LDA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTSF.

\section*{Description}

Routine LFTSF performs a \(U D U^{T}\) factorization of a real symmetric indefinite coefficient matrix. The \(U D U^{T}\) factorization is called the diagonal pivoting factorization.

LFTSF fails if \(A\) is singular or very close to a singular matrix.
The \(U D U^{T}\) factors are returned in a form that is compatible with routines LFISF, LFSSF and LFDSF. To solve systems of equations with multiple right-hand-side vectors, use LFTSF followed by either LFISF or LFSSF called once for each right-hand side. The routine LFDSF can be called to compute the determinant of the coefficient matrix after LFTSF has performed the factorization.
The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational error
Type Code
42 The input matrix is singular.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFTSF is called to factor the matrix and to check for singularity. LFSSF is called to determine the columns of the inverse.
```

USE LFTSF INT
USE LFSSF_INT

```
```

    USE WRRRN_INT
    ! Declare variables
PARAMETER (LDA=3, N=3)
INTEGER IPVT (N)
REAL A(LDA,LDA), AINV (N,N), FACT(LDA,LDA), RJ(N)
Set values for A
A = ( 1.0 -2.0 1.0)
(-2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
Factor A
CALL LFTSF (A, FACT, IPVT)
Set up the columns of the identity
matrix one at a time in RJ
RJ = 0.0EO
DO 10 J=1, N
RJ(J) = 1.0E0
RJ is the J-th column of the identity
matrix so the following LFSSF
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFSSF (FACT, IPVT, RJ, AINV(:,J))
RJ (J) = 0.0EO
1 0 ~ C O N T I N U E ~
CALL WRRRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{3}{c}{ AINV } \\
& 1 & 2 & 3 \\
1 & -2.500 & -2.000 & -0.500 \\
2 & -2.000 & -1.000 & 0.000 \\
3 & -0.500 & 0.000 & 0.500
\end{tabular}

\section*{LFSSF}

Solves a real symmetric system of linear equations given the \(U D U^{T}\) factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - N by N matrix containing the factorization of the coefficient matrix A as output from routine LFCSF/DLFCSF or LFTSF/DLFTSF. (Input) Only the upper triangle of \(F A C T\) is used.

IPVT - Vector of length \(N\) containing the pivoting information for the factorization of A as output from routine LFCSF/DLFCSF or LFTSF/DLFTSF. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(\mathbf{N}\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size \((\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSSF (FACT, IPVT, B, X [, ...])
Specific: The specific interface names are S_LFSSF and D_LFSSF.

\section*{FORTRAN 77 Interface}

Single: CALL LFSSF (N, FACT, LDFACT, IPVT, B, X)
Double: The double precision name is DLFSSF.

\section*{Description}

Routine LFSSF computes the solution of a system of linear algebraic equations having a real symmetric indefinite coefficient matrix.

To compute the solution, the coefficient matrix must first undergo a \(U D U^{T}\) factorization. This may be done by calling either LFCSF or LFTSF.

LFSSF and LFISF both solve a linear system given its \(U D U^{T}\) factorization. LFISF generally takes more time and produces a more accurate answer than LFSSF. Each iteration of the iterative refinement algorithm used by LFISF calls LFSSF.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Example}

A set of linear systems is solved successively. LFTSF is called to factor the coefficient matrix. LFSSF is called to compute the four solutions for the four right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCSF to perform the factorization, and LFISF to compute the solutions.
```

        USE LFSSF_INT
        USE LFTSF INT
        USE WRRRN_INT
    !
PARAMETER (LDA=3, N=3)
INTEGER IPVT(N)
REAL A(LDA,LDA), B(N,4), X(N,4), FACT (LDA,LDA)
Set values for A and B
A =( 1.0 -2.0 1.0)
(-2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
B =( (-1.0 3.6 -8.0 -9.4)
( -3.0 -4.2 11.0 17.6)
(-3.0 -5.2 -6.0 -23.4)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
DATA B/-1.0, -3.0, -3.0, 3.6, -4.2, -5.2, -8.0, 11.0, -6.0,\&
-9.4, 17.6, -23.4/
Factor A
CALL LFTSF (A, FACT, IPVT)
DO 10 I=1, 4
CALL LFSSF (FACT, IPVT, B(:,I), X(:,I))
1 0 ~ C O N T I N U E ~
!
CALL WRRRN ('X', X)
END

```

\section*{Output}
\begin{tabular}{rrrrr} 
& & & \(X\) & \\
& 1 & 2 & 3 & 4 \\
1 & 10.00 & 2.00 & 1.00 & 0.00 \\
2 & 5.00 & -3.00 & 5.00 & 1.20 \\
3 & -1.00 & -4.40 & 1.00 & -7.00
\end{tabular}

\section*{LFISF}

Uses iterative refinement to improve the solution of a real symmetric system of linear equations.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{N}\) by N matrix containing the coefficient matrix of the symmetric linear system. (Input) Only the upper triangle of \(A\) is referenced

FACT - N by N matrix containing the factorization of the coefficient matrix A as output from routine LFCSF/DLFCSF or LFTSF/DLFTSF. (Input) Only the upper triangle of \(F A C T\) is used.

IPVT - Vector of length N containing the pivoting information for the factorization of A as output from routine LFCSF/DLFCSF or LFTSF/DLFTSF. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
\(\boldsymbol{R E S}\) - Vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFISF (A, FACT, IPVT, B, X, RES [,...])

Specific: The specific interface names are S_LFISF and D_LFISF.

\section*{FORTRAN 77 Interface}

Single: CALL LFISF (N, A, LDA, FACT, LDFACT, IPVT, B, X, RES)
Double: The double precision name is DLFISF.

\section*{Description}

LFISF computes the solution of a system of linear algebraic equations having a real symmetric indefinite coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo a \(U D U^{T}\) factorization. This may be done by calling either LFCSF or LFTSF.

Iterative refinement fails only if the matrix is very ill-conditioned.

LFISF and LFSSF both solve a linear system given its \(U D U^{T}\) factorization. LFISF generally takes more time and produces a more accurate answer than LFSSF. Each iteration of the iterative refinement algorithm used by LFISF calls LFSSF .

\section*{Comments}

Informational error
Type Code
32 The input matrix is too ill-conditioned for iterative refinement to be effective.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding 0.2 to the second element.
```

USE LFISF INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE LFCSF_INT
USE WRRRN INT
PARAMETER (LDA=3,N=3)
INTEGER IPVT(N), NOUT
REAL A(LDA,LDA), B (N), X(N), FACT (LDA,LDA), RES (N), RCOND
Set values for A and B
A = ( 1.0 -2.0 1.0)
( -2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
B = ( 4.1 -4.7 6.5)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
DATA B/4.1, -4.7, 6.5/
Factor A and compute the estimate
of the reciprocal condition number
CALL LFCSF (A, FACT, IPVT, RCOND)
Print condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
! Solve, then perturb right-hand side
DO 10 I=1, 3
CALL LFISF (A, FACT, IPVT, B, X, RES)
! Print results
CALL WRRRN ('X', X, 1, N, 1)
CALL WRRRN ('RES', RES, 1, N, 1)
B(2) = B(2) + .20E0
1 0 CONTINUE
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```
!
!

\section*{Output}
```

RCOND < 0.035
L1 Condition number < 40.0
X

|  | X |  |  |  |
| ---: | ---: | ---: | :---: | :---: |
| 1 | 2 | 3 |  |  |
| -4.100 | -3.500 | 1.200 |  |  |

            RES
    ```

```

            X
    | X |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -4.500 | -3.700 | 1.200 |

            RES
    ```

```

            X
    1
l

```

\section*{LFDSF}

Computes the determinant of a real symmetric matrix given the \(U D U^{T}\) factorization of the matrix.

\section*{Required Arguments}
\(\boldsymbol{F A C T}-\mathrm{N}\) by N matrix containing the factored matrix \(A\) as output from subroutine LFTSF/DLFTSF or LFCSF/DLFCSF. (Input)

IPVT — Vector of length n containing the pivoting information for the \(U D U^{T}\) factorization as output from routine LFTSF/DLFTSF or LFCSF/DLFCSF. (Input)

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that, \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form, \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\text {DET2 }}\).

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDSF (FACT, IPVT, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDSF and D_LFDSF.

\section*{FORTRAN 77 Interface}

Single: CALL LFDSF (N, FACT, LDFACT, IPVT, DET1, DET2)
Double: \(\quad\) The double precision name is DLFDSF.

\section*{Description}

Routine LFDSF computes the determinant of a real symmetric indefinite coefficient matrix. To compute the determinant, the coefficient matrix must first undergo a \(U D U^{T}\) factorization. This may be done by calling either LFCSF or LFTSF. Since det \(U= \pm 1\), the formula \(\operatorname{det} A=\operatorname{det} U \operatorname{det} D \operatorname{det} U^{T}=\operatorname{det} D\) is used to compute the determinant. Next \(\operatorname{det} D\) is computed as the product of the determinants of its blocks.

LFDSF is based on the LINPACK routine SSIDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a real symmetric \(3 \times 3\) matrix.
```

USE LFDSF INT
USE LFTSE_INT
USE UMACH_INT
PARAMETER (LDA=3, N=3)
INTEGER IPVT(N), NOUT
REAL A(LDA,LDA), FACT (LDA,LDA), DET1, DET2
Set values for A
A =( 1.0 -2.0 1.0)
( -2.0 3.0 -2.0)
( 1.0 -2.0 3.0)
DATA A/1.0, -2.0, 1.0, -2.0, 3.0, -2.0, 1.0, -2.0, 3.0/
Factor A
CALL LFTSF (A, FACT, IPVT)
CALL LFDSF (FACT, IPVT, DET1, DET2)
Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2

```
```

99999 FORMAT (' The determinant of A is ', F6.3, ' * 10**', F2.0)

```
    END

\section*{Output}

The determinant of \(A\) is \(-2.000 * 10 * * 0\).

\section*{LSADH}

Solves a Hermitian positive definite system of linear equations with iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian positive definite linear system. (Input) Only the upper triangle of \(A\) is referenced.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution of the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSADH (A, B, X [, ...])
Specific: The specific interface names are S_LSADH and D_LSADH.

\section*{FORTRAN 77 Interface}

Single: CALL LSADH (N, A, LDA, B, X)
Double: The double precision name is DLSADH.

\section*{ScaLAPACK Interface}

Generic: CALL LSADH (A0, B0, X0 [,...])

Specific: The specific interface names are S_LSADH and D_LSADH.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSADH solves a system of linear algebraic equations having a complex Hermitian positive definite coefficient matrix. It first uses the routine LFCDH to compute an \(R^{H} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. The matrix \(R\) is upper triangular. The solution of the linear system is then found using the iterative refinement routine LFIDH.

LSADH fails if any submatrix of \(R\) is not positive definite, if \(R\) has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) either is very close to a singular matrix or is a matrix that is not positive definite.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSADH solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ADH/DL2ADH. The reference is:

CALL L2ADH (N, A, LDA, B, X, FACT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-\mathrm{N} \times \mathrm{N}\) work array containing the \(R^{H} R\) factorization of A on output.
WK - Complex work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.

44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ADH the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and IVAL(4) are temporarily replaced by \(\operatorname{IVAL}(1)\) and IVAL(2), respectively, in LSADH. Additional memory allocation for FACT and option value restoration are done automatically in LSADH. Users directly calling L2ADH can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSADH or L2ADH. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSADH temporarily replaces IVAL(2) by IVAL(1). The routine \(\operatorname{L2CDH}\) computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CDH skips this computation. LSADH restores the option. Default values for the option \(\operatorname{are} \operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
\(\boldsymbol{A O}\) - Complex MXLDA by mXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the Hermitian positive definite linear system. (Input)
Only the upper triangle of \(A\) is referenced.
B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector \(\mathrm{X} . \mathrm{X}\) contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of five linear equations is solved. The coefficient matrix has complex positive definite form and the right-hand-side vector \(b\) has five elements.
```

USE LSADH_INT
USE WRCRN_INT

```
```

! Declare variables
INTEGER LDA, N
PARAMETER (LDA=5, N=5)
COMPLEX A(LDA,LDA), B(N), X(N)
Set values for A and B
A = ( 2.0+0.0i -1.0+1.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
( 4.0+0.0i 1.0+2.0i 0.0+0.0i 0.0+0.0i )
( 10.0+0.0i 0.0+4.0i 0.0+0.0i )
(
B = ( 1.0+5.0i 12.0-6.0i 1.0-16.0i -3.0-3.0i 25.0+16.0i )
DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),\&
4* (0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0),\&
(0.0,4.0), (6.0,0.0), 4* (0.0,0.0), (1.0,1.0), (9.0,0.0)/
DATA B / (1.0,5.0), (12.0,-6.0), (1.0,-16.0), (-3.0,-3.0),\&
(25.0,16.0)/
CALL LSADH (A, B, X)
Print results
CALL WRCRN ('X', X, 1, N, 1)
!
END

```

\section*{Output}
\((2.000,1.000)^{1} \quad(3.000,0.000)^{2} \quad(-1.000,-1.000)^{3} \quad(0.000,-2.000)^{4}\)

\section*{ScaLAPACK Example}

The same system of five linear equations is solved as a distributed computing example. The coefficient matrix has complex positive definite form and the right-hand-side vector \(b\) has five elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSADH INT
USE WRCRN - INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA

```
```

    COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:)
    COMPLEX, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
    PARAMETER (LDA=5, N=5)
    MP NPROCS = MP SETUP()
    IF(MP RANK .EQ. 0) THEN
        ALLOCATE (A (LDA,N), B(N), X(N))
                Set values for A and B
    A(1,:) = (/ (2.0, 0.0),(-1.0, 1.0),( 0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
    A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0),(0.0, 0.0)/)
    A(3,:) = (/ (0.0, 0.0),( 0.0, 0.0),(10.0, 0.0),(0.0, 4.0),(0.0, 0.0)/)
    A(4,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
    A(5,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
    B = (/ (1.0, 5.0),(12.0, -6.0),(1.0, -16.0),(-3.0, -3.0),(25.0, 16.0)/)
    ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
    CALL SCALAPACK SETUP(N, N, .TRUE., .TRUE.)
    Get the array descriptor entities MXLDA,
    and MXCOL
    CALL SCALAPACK GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
    Sēt up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXXT, MXLDA, INFO)
    Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA), X0(MXLDA))
    Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, AO)
    CALL SCALAPACK MAP(B, DESCX, B0)
    Solve the system of equations
    CALL LSADH (A0, B0, X0)
        Unmap the results from the distributed
    arrays back to a non-distributed array.
    After the unmap, only Rank=0 has the full
    array.
    CALL SCALAPACK_UNMAP(X0, DESCX, X)
        Print results.
        Only Rank=0 has the solution, X.
    IF(MP_RANK .EQ. 0) CALL WRCRN ('X', X, 1, N, 1)
    IF (MP RANK .EQ. O) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, X0)
    CALL SCALAPACK_EXIT(MP_ICTXT)
        Shut down MPI
    MP_NPROCS = MP_SETUP(`FINAL')
    END
    ```

\section*{Output}
x
\((2.000,1.000)^{1} \quad(3.000,0.000)^{2} \quad(-1.000,-1.000)^{3} \quad(0.000,-2.000)^{4}\)
( \(3.000,2.000\) )

\section*{LSLDH}

\section*{勻MPI}

Solves a complex Hermitian positive definite system of linear equations without iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian positive definite linear system. (Input) Only the upper triangle of A is referenced.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLDH (A, B, X [, ...])
Specific: The specific interface names are S_LSLDH and D_LSLDH.

\section*{FORTRAN 77 Interface}

Single: CALL LSLDH (N, A, LDA, B, X)
Double: The double precision name is DLSLDH.

\section*{ScaLAPACK Interface}

Generic: CALL LSLDH (A0, B0, X0 [, ...])
Specific: The specific interface names are S_LSLDH and D_LSLDH.

See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLDH solves a system of linear algebraic equations having a complex Hermitian positive definite coefficient matrix. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. LSLDH first uses the routine LFCDH to compute an \(R^{H} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. The matrix \(R\) is upper triangular. The solution of the linear system is then found using the routine LFSDH.

LSLDH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSADH be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LDH/ DL2LDH. The reference is:
```

CALL L2LDH (N, A, LDA, B, X, FACT, WK)

```

The additional arguments are as follows:
FACT \(-\mathrm{N} \times \mathrm{N}\) work array containing the \(R^{H} R\) factorization of A on output. If A is not needed, A can share the same storage locations as FACT.

WK - Complex work vector of length n.
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix is not positive definite.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LDH the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by \(\operatorname{IVAL}(1)\) and IVAL(2), respectively, in LSLDH. Additional memory allocation for FACT and option value restoration are done automatically in LSLDH. Users directly calling L2LDH can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLDH or L2LDH. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).
17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLDH temporarily replaces IVAL(2) by IVAL(1). The routine L2CDH computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise \(\operatorname{L2CDH}\) skips this computation. LSLDH restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - Complex mXLDA by mXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the Hermitian positive definite linear system. (Input)
Only the upper triangle of A is referenced.
B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of five linear equations is solved. The coefficient matrix has complex Hermitian positive definite form and the right-hand-side vector \(b\) has five elements.
```

USE LSLDH INT
USE WRCRN_INT
! - Declare variables
INTEGER LDA, N
PARAMETER (LDA=5, N=5)
COMPLEX A(LDA,LDA), B(N), X(N)
!

```
```

! Set values for A and B
A = ( 2.0+0.0i -1.0+1.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
( 4.0+0.0i 1.0+2.0i 0.0+0.0i 0.0+0.0i )
( 10.0+0.0i 0.0+4.0i 0.0+0.0i )
( 6.0+0.0i 1.0+1.0i)
B = ( 1.0+5.0i 12.0-6.0i 1.0-16.0i -3.0-3.0i 25.0+16.0i )
DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),\&
4*(0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0),\&
(0.0,4.0), (6.0,0.0), 4*(0.0,0.0), (1.0,1.0), (9.0,0.0)/
DATA B / (1.0,5.0), (12.0,-6.0), (1.0,-16.0), (-3.0,-3.0),\&
(25.0,16.0)/
!
CALL LSLDH (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
!
END

```

\section*{Output}

\section*{X}
\((2.000,1.000)^{1}\)
\((3.000,2.000)^{2}\)\(\quad(3.000,0.000)^{2} \quad(-1.000,-1.000)^{3} \quad(0.000,-2.000)^{4}\)

\section*{ScaLAPACK Example}

The same system of five linear equations is solved as a distributed computing example. The coefficient matrix has complex positive definite form and the right-hand-side vector \(b\) has five elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LSLDH_INT
USE WRCRN-INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
INTEGER LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
COMPLEX, ALLOCATABLE :: A(:,:), B(:), X(:)
COMPLEX, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=5, N=5)
Set up for MPI

```
```

MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. O) THEN
ALLOCATE (A (LDA,N), B(N), X(N))
Set values for A and B
A(1,:) = (/ (2.0, 0.0),(-1.0, 1.0),( 0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0), (0.0, 0.0)/)
A(3,:)}=(/(0.0,0.0),(0.0, 0.0),(10.0, 0.0),(0.0, 4.0),(0.0, 0.0)/
A(4,:) = (/ (0.0, 0.0),( 0.0, 0.0), ( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
A(5,:)=(/(0.0, 0.0),(0.0, 0.0),(0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
B = (/ (1.0, 5.0),(12.0, -6.0),(1.0, -16.0), (-3.0, -3.0),(25.0, 16.0)/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Se}t up th\overline{e} array descriptors
CALL DESCINIT (DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
\overline{Allocate space for the local arrays}
ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA), X0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, AO)
CALL SCALAPACK_MAP(B, DESCX, BO)
Solve the system of equations
CALL LSLDH (A0, BO, X0)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP(X0, DESCX, X)
Print results.
Only Rank=0 has the solution, X.
IF(MP RANK .EQ. 0) CALL WRCRN ('X', X, 1, N, 1)
IF (MP_RANK .EQ. O) DEALLOCATE (A, B, X)
DEALLOCATE (A0, B0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

\section*{X}
\((2.000,1.000)^{1} \quad(3.000,0.000)^{2} \quad(-1.000,-1.000)^{3} \quad(0.000,-2.000)^{4}\)
( \(3.000,2.000\) )

\section*{LFCDH}

\section*{包MPI}

Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}

A-Complex n by n Hermitian positive definite matrix to be factored. (Input) Only the upper triangle of A is referenced.

FACT - Complex n by n matrix containing the upper triangular matrix \(R\) of the factorization of A in the upper triangle. (Output)
Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

RCOND - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(\mathbf{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT --- Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCDH (A, FACT, RCOND [, ...])

Specific: The specific interface names are S_LFCDH and D_LFCDH.

\section*{FORTRAN 77 Interface}

Single: CALL LFCDH (N, A, LDA, FACT, LDFACT, RCOND)
Double: The double precision name is DLFCDH.

\section*{ScaLAPACK Interface}

Generic: CALL LFCDH (A0, FACTO, RCOND [,...])
Specific: The specific interface names are S_LFCDH and D_LFCDH.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFCDH computes an \(R^{H} R\) Cholesky factorization and estimates the condition number of a complex Hermitian positive definite coefficient matrix. The matrix \(R\) is upper triangular.
The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCDH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

The \(R^{H} R\) factors are returned in a form that is compatible with routines LFIDH, LFSDH and LFDDH. To solve systems of equations with multiple right-hand-side vectors, use LFCDH followed by either LFIDH or LFSDH called once for each right-hand side. The routine LFDDH can be called to compute the determinant of the coefficient matrix after LFCDH has performed the factorization.
The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 C D H / D L 2 C D H\). The reference is:

CALL L2CDH (N, A, LDA, FACT, LDFACT, RCOND, WK)
The additional argument is
WK - Complex work vector of length N .
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.

4 The input matrix is not Hermitian.
42 The input matrix is not positive definite. It has a diagonal entry with an imaginary part

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - Complex mXLDA by mXCOL local matrix containing the local portions of the distributed matrix A. A contains the Hermitian positive definite matrix to be factored. (Input)
Only the upper triangle of \(A\) is referenced.
FACTO - Complex MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the upper triangular matrix \(R\) of the factorization of \(A\) in the upper triangle. (Output)
Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(5 \times 5\) Hermitian positive definite matrix is computed. LFCDH is called to factor the matrix and to check for nonpositive definiteness or ill-conditioning. LFIDH is called to determine the columns of the inverse.
```

USE LFCDH INT
USE LFIDH_INT
USE UMACH INT
USE WRCRN_INT
INTEGER LDA, LDFACT, N, NOUT
PARAMETER (LDA=5, LDFACT=5, N=5)
REAL RCOND
COMPLEX A(LDA,LDA), AINV(LDA,LDA), FACT(LDFACT,LDFACT),\&
RES(N), RJ(N)
A = ( 2.0+0.0i -1.0+1.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
( 4.0+0.0i 1.0+2.0i 0.0+0.0i 0.0+0.0i )
( 10.0+0.0i 0.0+4.0i 0.0+0.0i )
( 6.0+0.0i 1.0+1.0i )
DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),\&
4*(0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0),\&
(0.0,4.0), (6.0,0.0), 4*(0.0,0.0), (1.0,1.0), (9.0,0.0)/

```
```

! Factor the matrix A
CALL LFCDH (A, FACT, RCOND)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0EO, 0.0EO)
DO 10 J=1, N
RJ(J) = (1.0E0,0.0EO)
RJ is the J-th column of the identity
matrix so the following LFIDH
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFIDH (A, FACT, RJ, AINV(:,J), RES)
RJ(J) = (0.0E0,0.0E0)
1 0 ~ C O N T I N U E
! Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.OEO/RCOND
CALL WRCRN ('AINV', AINV)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.075
L1 Condition number < 25.0
\begin{tabular}{|c|c|c|c|c|}
\hline & \multicolumn{4}{|c|}{AINV} \\
\hline 1 & ( 0.7166, 0.0000) & ( 0.2166,-0.2166) & (-0.0899,-0.0300) & (-0.0207, 0.0622) \\
\hline 2 & ( 0.2166, 0.2166) & ( 0.4332, 0.0000) & (-0.0599,-0.1198) & (-0.0829, 0.0415) \\
\hline 3 & (-0.0899, 0.0300) & (-0.0599, 0.1198) & ( 0.1797, 0.0000) & ( 0.0000,-0.1244) \\
\hline 4 & (-0.0207, -0.0622) & (-0.0829,-0.0415) & ( 0.0000, 0.1244) & ( 0.2592, 0.0000) \\
\hline 5 & \[
(0.0092,0.0046)_{5}
\] & \[
(0.0138,-0.0046)
\] & (-0.0138,-0.0138) & (-0.0288, 0.0288) \\
\hline \multicolumn{5}{|l|}{\(1(0.0092,-0.0046)\)} \\
\hline \multicolumn{5}{|l|}{\(2(0.0138,0.0046)\)} \\
\hline \multicolumn{5}{|l|}{3 (-0.0138, 0.0138)} \\
\hline \multicolumn{5}{|l|}{\(4(-0.0288,-0.0288)\)} \\
\hline & ( 0.1175, 0.0000) & & & \\
\hline
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(5 \times 5\) Hermitian positive definite matrix in the preceding example is computed as a distributed computing example. LFCDH is called to factor the matrix and to check for nonpositive definiteness or ill-conditioning. LFIDH (page 187) is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFCDH INT
USE LFIDH_INT
USE WRCRN_INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
Declare variables
INTEGER J, LDA, N, NOUT, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
REAL RCOND
COMPLEX, ALLOCATABLE :: A(:,:), AINV(:,:), RJ(:), RJO(:)
COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), RESO(:), XO(:)
PARAMETER (LDA=5, N=5)
MP NPROCS = MP SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N), AINV(LDA,N))
Set values for A and B
A(1,:) = (/ (2.0, 0.0),(-1.0, 1.0),( 0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0),(0.0, 0.0)/)
A(3,:) = (/ (0.0, 0.0),( 0.0, 0.0),(10.0, 0.0),(0.0, 4.0),(0.0, 0.0)/)
A(4,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
A(5,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
Sèt up the array descriptors
CALL DESCINIT(DESCA, N, N, MP MB, MP_NB, 0, 0, MP ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE(A0 (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), RJ(N), \&
RJO(MXLDA), RESO (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, AO)
Factor the matrix A
CALL LFCDH (AO, FACTO, RCOND)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0EO, 0.0E0)
DO 10 J=1, N
RJ(J) = (1.0E0,0.0E0)
CALL SCALAPACK_MAP(RJ, DESCX, RJO)
RJ is the J-th column of the identity
matrix so the following LFIDH
reference solves for the J-th column of
the inverse of A
CALL LFIDH (A0, FACT0, RJ0, X0, RES0)
Unmap the results from the distributed
array back to a non-distributed array
CALL SCALAPACK_UNMAP(X0, DESCX, AINV(:,J))

```
```

                RJ(J) = (0.0E0,0.0E0)
    1 0 ~ C O N T I N U E
    !
!
IF(MP RANK .EQ. O) THEN
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0EO/RCOND
CALL WRCRN ('AINV', AINV)
ENDIF
IF (MP RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLO\overline{CATE (AO, FACTO, RJ, RJO, RESO, X0)}
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
99999 FORMAT (' RCON}D=',F5.3,/,' L1 Condition number = ',F6.3
END

```

\section*{Output}

RCOND < 0.075
L1 Condition number < 25.0
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{AINV} \\
\hline & 1 & 2 & & 4 \\
\hline 1 & ( 0.7166, 0.0000) & ( 0.2166,-0.2166) & (-0.0899,-0.0300) & \((-0.0207,0.0622)\) \\
\hline 2 & ( 0.2166, 0.2166) & ( 0.4332, 0.0000) & (-0.0599,-0.1198) & (-0.0829, 0.0415) \\
\hline 3 & (-0.0899, 0.0300) & (-0.0599, 0.1198) & ( 0.1797, 0.0000) & ( 0.0000, -0.1244) \\
\hline 4 & (-0.0207, -0.0622) & (-0.0829, -0.0415) & ( 0.0000, 0.1244) & ( 0.2592, 0.0000) \\
\hline 5 & \[
(0.0092,0.0046)_{5}
\] & \[
(0.0138,-0.0046)
\] & (-0.0138,-0.0138) & (-0.0288, 0.0288) \\
\hline 1 & ( 0.0092,-0.0046) & & & \\
\hline 2 & ( 0.0138, 0.0046) & & & \\
\hline 3 & (-0.0138, 0.0138) & & & \\
\hline 4 & (-0.0288, -0.0288) & & & \\
\hline 5 & ( 0.1175, 0.0000) & & & \\
\hline
\end{tabular}

\section*{LFTDH}

Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix.

\section*{Required Arguments}

A-Complex n by n Hermitian positive definite matrix to be factored. (Input) Only the upper triangle of A is referenced.

FACT - Complex N by N matrix containing the upper triangular matrix \(R\) of the factorization of A in the upper triangle. (Output)

Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTDH (A, FACT [,...])

Specific: The specific interface names are S_LFTDH and D_LFTDH.

\section*{FORTRAN 77 Interface}

Single: CALL LFTDH (N, A, LDA, FACT, LDFACT)
Double: The double precision name is DLFTDH.

\section*{ScaLAPACK Interface}

Generic: CALL LFTDH (A0, FACTO [, ...])
Specific: The specific interface names are S_LFTDH and D_LFTDH.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFTDH computes an \(R^{H} R\) Cholesky factorization of a complex Hermitian positive definite coefficient matrix. The matrix \(R\) is upper triangular.

LFTDH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element.
These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.
The \(R^{H} R\) factors are returned in a form that is compatible with routines LFIDH, LFSDH and LFDDH. To solve systems of equations with multiple right-hand-side vectors, use LFCDH followed
by either LFIDH or LFSDH called once for each right-hand side. The IMSL routine LFDDH can be called to compute the determinant of the coefficient matrix after LFCDH has performed the factorization.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational errors
\begin{tabular}{cll} 
Type & Code \\
3 & 4 & \begin{tabular}{l} 
The input matrix is not Hermitian. It has a diagonal entry with a small \\
imaginary part.
\end{tabular} \\
4 & 2 & The input matrix is not positive definite. \\
4 & 4 & \begin{tabular}{l} 
The input matrix is not Hermitian. It has a diagonal entry with an imaginary \\
part.
\end{tabular}
\end{tabular}

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - Complex mXLDA by mXcol local matrix containing the local portions of the distributed matrix A. A contains the Hermitian positive definite matrix to be factored. (Input)
Only the upper triangle of A is referenced.
FACTO - Complex MXLDA by MXCOL local matrix containing the local portions of the distributed matrix FACT. FACT contains the upper triangular matrix \(R\) of the factorization of A in the upper triangle. (Output)
Only the upper triangle of FACT will be used. If A is not needed, A and FACT can share the same storage locations.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

The inverse of a \(5 \times 5\) matrix is computed. LFTDH is called to factor the matrix and to check for nonpositive definiteness. LFSDH is called to determine the columns of the inverse.
```

USE LFTDH_INT
USE LFSDH INT
USE WRCRN_INT
! Declare variables
INTEGER LDA, LDFACT, N
PARAMETER (LDA=5, LDFACT=5, N=5)
COMPLEX A(LDA,LDA), AINV (LDA,LDA), FACT(LDFACT,LDFACT), RJ (N)

```
\(!\)
```

!
!
!
i
!
!
!
!
DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),\&
4* (0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0), \&
(0.0.4.0),(6.0,0.0), 4*(0.0,0.0), (1.0,1.0), (9.0,0.0)/
Factor the matrix A
CALL LFTDH (A, FACT)
RJ = (0.0E0,0.0EO)
DO 10 J=1, N
RJ (J) = (1.0E0,0.0E0)
RJ is the J-th column of the identity
matrix so the following LFSDH
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFSDH (FACT, RJ, AINV (:,J))
RJ (J) = (0.0E0,0.0EO)
10 CONTINUE
CALL WRCRN ('AINV', AINV, ITRING=1)
!
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline & \multicolumn{4}{|c|}{AINV} \\
\hline & 1 & & 3 & 4 \\
\hline 1 & \((0.7166,0.0000)\) & 0.2166,-0.2166) & (-0.0899, -0.0300) & (-0.0207, 0.0622) \\
\hline 2 & & 0.4332, 0.0000) & (-0.0599,-0.1198) & (-0.0829, 0.0415) \\
\hline 3 & & & ( 0.1797, 0.0000) & ( 0.0000, -0.1244) \\
\hline 4 & & & & ( 0.2592, 0.0000) \\
\hline & 5 & & & \\
\hline 1 & ( 0.0092,-0.0046) & & & \\
\hline 2 & ( 0.0138, 0.0046) & & & \\
\hline 3 & (-0.0138, 0.0138) & & & \\
\hline 4 & (-0.0288, -0.0288) & & & \\
\hline 5 & ( 0.1175, 0.0000) & & & \\
\hline
\end{tabular}

\section*{ScaLAPACK Example}

The inverse of the same \(5 \times 5\) Hermitian positive definite matrix in the preceding example is computed as a distributed computing example. LFTDH is called to factor the matrix and to check for nonpositive definiteness. LFSDH (page 192) is called to determine the columns of the inverse. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities")
used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFTDDH INT
USE LFSDH_INT
USE WRCRN }\mp@subsup{}{}{-}\mathrm{ INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER J, LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
COMPLEX, ALLOCATABLE :: A(:,:), AINV(:,:), RJ(:), RJO(:)
COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), XO(:)
PARAMETER (LDA=5, N=5)
Set up for MPI
MP NPROCS = MP SETUP()
IF(MP_RANK .EQ. O) THEN
A\overline{LLOCATE (A (LDA,N), AINV (LDA,N))}
Set values for A and B
A(1,:) = (/ (2.0, 0.0),(-1.0, 1.0),(0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0), (0.0, 0.0)/)
A(3,:) = (/ (0.0, 0.0),( 0.0, 0.0),(10.0, 0.0),(0.0, 4.0), (0.0, 0.0)/)
A(4,:)=(/(0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
A(5,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
ENDIF
Set up a 1D processor grid and define
its context ID, MP ICTXT
CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
S\overline{e}t up the array descriptors
CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
CALL DESCINIT(DESCX, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA), FACTO (MXLDA,MXCOL), RJ(N), \&
RJ0 (MXLDA))
Map input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
Factor the matrix A
CALL LFTDH (A0, FACTO)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0E0, 0.0EO)
DO 10 J=1, N
RJ (J) = (1.0E0,0.0E0)
CALL SCALAPACK_MAP(RJ, DESCX, RJO)
RJ is the J-th column of the identity
matrix so the following LFIDH
reference solves for the J-th column of
the inverse of A
CALL LFSDH (FACTO, RJO, X0)
Unmap the results from the distributed

```
```

! array back to a non-distributed array
CALL SCALAPACK_UNMAP (X0, DESCX, AINV (:, J))
RJ(J) = (0.0EO-0.0E0)
1 0 ~ C O N T I N U E
Print the results.
After the unmap, only Rank=0 has the full
array.
IF(MP RANK .EQ. O) CALL WRCRN ('AINV', AINV)
IF (MP RANK .EQ. O) DEALLOCATE (A, AINV)
DEALLOCATE (A0, FACT0, RJ, RJ0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{AINV} \\
\hline & 1 & & & 4 \\
\hline 1 & ( 0.7166, 0.0000) & ( 0.2166,-0.2166) & (-0.0899,-0.0300) & (-0.0207, 0.0622) \\
\hline 2 & ( 0.2166, 0.2166) & ( 0.4332, 0.0000) & (-0.0599,-0.1198) & (-0.0829, 0.0415) \\
\hline 3 & (-0.0899, 0.0300) & (-0.0599, 0.1198) & ( 0.1797, 0.0000) & ( 0.0000, -0.1244) \\
\hline 4 & (-0.0207, -0.0622) & (-0.0829, -0.0415) & ( 0.0000, 0.1244) & ( 0.2592, 0.0000) \\
\hline 5 & \[
(0.0092,0.0046)_{5}
\] & \[
(0.0138,-0.0046)
\] & (-0.0138,-0.0138) & (-0.0288, 0.0288) \\
\hline 1 & ( 0.0092,-0.0046) & & & \\
\hline 2 & ( 0.0138, 0.0046) & & & \\
\hline 3 & (-0.0138, 0.0138) & & & \\
\hline 6 & (-0.0288, -0.0288) & & & \\
\hline 7 & ( 0.1175, 0.0000) & & & \\
\hline
\end{tabular}

\section*{LFSDH}

Solves a complex Hermitian positive definite system of linear equations given the \(R^{H} R\) factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - Complex n by n matrix containing the factorization of the coefficient matrix A as output from routine LFCDH/DLFCDH or LFTDH/DLFTDH. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSDH (FACT, B, X [,...])
Specific: The specific interface names are S_LFSDH and D_LFSDH.

\section*{FORTRAN 77 Interface}

Single: CALL LFSDH (N, FACT, LDFACT, B, X)

Double: The double precision name is DLFSDH.

\section*{ScaLAPACK Interface}

Generic: CALL LFSDH (FACT0, B0, X0 [, ...])
Specific: The specific interface names are S_LFSDH and D_LFSDH.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

This routine computes the solution for a system of linear algebraic equations having a complex Hermitian positive definite coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either LFCDH or LFTDH. \(R\) is an upper triangular matrix.
The solution to \(A x=b\) is found by solving the triangular systems \(R^{H} y=b\) and \(R x=y\).
LFSDH and LFIDH both solve a linear system given its \(R^{H} \mathrm{R}\) factorization. LFIDH generally takes more time and produces a more accurate answer than LFSDH. Each iteration of the iterative refinement algorithm used by LFIDH calls LFSDH.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational error
\(\begin{array}{cc}\text { Type } & \text { Code } \\ 4 & 1\end{array}\)

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT as output from routine LFCDH / DLFCDH or LFTDH/DLFTDH. FACT contains the factorization of the matrix A. (Input)

BO - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector \(\mathrm{X} . \mathrm{X}\) contains the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. LFTDH is called to factor the coefficient matrix. LFSDH is called to compute the four solutions for the four right-hand sides. In this case, the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCDH to perform the factorization, and LFIDH to compute the solutions.
```

USE LFSDH_INT
USE LFTDH_INT
USE WRCRN_INT

```
INTEGER LDA, LDFACT, N
PARAMETER (LDA=5, LDFACT=5, N=5)
COMPLEX \(A(L D A, L D A), B(N, 3), ~ F A C T(L D F A C T, L D F A C T), X(N, 3)\)
                                    Set values for A and B
\(\left.A=\begin{array}{rrrrr}(2.0+0.0 i & -1.0+1.0 i & 0.0+0.0 i & 0.0+0.0 i & 0.0+0.0 i\end{array}\right)\)
    \(B=\left(\begin{array}{ll} & 3.0+3.0 i\end{array} 4.0+0.0 i \quad 29.0-9.0 i\right)\)
    ( 5.0-5.0i 15.0-10.0i -36.0-17.0i )
    ( \(5.0+4.0 i-12.0-56.0 i \quad-15.0-24.0 i)\)
```

        ( 9.0+7.0i -12.0+10.0i -23.0-15.0i )
        (-22.0+1.0i 3.0-1.0i -23.0-28.0i )
    DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),&
        4* (0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0),&
        (0.0,4.0), (6.0,0.0), 4* (0.0,0.0), (1.0,1.0), (9.0,0.0)/
        DATA B / (3.0,3.0), (5.0,-5.0), (5.0,4.0), (9.0,7.0), (-22.0,1.0),&
        (4.0,0.0), (15.0,-10.0), (-12.0,-56.0), (-12.0,10.0), &
        (3.0,-1.0), (29.0,-9.0), (-36.0,-17.0), (-15.0,-24.0),&
                (-23.0,-15.0), (-23.0,-28.0)/
    CALL LFTDH (A, FACT)
        DO 10 I=1, 3
        CALL LFSDH (FACT, B(:,I), X(:,I))
    1 0 ~ C O N T I N U E
    CALL WRCRN ('X', X)
    END
    ```

\section*{Output}

X


\section*{ScaLAPACK Example}

The same set of linear systems as in in the preceding example is solved successively as a distributed computing example. LFTDH is called to factor the matrix. LFSDH is called to compute the four solutions for the four right-hand sides. In this case, the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCDH to perform the factorization, and LFIDH to compute the solutions.

SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LFTDH_INT
USE LFSDH INT
USE WRCRN_INT
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
INHEGER J, LDA, N, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA

```
```

    COMPLEX, ALLOCATABLE :: A(:,:), B(:,:), BO(:), X(:,:)
    COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), XO(:)
    PARAMETER (LDA=5, N=5)
    MP_NPROCS = MP_SETUP()
    IF(MP_RANK .EQ- 0) THEN
    ALLO\overline{CATE (A (LDA,N), B(LDA,3), X(LDA,3))}
                                    Set values for A and B
    A(1,:) = (/(2.0, 0.0),(-1.0, 1.0),( 0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
    A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0),(0.0, 0.0)/)
    A(3,:) = (/ (0.0, 0.0),( 0.0, 0.0),(10.0, 0.0),(0.0, 4.0),(0.0, 0.0)/)
    A(4,:) = (/(0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
    A(5,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
    B(1,:) = (/(3.0, 3.0), ( 4.0, 0.0), ( 29.0, -9.0)/)
    B(2,:) = (/ (5.0, -5.0), ( 15.0,-10.0), (-36.0,-17.0)/)
    B(3,:) = (/ (5.0, 4.0), (-12.0,-56.0), (-15.0,-24.0)/)
    B(4,:) = (/ (9.0, 7.0), (-12.0, 10.0), (-23.0,-15.0)/)
    B(5,:) = (/(-22.0.1.0), ( 3.0, -1.0), (-23.0,-28.0)/)
    ENDIF
                                    Set up a 1D processor grid and define
                                    its context ID, MP_ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
        Get the array descriptor entities MXLDA,
        and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
                    Sèt up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCX, N, 1, MP_MB, 1, -0, 0, MP_ICTXXT, MXLDA, INFO)
                A
    ALLOCATE (AO (MXLDA,MXCOL), XO (MXLDA),FACTO (MXLDA,MXCOL), &
            B0 (MXLDA))
    CALL SCALAPACK_MAP(A, DESCA, A0)
            Factor the matrix A
    CALL LFTDH (A0, FACTO)
    DO 10 J=1, 3
        CALL SCALAPACK_MAP(B(:,J), DESCX, BO)
        CALL LFSDH (FACTO, BO, XO)
                                    Unmap the results from the distributed
                    array back to a non-distributed array
        CALL SCALAPACK_UNMAP(X0, DESCX, X(:,J))
    10 CONTINUE
                            Print the results.
                    After the unmap, only Rank=0 has the full
                    array.
    IF(MP_RANK .EQ. O) CALL WRCRN ('X', X)
    IF (M\overline{P}}\mathrm{ RANK .EQ. O) DEALLOCATE (A, B, X)
    DEALLOCATE (A0, B0, FACTO, X0)
            Exit ScaLAPACK usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
        Shut down MPI
    MP_NPROCS = MP_SETUP(`FINAL')
    END
    ```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|c|}{x} \\
\hline & 1 & & & 2 & & 3 \\
\hline 1.00, & \(0.00)\) & ( & 3.00 , & -1.00) & ( 11.00, & -1.00) \\
\hline 1.00, & -2.00) & ( & 2.00 , & \(0.00)\) & ( -7.00, & \(0.00)\) \\
\hline 2.00, & \(0.00)\) & ( & -1.00, & -6.00) & ( -2.00, & -3.00) \\
\hline 2.00, & \(3.00)\) & & 2.00, & 1.00) & ( -2.00, & -3.00) \\
\hline ( -3.00, & \(0.00)\) & ( & 0.00, & \(0.00)\) & ( -2.00, & -3.00) \\
\hline
\end{tabular}

\section*{LFIDH}

Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex n by n matrix containing the coefficient matrix of the linear system. (Input) Only the upper triangle of A is referenced.

FACT - Complex N by N matrix containing the factorization of the coefficient matrix A as output from routine LFCDH/DLFCDH or LFTDH/DLFTDH. (Input) Only the upper triangle of FACT is used.

B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution. (Output)
RES - Complex vector of length \(N\) containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIDH (A, FACT, B, X, RES [,...])

Specific: The specific interface names are S_LFIDH and D_LFIDH.

\section*{FORTRAN 77 Interface}

Single: CALL LFIDH (N, A, LDA, FACT, LDFACT, B, X, RES)
Double: The double precision name is DLFIDH.

\section*{ScaLAPACK Interface}

Generic: CALL LFIDH (A0, FACTO, B0, X0, RESO [, ...])
Specific: The specific interface names are S_LFIDH and D_LFIDH.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LFIDH computes the solution of a system of linear algebraic equations having a complex Hermitian positive definite coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either LFCDH or LFTDH.

Iterative refinement fails only if the matrix is very ill-conditioned.
LFIDH and LFSDH both solve a linear system given its \(R^{H} R\) factorization. LFIDH generally takes more time and produces a more accurate answer than LFSDH. Each iteration of the iterative refinement algorithm used by LFIDH calls LFSDH.

\section*{Comments}

Informational error
Type Code
33 The input matrix is too ill-conditioned for iterative refinement to be effective.

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:

AO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the linear system. (Input) Only the upper triangle of \(A\) is referenced.

FACTO - MXLDA by MXCOL complex local matrix containing the local portions of the distributed matrix FACT as output from routine LFCDH or LFTDH. FACT contains the factorization of the matrix \(A\). (Input)
Only the upper triangle of FACT is referenced.
B0 - Complex local vector of length MXLDA containing the local portions of the distributed vector B. B contains the right-hand side of the linear system. (Input)
\(\boldsymbol{X 0}\) - Complex local vector of length MXLDA containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)

RESO - Complex local vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the residual vector at the improved solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (Chapter 11, "Utilities") after a call to SCALAPACK_SETUP (Chapter 11,"Utilities") has been made. See the ScaLAPACK Example below.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed by adding \((1+i) / 2\) to the second element after each call to LFIDH.
```

USE LFIDH_INT
USE LFCDH_INT
USE UMACH_INT
USE WRCRN_INT
INTEGER LDA, LDFACT, N
PARAMETER (LDA=5, LDFACT=5, N=5)
REAL RCOND
COMPLEX A(LDA,LDA), B(N), FACT(LDFACT,LDFACT), RES (N,3), X(N,3)
Set values for A and B
A = ( 2.0+0.0i -1.0+1.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
( 4.0+0.0i 1.0+2.0i 0.0+0.0i 0.0+0.0i )
( 10.0+0.0i 0.0+4.0i 0.0+0.0i )
( 6.0+0.0i 1.0+1.0i )
( 9.0+0.0i )
B = ( 3.0+3.0i 5.0-5.0i 5.0+4.0i 9.0+7.0i -22.0+1.0i )
DATA A / (2.0,0.0), 4* (0.0,0.0), (-1.0,1.0), (4.0,0.0),\&
4* (0.0,0.0), (1.0,2.0), (10.0,0.0), 4* (0.0,0.0),\&

```
```

                (0.0,4.0), (6.0,0.0), 4* (0.0,0.0), (1.0,1.0), (9.0,0.0)/
    DATA B / (3.0,3.0), (5.0,-5.0), (5.0,4.0), (9.0,7.0), (-22.0,1.0)/
    !
CALL LFCDH (A, FACT, RCOND)
! Print the estimated condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
Compute the solutions, then perturb B
DO 10 I=1, 3
CALL LFIDH (A, FACT, B, X(:,I), RES(:,I))
B(2) = B (2) + (0.5E0,0.5E0)
10 CONTINUE
! Print solutions and residuals
CALL WRCRN ('X', X)
CALL WRCRN ('RES', RES)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < 0.07
L1 Condition number < 25.0

X
\begin{tabular}{rrrrr}
1 \\
1 & \((1.000,0.000)\) & \((1.217,0.000)\) & \((1.433,0.000)\) \\
2 & \((1.000,-2.000)\) & \((1.217,-1.783)\) & \((1.433,-1.567)\) \\
3 & \((2.000,0.000)\) & \((1.910,0.030)\) & \((1.820,0.060)\) \\
4 & \((2.000,3.000)\) & \((1.979,2.938)\) & \((1.959,2.876)\) \\
5 & \((-3.000,0.000)\) & \((-2.991,0.005)\) & \((-2.982,0.009)\)
\end{tabular}

RES
\begin{tabular}{|c|c|c|}
\hline 1 & 2 & 3 \\
\hline \(1(1.192 \mathrm{E}-07,0.000 \mathrm{E}+00)\) & ( 6.592E-08, 1.686E-07) & ( 1.318E-07, 2.010E-14) \\
\hline 2 ( 1.192E-07,-2.384E-07) & (-5.329E-08, -5.329E-08) & ( 1.318E-07,-2.258E-07) \\
\hline 3 ( 2.384E-07, 8.259E-08) & ( 2.390E-07, -3.309E-08) & ( 2.395E-07, 1.015E-07) \\
\hline \(4(-2.384 \mathrm{E}-07,2.814 \mathrm{E}-14)\) & (-8.240E-08, -8.790E-09) & (-1.648E-07, -1.758E-08) \\
\hline 5 (-2.384E-07,-1.401E-08) & (-2.813E-07, 6.981E-09) & (-3.241E-07,-2.795E-08) \\
\hline
\end{tabular}

\section*{ScaLAPACK Example}

As in the preceding example, a set of linear systems is solved successively as a distributed computing example. The right-hand-side vector is perturbed by adding \((1+i) / 2\) to the second element after each call to LFIDH. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI_SETUP_INT
USE LFCDH_INT
USE LFIDH INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRCRN_INT

```
```

    USE SCALAPACK_SUPPORT
    IMPLICIT NONE
    INCLUDE `mpif.h'
    | Declare variables |  |
| :--- | :--- |
| INTEGER | J, LDA, N, NOUT, DESCA(9), DESCX(9) |
| INTEGER | INFO, MXCOL, MXLDA |

    IN INEO, MXCOL, MXLDA
    REAL RCOND
    COMPLEX, ALLOCATABLE :: A(:,:), B(:), BO(:), RES(:,:), X(:,:)
    COMPLEX, ALLOCATABLE :: AO(:,:), FACTO(:,:), XO(:), RESO(:)
    PARAMETER (LDA=5, N=5)
    MP NPROCS = MP SETUP()
    IF(MP RANK .EQ. O) THEN
        ALLO\overline{CATE (A(LDA,N), B(N), RES(N,3), X(N,3))}
                            Set values for A and B
        A(1,:) = (/ (2.0, 0.0),(-1.0, 1.0),( 0.0, 0.0),(0.0, 0.0),(0.0, 0.0)/)
        A(2,:) = (/ (0.0, 0.0),( 4.0, 0.0),( 1.0, 2.0),(0.0, 0.0),(0.0, 0.0)/)
        A(3,:) = (/ (0.0, 0.0),( 0.0, 0.0),(10.0, 0.0),(0.0, 4.0),(0.0, 0.0)/)
        A(4,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(6.0, 0.0),(1.0, 1.0)/)
        A(5,:) = (/ (0.0, 0.0),( 0.0, 0.0),( 0.0, 0.0),(0.0, 0.0),(9.0, 0.0)/)
        B = (/ (3.0, 3.0),( 5.0,-5.0),( 5.0, 4.0),(9.0, 7.0),(-22.0,1.0)/)
    ENDIF
                            Set up a 1D processor grid and define
                            its context ID, MP ICTXT
    CALL SCALAPACK_SETUP(N, N, .TRUE., .TRUE.)
    Get the array descriptor entities MXLDA,
    and MXCOL
    CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
    Set up the array descriptors
    CALL DESCINIT(DESCA, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO)
    CALL DESCINIT(DESCX, N, 1, MP MB, 1, 0, 0, MP ICTXT, MXLDA, INFO)
    A
    ALLOCATE (A0 (MXLDA,MXCOL), X0 (MXLDA),FACTO (MXLDA,MXCOL), &
                B0 (MXLDA), RESO(MXLDA))
                            Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, AO)
                            Factor the matrix A
    CALL LFCDH (A0, FACTO, RCOND)
                                    Print the estimated condition number
    IF(MP RANK .EQ. 0) THEN
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) RCOND, 1.OEO/RCOND
    ENDIF
    DO 10 J=1, 3
    CALL SCALAPACK MAP(B, DESCX, B0)
    CALL LFIDH (AO, FACTO, BO, XO, RESO)
                                    Unmap the results from the distributed
                    array back to a non-distributed array
    CALL SCALAPACK UNMAP(XO, DESCX, X(:,J))
    CALL SCALAPACK UNMAP(RESO, DESCX, RES(:,J))
    IF(MP_RANK .EQ. 0) B(2) = B(2) + (0.5E0, 0.5E0)
    10 CONTINUE
    Print the results.

```
```

! After the unmap, only Rank=0 has the full
IF(MP RANK .EQ. O) THEN
CALL WRCRN ('X', X)
CALL WRCRN ('RES', RES)
ENDIF
IF (MP_RANK .EQ. O) DEALLOCATE (A, B, RES, X)
DEALLOCATE (A0, B0, FACT0, RES0, X0)
Exit ScaLAPACK usage
CALL SCALAPACK EXIT(MP ICTXT)
Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
99999 FO\overline{RMAT (' RCON}D=',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND < 0.07
L1 Condition number < 25.0

```

X
\begin{tabular}{rrrrrr}
1 \\
1 & \((1.000,0.000)\) & \((1.217,0.000)\) & \((1.433,0.000)\) \\
2 & \((1.000,-2.000)\) & \((1.217,-1.783)\) & \((1.433,-1.567)\) \\
3 & \((2.000,0.000)\) & \((1.910,0.030)\) & \((1.820,0.060)\) \\
4 & \((2.000,3.000)\) & \((1.979,2.938)\) & \((1.959,2.876)\) \\
5 & \((-3.000,0.000)\) & \((-2.991,0.005)\) & \((-2.982,0.009)\)
\end{tabular}
\begin{tabular}{rrrrr} 
\\
1 & \((1.192 \mathrm{E}-07,-0.000 \mathrm{E}+00)\) & \(\left(\begin{array}{l}\text { RES } \\
2\end{array}(1.192 \mathrm{E}-07,-2.384 \mathrm{E}-07)\right.\) & \((-5.329 \mathrm{E}-08,-5.329 \mathrm{E}-08)\) & \((1.318 \mathrm{E}-07,-2.258 \mathrm{E}-07)\) \\
3 & \((2.384 \mathrm{E}-07,8.259 \mathrm{E}-08)\) & \((2.390 \mathrm{E}-07,-3.309 \mathrm{E}-08)\) & \((2.395 \mathrm{E}-07,1.015 \mathrm{E}-07)\) \\
4 & \((-2.384 \mathrm{E}-07,2.814 \mathrm{E}-14)\) & \((-8.240 \mathrm{E}-08,-8.790 \mathrm{E}-09)\) & \((-1.648 \mathrm{E}-07,-1.758 \mathrm{E}-08)\) \\
5 & \((-2.384 \mathrm{E}-07,-1.401 \mathrm{E}-08)\) & \((-2.813 \mathrm{E}-07,6.981 \mathrm{E}-09)\) & \((-3.241 \mathrm{E}-07,-2.795 \mathrm{E}-08)\)
\end{tabular}

\section*{LFDDH}

Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\)
Cholesky factorization of the matrix.

\section*{Required Arguments}

FACT - Complex N by N matrix containing the \(R^{H} R\) factorization of the coefficient matrix A as output from routine LFCDH/DLFCDH or LFTDH/DLFTDH. (Input)

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq|\operatorname{DET} 1|<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output) The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\mathrm{DET} 2}\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDDH (FACT, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDDH and D_LFDDH.

\section*{FORTRAN 77 Interface}

Single: CALL LFDDH (N, FACT, LDFACT, DET1, DET2)
Double: The double precision name is DLFDDH.

\section*{Description}

Routine LFDDH computes the determinant of a complex Hermitian positive definite coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either LFCDH or LFTDH. The formula \(\operatorname{det} A=\operatorname{det} R^{H} \operatorname{det}\) \(R=(\operatorname{det} R)^{2}\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} R=\prod_{i=1}^{N} R_{i i}
\]
(The matrix \(R\) is stored in the upper triangle of FACT.)
LFDDH is based on the LINPACK routine CPODI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a complex Hermitian positive definite \(3 \times 3\) matrix.
```

USE LFDDH_INT
USE LFTDH_INT
USE UMACH_INT
! - Declare variables
INTEGER LDA, LDFACT, NOUT
PARAMETER (LDA=3, LDFACT=3)
REAL DET1, DET2
COMPLEX A(LDA,LDA), FACT (LDFACT,LDFACT)
Set values for A

```
```

! A = ( 6.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 7.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i 11.0+0.0i )
DATA A / (6.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (7.0,0.0),\&
(-5.0,-1.0),(4.0,0.0),(-5.0,1.0), (11.0,0.0)/
CALL LFTDH (A, FACT)
CALL LFDDH (FACT, DET1, DET2)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is ',F6.3,' * 10**',F2.0)
END

```

\section*{Output}
```

The determinant of A is 1.400 * 10**2.

```

\section*{LSAHF}

Solves a complex Hermitian system of linear equations with iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian linear system. (Input)
Only the upper triangle of \(A\) is referenced.
B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSAHF (A, B, X [, ...])

Specific: The specific interface names are S_LSAHF and D_LSAHF.

\section*{FORTRAN 77 Interface}

Single: CALL LSAHF (N, A, LDA, B, X)
Double: \(\quad\) The double precision name is DLSAHF.

\section*{Description}

Routine LSAHF solves systems of linear algebraic equations having a complex Hermitian indefinite coefficient matrix. It first uses the routine LFCHF to compute a \(U D U^{H}\) factorization of the coefficient matrix and to estimate the condition number of the matrix. \(D\) is a block diagonal matrix with blocks of order 1 or 2 and \(U\) is a matrix composed of the product of a permutation matrix and a unit upper triangular matrix. The solution of the linear system is then found using the iterative refinement routine LFIHF.

LSAHF fails if a block in \(D\) is singular or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSAHF solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2AHF/DL2AHF. The reference is:

CALL L2AHF (N, A, LDA, B, X, FACT, IPVT, CWK)
The additional arguments are as follows:
FACT - Complex work vector of length \(\mathrm{N}^{2}\) containing information about the \(U D U^{H}\) factorization of \(A\) on output.

IPVT - Integer work vector of length \(N\) containing the pivoting information for the factorization of A on output.

CWK - Complex work vector of length n.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix singular.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2AHF the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSAHF. Additional memory allocation for FACT and option value restoration are done automatically in LSAHF. Users directly calling L2AHF can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSAHF or L2AHF. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSAHF temporarily replaces IVAL(2) by IVAL(1). The routine L2CHF computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CHF skips this computation. LSAHF restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has complex Hermitian form and the right-hand-side vector \(b\) has three elements.
```

USE LSAHF_INT
USE WRCRN__INT
INTEGER LDA, N
PARAMETER (LDA=3,N=3)
COMPLEX A(LDA,LDA), B(N), X(N)
Set values for A and B
A =( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
(4.0+0.0i -5.0-1.0i -2.0+0.0i )
B = ( 7.0+32.0i -39.0-21.0i 51.0+9.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/
DATA B/(7.0,32.0), (-39.0,-21.0), (51.0,9.0)/
CALL LSAHF (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```
!
!

\section*{Output}

X
\((2.00,1.00)^{1}(-10.00,-1.00)^{2} \quad(3.00,5.00)^{3}\)

\section*{LSLHF}

Solves a complex Hermitian system of linear equations without iterative refinement.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian linear system. (Input)
Only the upper triangle of A is referenced.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLHF (A, B, X [,...])
Specific: The specific interface names are S_LSLHF and D_LSLHF.

\section*{FORTRAN 77 Interface}

Single: CALL LSLHF (N, A, LDA, B, X)
Double: The double precision name is DLSLHF.

\section*{Description}

Routine LSLHF solves systems of linear algebraic equations having a complex Hermitian indefinite coefficient matrix. It first uses the routine LFCHF to compute a \(U D U^{H}\) factorization of the coefficient matrix. \(D\) is a block diagonal matrix with blocks of order 1 or 2 and \(U\) is a matrix composed of the product of a permutation matrix and a unit upper triangular matrix.

The solution of the linear system is then found using the routine LFSHF. LSLHF fails if a block in \(D\) is singular. This occurs only if \(A\) is singular or very close to a singular matrix. If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSAHF be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LHF/DL2LHF. The reference is:

CALL L2LHF (N, A, LDA, B, X, FACT, IPVT, CWK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}\) - Complex work vector of length \(\mathrm{N}^{2}\) containing information about the \(U D U^{H}\) factorization of A on output.

IPVT - Integer work vector of length N containing the pivoting information for the factorization of A on output.

CWK - Complex work vector of length N.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix singular.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LHF the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSLHF. Additional memory allocation for FACT and option value restoration are done automatically in LSLHF. Users directly calling L2 LHF can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLHF or L2LHF. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLHF temporarily replaces IVAL(2) by IVAL(1). The routine L2CHF computes the condition number if IVAL \((2)=2\). Otherwise L2CHF skips this computation. LSLHF restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of three linear equations is solved. The coefficient matrix has complex Hermitian form and the right-hand-side vector \(b\) has three elements.
```

USE LSLHF INT
USE WRCRN_INT
INTEGER }\quad\mathrm{ LDA, N
COMPLEX A(LDA,LDA), B(N), X(N)
Set values for A and B
A = ( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i -2.0+0.0i )
B = ( 7.0+32.0i -39.0-21.0i 51.0+9.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/
DATA B/(7.0,32.0), (-39.0,-21.0), (51.0,9.0)/
CALL LSLHF (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{ccc} 
\\
\\
\((2.00, ~ 1.00)^{1}\) & \((-10.00,-1.00)^{2}\) & \((3.00,5.00)^{3}\)
\end{tabular}

\section*{LFCHF}

Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian linear system. (Input)
Only the upper triangle of A is referenced.
FACT - Complex \(n\) by n matrix containing the information about the factorization of the Hermitian matrix A. (Output)
Only the upper triangle of FACT is used. If A is not needed, A and FACT can share the same storage locations.

IPVT - Vector of length N containing the pivoting information for the factorization. (Output)
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCHF (A, FACT, IPVT, RCOND [, ...])

Specific: The specific interface names are \(S_{\text {_ }}\) LFCHF and \(D_{-}\)LFCHF.

\section*{FORTRAN 77 Interface}

Single: CALL LFCHF (N, A, LDA, FACT, LDFACT, IPVT, RCOND)
Double: The double precision name is DLFCHF.

\section*{Description}

Routine LFCHF performs a \(U D U^{H}\) factorization of a complex Hermitian indefinite coefficient matrix. It also estimates the condition number of the matrix. The \(U D U^{H}\) factorization is called the diagonal pivoting factorization.

The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCHF fails if \(A\) is singular or very close to a singular matrix.
The \(U D U^{H}\) factors are returned in a form that is compatible with routines LFIHF, LFSHF and LFDHF. To solve systems of equations with multiple right-hand-side vectors, use LFCHF followed
by either LFIHF or LFSHF called once for each right-hand side. The routine LFDHF can be called to compute the determinant of the coefficient matrix after LFCHF has performed the factorization.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CHF/DL2CHF. The reference is:

CALL L2CHF (N, A, LDA, FACT, LDFACT, IPVT, RCOND, CWK)
The additional argument is:

CWK - Complex work vector of length n.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix is singular.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.

\section*{Example}

The inverse of a \(3 \times 3\) complex Hermitian matrix is computed. LFCHF is called to factor the matrix and to check for singularity or ill-conditioning. LFIHF is called to determine the columns of the inverse.
```

USE LFCHF INT
USE UMACH_INT
USE LFIHF INT
USE WRCRN INT
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
INTEGER IPVT(N), NOUT
REAL RCOND
COMPLEX A(LDA,LDA), AINV (LDA,N), FACT(LDA,LDA), RJ (N), RES (N)
Set values for A
A = ( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i -2.0+0.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0.0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/

```
\(!\)
```

! Set output unit number
CALL UMACH (2, NOUT)
!
!
CALL LFCHF (A, FACT, IPVT, RCOND)
Print the estimate of the condition
number
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0E0,0.0E0)
DO 10 J=1, N
RJ(J) = (1.0E0, 0.0E0)
RJ is the J-th column of the identity
matrix so the following LFIHF
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFIHF (A, FACT, IPVT, RJ, AINV(:,J), RES)
RJ (J) = (0.0E0, 0.0E0)
1 0 ~ C O N T I N U E
! Print the inverse
CALL WRCRN ('AINV', AINV)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND < 0.25
L1 Condition number < 6.0
AINV

|  |  |  |  |
| ---: | ---: | ---: | ---: | ---: |
| 1 | $(0.2000,0.0000)$ | $(0.1200,0.0400)$ | $(0.0800,-0.0400)$ |
| 2 | $(0.1200,-0.0400)$ | $(0.1467,0.0000)$ | $(-0.1267,-0.0067)$ |
| 3 | $(0.0800,0.0400)$ | $(-0.1267,0.0067)$ | $(-0.0267,0.0000)$ |

```

\section*{LFTHF}

Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian linear system. (Input)
Only the upper triangle of \(A\) is referenced.

FACT - Complex n by n matrix containing the information about the factorization of the Hermitian matrix A. (Output)
Only the upper triangle of FACT is used. If A is not needed, A and FACT can share the same storage locations.
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the factorization. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTHF (A, FACT, IPVT [,...])

Specific: The specific interface names are S_LFTHF and D_LFTHF.

\section*{FORTRAN 77 Interface}

Single: CALL LFTHF (N, A, LDA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTHF.

\section*{Description}

Routine LFTHF performs a \(U D U^{H}\) factorization of a complex Hermitian indefinite coefficient matrix. The \(U D U^{H}\) factorization is called the diagonal pivoting factorization.

LFTHF fails if \(A\) is singular or very close to a singular matrix.
The \(U D U^{H}\) factors are returned in a form that is compatible with routines LFIHF, LFSHF and LFDHF. To solve systems of equations with multiple right-hand-side vectors, use LFTHF followed by either LFIHF or LFSHF called once for each right-hand side. The routine LFDHF can be called to compute the determinant of the coefficient matrix after LFTHF has performed the factorization.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}

Informational errors
Type Code

4 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
\(4 \quad 2\) The input matrix is singular.
\(4 \quad 4\) The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFTHF is called to factor the matrix and check for singularity. LFSHF is called to determine the columns of the inverse.
```

    USE LFTHF_INT
    USE LFSHF INT
    USE WRCRN_INT
    !
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
INTEGER IPVT(N)
COMPLEX A(LDA,LDA), AINV (LDA,N), FACT (LDA,LDA), RJ (N)
Set values for A
A = ( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i -2.0+0.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0),(4.0,0.0),(-5.0,1.0), (-2.0,0.0)/
Factor A
CALL LFTHF (A, FACT, IPVT)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0E0,0.0EO)
DO 10 J=1, N
RJ (J) = (1.0E0, 0.0EO)
RJ is the J-th column of the identity
matrix so the following LFSHF
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFSHF (FACT, IPVT, RJ, AINV(:,J))
RJ (J) = (0.0E0, 0.0EO)
10 CONTINUE
CALL WRCRN ('AINV', AINV)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{4}{|c|}{AINV} \\
\hline 1 & & 2 & 3 \\
\hline ( 0.2000, 0.0000) & ( 0.1200, & \(0.0400)\) & ( 0.0800, -0.0400) \\
\hline ( 0.1200, -0.0400) & ( 0.1467, & \(0.0000)\) & (-0.1267, -0.0067) \\
\hline ( 0.0800, 0.0400) & (-0.1267, & \(0.0067)\) & (-0.0267, 0.0000) \\
\hline
\end{tabular}

\section*{LFSHF}

Solves a complex Hermitian system of linear equations given the \(U D U^{H}\) factorization of the coefficient matrix.

\section*{Required Arguments}

FACT - Complex n by n matrix containing the factorization of the coefficient matrix A as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input)
Only the upper triangle of FACT is used.
IPVT - Vector of length \(N\) containing the pivoting information for the factorization of A as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSHF (FACT, IPVT, B, X [, ...])

Specific: The specific interface names are S_LFSHF and D_LFSHF.

\section*{FORTRAN 77 Interface}

Single: CALL LFSHF (N, FACT, LDFACT, IPVT, B, X)
Double: The double precision name is DLFSHF.

\section*{Description}

Routine LFSHF computes the solution of a system of linear algebraic equations having a complex Hermitian indefinite coefficient matrix.
To compute the solution, the coefficient matrix must first undergo a \(U D U^{H}\) factorization. This may be done by calling either LFCHF or LFTHF.

LFSHF and LFIHF both solve a linear system given its \(U D U^{H}\) factorization. LFIHF generally takes more time and produces a more accurate answer than LFSHF. Each iteration of the iterative refinement algorithm used by LFIHF calls LFSHF.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Example}

A set of linear systems is solved successively. LFTHF is called to factor the coefficient matrix. LFSHF is called to compute the three solutions for the three right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCHF to perform the factorization, and LFIHF to compute the solutions.
```

USE LFSHF INT
USE WRCRN }\mp@subsup{}{}{-}\mathrm{ INT
USE LFTHF_INT
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
INTEGER IPVT (N), I
COMPLEX A(LDA,LDA), B(N,3), X(N,3), FACT(LDA,LDA)
Set values for A and B
A = ( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
(4.0+0.0i -5.0-1.0i -2.0+0.0i )
B = ( 7.0+32.0i -6.0+11.0i -2.0-17.0i )
(-39.0-21.0i -5.5-22.5i 4.0+10.0i )
( 51.0+ 9.0i 16.0+17.0i -2.0+12.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/
DATA B/ (7.0,32.0), (-39.0,-21.0), (51.0,9.0), (-6.0,11.0),\&
(-5.5,-22.5), (16.0,17.0), (-2.0,-17.0), (4.0,10.0),\&
(-2.0,12.0)/
CALL LFTHF (A, FACT, IPVT)
DO 10 I=1, 3
CALL LFSHF (FACT, IPVT, B(:,I), X(:,I))
1 0 ~ C O N T I N U E
CALL WRCRN ('X', X)
END

```
!

\section*{Output}

X
\(1(2.00,1.00)^{1}(1.00,0.00)^{2}(0.00,-1.00)^{3}\)
```

(-10.00, -1.00) ( -3.00, -4.00) ( 0.00, -2.00)
( 3.00, 5.00) ( -0.50, 3.00) ( 0.00, -3.00)

```

\section*{LFIHF}

Uses iterative refinement to improve the solution of a complex Hermitian system of linear equations.

\section*{Required Arguments}

A-Complex n by n matrix containing the coefficient matrix of the Hermitian linear system. (Input)
Only the upper triangle of A is referenced.
FACT - Complex \(n\) by N matrix containing the factorization of the coefficient matrix A as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input) Only the upper triangle of FACT is used.

IPVT - Vector of length \(N\) containing the pivoting information for the factorization of \(A\) as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution. (Output)
RES - Complex vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIHF (A, FACT, IPVT, B, X, RES [,...])
Specific: The specific interface names are S_LFIHF and D_LFIHF.

\section*{FORTRAN 77 Interface}

Single: CALL LFIHF (N, A, LDA, FACT, LDFACT, IPVT, B, X, RES)
Double: The double precision name is DLFIHF.

\section*{Description}

Routine LFIHF computes the solution of a system of linear algebraic equations having a complex Hermitian indefinite coefficient matrix.

Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo a \(U D U^{H}\) factorization. This may be done by calling either LFCHF or LFTHF.
Iterative refinement fails only if the matrix is very ill-conditioned.
LFIHF and LFSHF both solve a linear system given its \(U D U^{H}\) factorization. LFIHF generally takes more time and produces a more accurate answer than LFSHF. Each iteration of the iterative refinement algorithm used by LFIHF calls LFSHF.

\section*{Comments}

Informational error
Type Code
33 The input matrix is too ill-conditioned for iterative refinement to be effective.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding \(0.2+0.2 i\) to the second element.
```

USE LFIHF INT
USE UMACH_INT
USE LFCHF INT
USE WRCRN_INT
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
INTEGER IPVT(N), NOUT
REAL RCOND
COMPLEX A(LDA,LDA), B(N), X(N), FACT(LDA,LDA), RES(N)
Set values for A and B
A = ( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i -2.0+0.0i )

```
!
```

! B = ( 7.0+32.0i -39.0-21.0i 51.0+9.0i )
DATA A/ (3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/
DATA B/(7.0,32.0), (-39.0,-21.0), (51.0,9.0)/
Set output unit number
CALL UMACH (2, NOUT)
Factor A and compute the estimate
of the reciprocal condition number
CALL LFCHF (A, FACT, IPVT, RCOND)
WRITE (NOUT,99998) RCOND, 1.OEO/RCOND
Solve, then perturb right-hand side
DO 10 I=1, 3
CALL LFIHF (A, FACT, IPVT, B, X, RES)
Print results
WRITE (NOUT,99999) I
CALL WRCRN ('X', X, 1, N, 1)
CALL WRCRN ('RES', RES, 1, N, 1)
B(2) = B(2) + (0.2E0, 0.2E0)
10 CONTINUE
!
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
99999 FORMAT (//,' For problem ', I1)
END

```

\section*{Output}
```

RCOND < 0.25
L1 Condition number < 5.0
For problem 1
X
$(2.00,1.00)^{1}(-10.00,-1.00)^{2}(3.00,5.00)^{3}$
RES
(2.384E-07,-4.768E-07) ( ( 0.000E+00,-3.576E-07) 2 (-1.421E-14, 1.421E-14)
For problem 2
X
$(2.016,1.032)^{1} \quad(-9.971,-0.971)^{2} \quad(2.973,4.976)^{3}$
RES
(2.098E-07,-1.764E-07) ( 6.231E-07,-1.518E-07) 2)
For problem 3
X
$(2.032,1.064)^{1} \quad(-9.941,-0.941)^{2} \quad(2.947,4.952)^{3}$
RES
$(4.196 \mathrm{E}-07,-3.529 \mathrm{E}-07)^{1} \quad(2.925 \mathrm{E}-07,-3.632 \mathrm{E}-07)^{2} \quad(2.543 \mathrm{E}-07,3.242 \mathrm{E}-07)^{3}$

```

\section*{LFDHF}

Computes the determinant of a complex Hermitian matrix given the \(U D U^{H}\) factorization of the matrix.

\section*{Required Arguments}

FACT - Complex n by n matrix containing the factorization of the coefficient matrix A as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input) Only the upper triangle of FACT is used.

IPVT - Vector of length N containing the pivoting information for the factorization of \(A\) as output from routine LFCHF/DLFCHF or LFTHF/DLFTHF. (Input)

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\text {DET2 }}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDHF (FACT, IPVT, DET1, DET2 [, ...])
Specific: The specific interface names are S_LFDHF and D_LFDHF.

\section*{FORTRAN 77 Interface}

Single: CALL LFDHF (N, FACT, LDFACT, IPVT, DET1, DET2)
Double: The double precision name is DLFDHF.

\section*{Description}

Routine LFDHF computes the determinant of a complex Hermitian indefinite coefficient matrix. To compute the determinant, the coefficient matrix must first undergo a \(U D U^{H}\) factorization. This may be done by calling either LFCHF or LFTHF since \(\operatorname{det} U= \pm 1\), the formula
\(\operatorname{det} A=\operatorname{det} U \operatorname{det} D \operatorname{det} U^{H}=\operatorname{det} D\) is used to compute the \(\operatorname{determinant}\). \(\operatorname{det} D\) is computed as the product of the determinants of its blocks.

LFDHF is based on the LINPACK routine CSIDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a complex Hermitian \(3 \times 3\) matrix.
```

        USE LFDHF_INT
        USE LFTHF_INT
        USE UMACH_INT
    ! - Declare variables
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
INTEGER IPVT (N), NOUT
REAL DET1, DET2
COMPLEX A(LDA,LDA), FACT (LDA,LDA)
Set values for A
A =( 3.0+0.0i 1.0-1.0i 4.0+0.0i )
( 1.0+1.0i 2.0+0.0i -5.0+1.0i )
( 4.0+0.0i -5.0-1.0i -2.0+0.0i )
DATA A/(3.0,0.0), (1.0,1.0), (4.0,0.0), (1.0,-1.0), (2.0,0.0),\&
(-5.0,-1.0), (4.0,0.0), (-5.0,1.0), (-2.0,0.0)/
Factor A
CALL LFTHF (A, FACT, IPVT)
CALL LFDHF (FACT, IPVT, DET1, DET2)
Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant is', F5.1, ' * 10**', F2.0)
END

```

\section*{Output}

The determinant is -1.5 * 10**2.

\section*{LSLTR}

Solves a real tridiagonal system of linear equations.

\section*{Required Arguments}
\(\boldsymbol{C}\) - Vector of length N containing the subdiagonal of the tridiagonal matrix in \(\mathrm{C}(2)\) through C(N). (Input/Output)
On output C is destroyed.
\(\boldsymbol{D}\) - Vector of length \(N\) containing the diagonal of the tridiagonal matrix. (Input/Output) On output D is destroyed.
\(\boldsymbol{E}\) - Vector of length N containing the superdiagonal of the tridiagonal matrix in \(\mathrm{E}(1)\) through \(\mathrm{E}(\mathrm{N}-1)\). (Input/Output)
On output E is destroyed.
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system on entry and the solution vector on return. (Input/Output)

\section*{Optional Arguments}
\(N\) - Order of the tridiagonal matrix. (Input) Default: \(\mathrm{N}=\operatorname{size}(\mathrm{C}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLTR (C, D, E, B [, ...])
Specific: The specific interface names are S_LSLTR and D_LSLTR.

\section*{FORTRAN 77 Interface}

Single: CALL LSLTR (N, C, D, E, B)
Double: \(\quad\) The double precision name is DLSLTR.

\section*{Description}

Routine LSLTR factors and solves the real tridiagonal linear system \(A x=b\). LSLTR is intended just for tridiagonal systems. The coefficient matrix does not have to be symmetric. The algorithm is Gaussian elimination with partial pivoting for numerical stability. See Dongarra (1979), LINPACK subprograms SGTSL/DGTSL, for details. When computing on vector or parallel computers the cyclic reduction algorithm, LSLCR, should be considered as an alternative method to solve the system.

\section*{Comments}

Informational error
Type Code
42 An element along the diagonal became exactly zero during execution.

\section*{Example}

A system of \(n=4\) linear equations is solved.
```

USE LSLTR_INT
USE WRRRL__INT

```
```

! Declaration of variables
INTEGER N
PARAMETER (N=4)
REAL B(N), C(N), D(N), E(N)
CHARACTER CLABEL(1)*6, FMT*8, RLABEL(1)*4
DATA FMT/'(E13.6)'/
DATA CLABEL/'NUMBER'/
DATA RLABEL/'NONE' /
C(*), D(*), E(*), and B(*)
contain the subdiagonal, diagonal,
superdiagonal and right hand side.
DATA C/0.0, 0.0, -4.0, 9.0/, D/6.0, 4.0, -4.0, -9.0/
DATA E/-3.0, 7.0, -8.0, 0.0/, B/48.0, -81.0, -12.0, -144.0/
CALL LSLTR (C, D, E, B)
CALL WRRRL ('Solution:', B, RLABEL, CLABEL, 1, N, 1, FMT=FMT)
END

```

\section*{Output}

Solution:
\begin{tabular}{cccc}
1 & 2 & 3 & 4 \\
\(0.400000 \mathrm{E}+01\) & \(-0.800000 \mathrm{E}+01\) & \(-0.700000 \mathrm{E}+01\) & \(0.900000 \mathrm{E}+01\)
\end{tabular}

\section*{LSLCR}

Computes the \(L D U\) factorization of a real tridiagonal matrix \(A\) using a cyclic reduction algorithm.

\section*{Required Arguments}
\(\boldsymbol{C}\) - Array of size 2 N containing the upper codiagonal of the N by N tridiagonal matrix in the entries \(\mathrm{C}(1), \ldots, \mathrm{C}(\mathrm{N}-1)\). (Input/Output)

A-Array of size 2 N containing the diagonal of the N by N tridiagonal matrix in the entries \(A(1), \ldots, A(N)\). (Input/Output)
\(\boldsymbol{B}\) - Array of size 2 N containing the lower codiagonal of the N by N tridiagonal matrix in the entries \(B(1), \ldots, B(N-1)\). (Input/Output)
\(\boldsymbol{Y}\) - Array of size 2 N containing the right hand side for the system \(A x=y\) in the order \(\mathrm{Y}(1)\), \(\ldots, \mathrm{Y}(\mathrm{N})\). (Input/Output) The vector x overwrites Y in storage.
\(\boldsymbol{U}\) - Array of size 2 N of flags that indicate any singularities of A. (Output) A value \(U(I)=1\). means that a divide by zero would have occurred during the factoring. Otherwise \(U(I)=0\).

IR - Array of integers that determine the sizes of loops performed in the cyclic reduction algorithm. (Output)

IS - Array of integers that determine the sizes of loops performed in the cyclic reduction algorithm. (Output)
The sizes of IR and IS must be at least \(\log _{2}(\mathrm{~N})+3\).

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
N must be greater than zero
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{C}, 1)\).
\(\boldsymbol{I J O B}\) - Flag to direct the desired factoring or solving step. (Input)
Default: \(\mathrm{IJOB}=1\).
IJOB Action
\(1 \quad\) Factor the matrix \(A\) and solve the system \(A x=y\), where \(y\) is stored in array Y.

2 Do the solve step only. Use \(y\) from array Y. (The factoring step has already been done.)

3 Factor the matrix \(A\) but do not solve a system.
\(4,5,6\) Same meaning as with the value \(I J O B=3\). For efficiency, no error checking is done on the validity of any input value.

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCR (C, A, B, Y, U, IR, IS [, ...])
Specific: The specific interface names are \(S_{-}\)LSLCR and \(D_{-}\)LSLCR.

\section*{FORTRAN 77 Interface}

Single: CALL LSLCR (N, C, A, B, IJOB, Y, U, IR, IS)
Double: The double precision name is DLSLCR.

\section*{Description}

Routine LSLCR factors and solves the real tridiagonal linear system \(A x=y\). The matrix is decomposed in the form \(\mathrm{A}=L D U\), where \(L\) is unit lower triangular, \(U\) is unit upper triangular, and \(D\) is diagonal. The algorithm used for the factorization is effectively that described in Kershaw (1982). More details, tests and experiments are reported in Hanson (1990).

LSLCR is intended just for tridiagonal systems. The coefficient matrix does not have to be symmetric. The algorithm amounts to Gaussian elimination, with no pivoting for numerical stability, on the matrix whose rows and columns are permuted to a new order. See Hanson (1990) for details. The expectation is that LSLCR will outperform either LSLTR or LSLPB on vector or parallel computers. Its performance may be inferior for small values of \(n\), on scalar computers, or high-performance computers with non-optimizing compilers.

\section*{Example}

A system of \(n=1000\) linear equations is solved. The coefficient matrix is the symmetric matrix of the second difference operation, and the right-hand-side vector \(y\) is the first column of the identity matrix. Note that \(a_{n, n}=1\). The solution vector will be the first column of the inverse matrix of \(A\). Then a new system is solved where \(y\) is now the last column of the identity matrix. The solution vector for this system will be the last column of the inverse matrix.
```

    USE LSLCR_INT
    USE UMACH_INT
    !
INTEGER LP, N, N2
PARAMETER (LP=12,N=1000, N2=2*N)
!
INTEGER I, IJOB, IR(LP), IS(LP), NOUT
REAL A(N2), B(N2), C(N2), U(N2), Y1 (N2), Y2 (N2)
!
DO 10 I=1, N - 1
C(I) = -1.E0
A(I) = 2.E0
B(I) = -1.E0
Y1(I+1) = 0.E0
Y2(I) = 0.E0
10 CONTINUE
A(N) = 1.EO
Y1(1) = 1.E0
Y2(N) = 1.E0
!
!
Obtain decomposition of matrix and
solve the first system:
IJOB = 1
CALL LSLCR (C, A, B, Y1, U, IR, IS, IJOB=IJOB)
Solve the second system with the
decomposition ready:
IJOB = 2
CALL LSLCR (C, A, B, Y2, U, IR, IS, IJOB=IJOB)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The value of n is: ', N
WRITE (NOUT,*) ' Elements 1, n of inverse matrix columns 1 '//\&
END

```

\section*{Output}
```

The value of n is: 1000
Elements 1, n of inverse matrix columns 1 and n: 1.00000 1000.000

```

\section*{LSARB}

Solves a real system of linear equations in band storage mode with iterative refinement.

\section*{Required Arguments}
\(\boldsymbol{A}-(\mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the N by N banded coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A,1).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} \mathrm{X}=\mathrm{B}\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSARB (A, NLCA, NUCA, B, X [, ...])

Specific: The specific interface names are S_LSARB and D_LSARB.

\section*{FORTRAN 77 Interface}

Single: CALL LSARB (N, A, LDA, NLCA, NUCA, B, IPATH, X)
Double: The double precision name is DLSARB.

\section*{Description}

Routine LSARB solves a system of linear algebraic equations having a real banded coefficient matrix. It first uses the routine LFCRB to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using the iterative refinement routine LFIRB.

LSARB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if A is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSARB solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2ARB/DL2ARB. The reference is:

CALL L2ARB (N, A, LDA, NLCA, NUCA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-\) Work vector of length \((2 * \mathrm{NLCA}+\mathrm{NUCA}+1) \times \mathrm{N}\) containing the \(L U\) factorization of A on output.

IPVT - Work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.

WK - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ARB the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSARB.
Additional memory allocation for FACT and option value restoration are done automatically in LSARB. Users directly calling L2ARB can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing
applications that use LSARB or L2ARB. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSARB temporarily replaces IVAL(2) by IVAL(1). The routine L2CRB computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CRB skips this computation. LSARB restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has real banded form with 1 upper and 1 lower codiagonal. The right-hand-side vector \(b\) has four elements.
```

USE LSARB_INT
USE WRRRN_INT
INTEGER LDA, N, NLCA, NUCA
PARAMETER (LDA=3, N=4, NLCA=1, NUCA=1)
REAL A(LDA,N), B (N), X(N)
Set values for A in band form, and B
A =( lllll
( 2.0 1.0 -1.0 1.0)
($$
\begin{array}{llll}{-3.0}&{0.0}&{2.0}&{0.0}\end{array}
$$)
B = ( 3.0 1.0 11.0 -2.0)
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0,\&
2.0, 1.0, 0.0/
DATA B/3.0, 1.0, 11.0, -2.0/
CALL LSARB (A, NLCA, NUCA, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrrrr} 
& & X & \\
2.000 & 1.000 & -3.000 & 4.000
\end{tabular}

\section*{LSLRB}


Solves a real system of linear equations in band storage mode without iterative refinement.

\section*{Required Arguments}
\(\boldsymbol{A}-(\mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the N by N banded coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
\[
\text { IPATH }=1 \text { means the system } A X=B \text { is solved. }
\]

IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: \(\operatorname{IPATH}=1\)

\section*{FORTRAN 90 Interface}

Generic: CALL LSLRB (A, NLCA, NUCA, B, X [,...])

Specific: The specific interface names are S_LSLRB and D_LSLRB.

\section*{FORTRAN 77 Interface}

Single: CALL LSLRB (N, A, LDA, NLCA, NUCA, B, IPATH, X)
Double: The double precision name is DLSLRB

\section*{ScaLAPACK Interface}

Generic: CALL LSLRB (A0, NLCA, NUCA, B0, X0 [, ...])

Specific: The specific interface names are \(S_{\_}\)LSLRB and D_LSLRB.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

\section*{Description}

Routine LSLRB solves a system of linear algebraic equations having a real banded coefficient matrix. It first uses the routine LFCRB to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using LFSRB. LSLRB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This occurs only if \(A\) is singular or very close to a singular matrix. If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSARB be used.
The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LRB/DL2LRB. The reference is:

CALL L2LRB (N, A, LDA, NLCA, NUCA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-(2 \times \mathrm{NLCA}+\mathrm{NUCA}+1) \times \mathrm{N}\) containing the \(L U\) factorization of A on output. If A is not needed, A can share the first (NLCA \(+\mathrm{NUCA}+1) * N\) storage locations with FACT.

IPVT - Work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(W K-\) Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LRB the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSLRB.
Additional memory allocation for FACT and option value restoration are done automatically in LSLRB. Users directly calling L2 LRB can allocate additional space for \(\operatorname{FACT}\) and set \(\operatorname{IVAL}(3)\) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing
applications that use LSLRB or L2LRB. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLRB temporarily replaces IVAL(2) by IVAL(1). The routine L2CRB computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise \(\operatorname{L2CRB}\) skips this computation. LSLRB restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{ScaLAPACK Usage Notes}

The arguments which differ from the standard version of this routine are:
A0 - \(\left(2 *_{N L C A}+2 * N U C A+1\right)\) by MXCOL local matrix containing the local portions of the distributed matrix \(A\). A contains the N by N banded coefficient matrix in band storage mode. (Input)

B0 - Local vector of length MXCOL containing the local portions of the distributed vector B. \(B\) contains the right-hand side of the linear system. (Input)

X0 - Local vector of length MXCOL containing the local portions of the distributed vector X . X contains the solution to the linear system. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has real banded form with 1 upper and 1 lower codiagonal. The right-hand-side vector \(b\) has four elements.
```

USE LSLRB INT
USE WRRRN_INT
INTEGER LDA, N, NLCA, NUCA
PARAMETER (LDA=3, N=4, NLCA=1, NUCA=1)
REAL A(LDA,N), B(N), X(N)
Set values for A in band form, and B
A =( ( 0.0 -1.0 -2.0 2.0)
( 2.0 1.0 -1.0 1.0)
($$
\begin{array}{llll}{-3.0}&{0.0}&{2.0}&{0.0}\end{array}
$$)
B = ($$
\begin{array}{lllll}{3.0}&{1.0}&{11.0}&{-2.0)}\end{array}
$$)
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0,\&
2.0, 1.0, 0.0/
DATA B/3.0, 1.0, 11.0, -2.0/
CALL LSLRB (A, NLCA, NUCA, B, X)

```
!
```

! Print results
CALL WRRRN ('X', X, 1, N, 1)
!
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{4}{|c|}{X} \\
\hline 1 & 2 & 3 & 4 \\
\hline 2.000 & 1.000 & -3.000 & 4.000 \\
\hline
\end{tabular}

\section*{ScaLAPACK Example}

The same system of four linear equations is solved as a distributed computing example. The coefficient matrix has real banded form with 1 upper and 1 lower codiagonal. The right-hand-side vector \(b\) has four elements. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.
```

USE MPI SETUP INT
USE LSLRBB_INT
USE WRRRN INT
USE SCALAP\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER LDA, M, N, NLCA, NUCA, NRA, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA
REAL, ALLOCATABLE :: A(:,:), B(:), X(:)
REAL, ALLOCATABLE :: A0(:,:), B0(:), X0(:)
PARAMETER (LDA=3, N=6, NLCA=1, NUCA=1)
Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
ALLOCATE (A (LDA,N), B(N), X(N))
Set values for A and B
A(1,:) = (/ 0.0, 0.0, -3.0, 0.0, -1.0, -3.0/)
A(2,:) = (/ 10.0, 10.0, 15.0, 10.0, 1.0, 6.0/)
A(3,:) = (/ 0.0, 0.0, 0.0, -5.0, 0.0, 0.0/)
B = (/ 10.0, 7.0, 45.0, 33.0, -34.0, 31.0/)
ENDIF
NRA = NLCA + NUCA + 1
M = 2*NLCA + 2*NUCA + 1
Set up a 1D processor grid and define
its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(M, N, .FALSE., .TRUE.)
Get the array descriptor entities MXLDA,
and MXCOL
CALL SCALAPACK_GETDIM(M, N, MP_MB, MP_NB, MXLDA, MXCOL)
Reset MXIDDA to M
MXLDA = M

```
! Set up the array descriptors
CALL DESCINIT (DESCA, NRA, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, INFO) CALL DESCINIT (DESCX, 1, N, 1, MP_NB, 0, 0, MP_ICTXT, 1, INFO)

Allocate space for the local arrays
ALLOCATE (A0 (MXLDA, MXCOL), B0 (MXCOL), X0 (MXCOL)) Map input arrays to the processor grid
CALL SCALAPACK_MAP (A, DESCA, A0)
CALL SCALAPACK MAP (B, DESCX, B0, 1, .FALSE.)
Solve the system of equations
CALL LSLRB (A0, NLCA, NUCA, B0, X0)
Unmap the results from the distributed arrays back to a non-distributed array. After the unmap, only Rank=0 has the full array.
CALL SCALAPACK_UNMAP (X0, DESCX, X, 1, .FALSE.)
Print results.
Only Rank=0 has the solution, X.
IF (MP_RANK .EQ. O) CALL WRRRN (' \(\mathrm{X}^{\prime}, \mathrm{X}, \mathrm{I}, \mathrm{N}, 1\) )
IF (M \(\bar{P}\) RANK .EQ. O) DEALLOCATE (A, B, X) DEALLOCATE (A0, B0, X0)
```

! Exit ScaLAPACK usage

```

CALL SCALAPACK_EXIT (MP_ICTXT)
Shut down MPI
\(\mathrm{MP}_{-}\)NPROCS \(=\mathrm{MP}_{-} \operatorname{SETUP}(\) 'FINAL')
END

\section*{Output}
```

            X
    ```
\begin{tabular}{rrrrrr}
1 & 2 & 3 & 4 & 5 & 6 \\
1.000 & 1.600 & 3.000 & 2.900 & -4.000 & 5.167
\end{tabular}

\section*{LFCRB}

Computes the \(L U\) factorization of a real matrix in band storage mode and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-(\mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the N by N matrix in band storage mode to be factored. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-(2 * \mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, A can share the first (NLCA \(+\mathrm{NUCA}+1)\) * N locations with FACT.

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCRB (A, NLCA, NUCA, FACT, IPVT, RCOND [, ...])
Specific: The specific interface names are \(S\) _LFCRB and D_LFCRB.

\section*{FORTRAN 77 Interface}

Single: CALL LFCRB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, RCOND)
Double: The double precision name is DLFCRB.

\section*{Description}

Routine LFCRB performs an \(L U\) factorization of a real banded coefficient matrix. It also estimates the condition number of the matrix. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm.

The \(\mathrm{L}_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LSCRB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) is singular or very close to a singular matrix. The \(L U\) factors are returned in a
form that is compatible with routines LFIRB, LFSRB and LFDRB. To solve systems of equations with multiple right-hand-side vectors, use LFCRB followed by either LFIRB or LFSRB called once for each right-hand side. The routine LFDRB can be called to compute the determinant of the coefficient matrix after LFCRB has performed the factorization.

Let \(F\) be the matrix FACT, let \(m_{l}=\) NLCA and let \(m_{u}=\) NUCA. The first \(m_{l}+m_{u}+1\) rows of F contain the triangular matrix \(U\) in band storage form. The lower \(\mathrm{m}_{l}\) rows of F contain the multipliers needed to reconstruct \(L^{-1}\).

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CRB/DL2CRB. The reference is:

CALL L2CRB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, RCOND, WK)

The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
42 The input matrix is singular.

\section*{Example}

The inverse of a \(4 \times 4\) band matrix with one upper and one lower codiagonal is computed. LFCRB is called to factor the matrix and to check for singularity or ill-conditioning. LFIRB is called to determine the columns of the inverse.
```

USE LFCRB INT
USE UMACH_INT
USE LFIRB_INT
USE WRRRN_INT
! Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
REAL A(LDA,N), AINV (N,N), FACT(LDFACT,N), RCOND, RJ(N), RES(N)
Set values for A in band form
A =( ( 0.0 -1.0 -2.0 2.0)
( 2.0 1.0 -1.0 1.0)
( -3.0 0.0 2.0 0.0)
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0,\&

```
```

        2.0, 1.0, 0.0/
    !
CALL LFCRB (A, NLCA, NUCA, FACT, IPVT, RCOND)
Print the reciprocal condition number
and the L1 condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
Set up the columns of the identity
matrix one at a time in RJ
RJ = 0.OEO
DO 10 J=1, N
RJ(J) = 1.0EO
RJ is the J-th column of the identity
matrix so the following LFIRB
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFIRB (A, NLCA, NUCA, FACT, IPVT, RJ, AINV(:,J), RES)
RJ(J) = 0.0E0
1 0 ~ C O N T I N U E
! Print results
CALL WRRRN ('AINV', AINV)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < . 07
L1 Condition number \(=25.0\)
\begin{tabular}{rrrrr} 
& & \multicolumn{4}{c}{ AINV } \\
& 1 & 2 & 3 & 4 \\
1 & -1.000 & -1.000 & 0.400 & -0.800 \\
2 & -3.000 & -2.000 & 0.800 & -1.600 \\
3 & 0.000 & 0.000 & -0.200 & 0.400 \\
4 & 0.000 & 0.000 & 0.400 & 0.200
\end{tabular}

\section*{LFTRB}

Computes the \(L U\) factorization of a real matrix in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{A}-(\mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the N by N matrix in band storage mode to be factored. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-(2 *\) NLCA \(+\mathrm{NUCA}+1)\) by N array containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, \(A\) can share the first \((N L C A+N U C A+1) * N\) locations with FACT.
IPVT - Vector of length n containing the pivoting information for the \(L U\) factorization. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTRB (A, NLCA, NUCA, FACT,,
Specific: The specific interface names are \(S\) _LFTRB and D_LFTRB.

\section*{FORTRAN 77 Interface}

Single: CALL LFTRB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTRB.

\section*{Description}

The routine LFTRB performs an \(L U\) factorization of a real banded coefficient matrix using Gaussian elimination with partial pivoting. \(A\) failure occurs if \(U\), the upper triangular factor, has a zero diagonal element. This can happen if \(A\) is close to a singular matrix. The \(L U\) factors are returned in a form that is compatible with routines LFIRB, LFSRB and LFDRB. To solve systems of equations with multiple right-hand-side vectors, use LFTRB followed by either LFIRB or LFSRB called once for each right-hand side. The routine LFDRB can be called to compute the determinant of the coefficient matrix after LFTRB has performed the factorization

Let \(m_{l}=\) NLCA, and let \(\mathrm{m}_{u}=\) NUCA. The first \(m_{l}+m_{u}+1\) rows of FACT contain the triangular matrix \(U\) in band storage form. The next \(\mathrm{m}_{l}\) rows of FACT contain the multipliers needed to produce \(L\).

The routine LFTRB is based on the the blocked \(L U\) factorization algorithm for banded linear systems given in Du Croz, et al. (1990). Level-3 BLAS invocations were replaced by in-line loops.

The blocking factor \(n b\) has the default value 1 in LFTRB. It can be reset to any positive value not exceeding 32 .

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2\) TRB/DL2TRB. The reference is:

CALL L2TRB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, WK)
The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N used for scaling.
2 Informational error
Type Code
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

21 The performance of the \(L U\) factorization may improve on high-performance computers if the blocking factor, NB, is increased. The current version of the routine allows NB to be reset to a value no larger than 32. Default value is \(\mathrm{NB}=1\).

\section*{Example}

A linear system with multiple right-hand sides is solved. LFTRB is called to factor the coefficient matrix. LFSRB is called to compute the two solutions for the two right-hand sides. In this case the coefficient matrix is assumed to be appropriately scaled. Otherwise, it may be better to call routine LFCRB to perform the factorization, and LFIRB to compute the solutions.
```

USE LFTRB INT
USE LFSRB - INT
USE WRRRN_INT
INTEGER LDA, LDFACT, N, NLCA, NUCA
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
REAL A(LDA,N), B(N,2), FACT(LDFACT,N), X(N,2)
Set values for A in band form, and B
A =( }$$
\begin{array}{lllll}{0.0}&{-1.0}&{-2.0}&{2.0}\end{array}
$$
( 2.0 1.0 -1.0 1.0)
($$
\begin{array}{llll}{-3.0}&{0.0}&{2.0}&{0.0}\end{array}
$$)
B = ( 12.0 -17.0)
(-19.0 23.0)
( 6.0 5.0)
( 8.0 5.0)
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0.\&

```
```

            2.0, 1.0, 0.0/
    DATA B/12.0, -19.0, 6.0, 8.0, -17.0, 23.0, 5.0, 5.0/
    CALL LFTRB (A, NLCA, NUCA, FACT, IPVT)
    Solve for the two right-hand sides
    DO 10 J=1, 2
        CALL LFSRB (FACT, NLCA, NUCA, IPVT, B (:,J), X(:,J))
    1 0 ~ C O N T I N U E ~
    ! Print results
CALL WRRRN ('X', X)
END

```

\section*{Output}
\begin{tabular}{rrr} 
& X & \\
& 1 & 2 \\
1 & 3.000 & -8.000 \\
2 & -6.000 & 1.000 \\
3 & 2.000 & 1.000 \\
4 & 4.000 & 3.000
\end{tabular}

\section*{LFSRB}

Solves a real system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{F A C T}-(2 * \mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the \(L U\) factorization of the coefficient matrix A as output from routine LFCRB/DLFCRB or LFTRB/DLFTRB. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCRB/DLFCRB or LFTRB/DLFTRB. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(N=\operatorname{size}(F A C T, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: ІРАТН \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSRB (FACT, NLCA, NUCA, IPVT, B, X [, ...])

Specific: The specific interface names are \(S_{\text {_LFSRB }}\) and \(D_{-}\)LFSRBB.

\section*{FORTRAN 77 Interface}

Single: CALL LFSRB (N, FACT, LDFACT, NLCA, NUCA, IPVT, B, IPATH, X)
Double: The double precision name is DLFSRB.

\section*{Description}

Routine LFSRB computes the solution of a system of linear algebraic equations having a real banded coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCRB or LFTRB. The solution to \(A x=b\) is found by solving the banded triangular systems \(L y=b\) and \(U x=y\). The forward elimination step consists of solving the system \(L y=b\) by applying the same permutations and elimination operations to \(b\) that were applied to the columns of \(A\) in the factorization routine. The backward substitution step consists of solving the banded triangular system \(U x=y\) for \(x\).

LFSRB and LFIRB both solve a linear system given its \(L U\) factorization. LFIRB generally takes more time and produces a more accurate answer than LFSRB. Each iteration of the iterative refinement algorithm used by LFIRB calls LFSRB.

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

\section*{Example}

The inverse is computed for a real banded \(4 \times 4\) matrix with one upper and one lower codiagonal. The input matrix is assumed to be well-conditioned, hence LFTRB is used rather than LFCRB.
```

USE LFSRB INT
USE LFTRBB INT
USE WRRRN_INT
INTEGER LDA, LDFACT, N, NLCA, NUCA
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)

```
```

    INTEGER IPVT(N)
    REAL A(LDA,N), AINV (N,N), FACT (LDFACT,N), RJ(N)
                                Set values for A in band form
                                A =( ( 0.0 -1.0 -2.0 2.0)
                                ( 2.0 1.0 -1.0 1.0)
                            (\begin{array}{llll}{-3.0}&{0.0}&{2.0}&{0.0}\end{array})
    DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0.&
        2.0, 1.0, 0.0/
    CALL LFTRB (A, NLCA, NUCA, FACT, IPVT)
                        Set up the columns of the identity
        matrix one at a time in RJ
    RJ = 0.0E0
    DO 10 J=1, N
        RJ(J) = 1.0E0
                            RJ is the J-th column of the identity
                        matrix so the following LFSRB
                        reference places the J-th column of
                        the inverse of A in the J-th column
                            of AINV
        CALL LFSRB (FACT, NLCA, NUCA, IPVT, RJ, AINV(:,J))
        RJ(J) = 0.0E0
    1 0 ~ C O N T I N U E ~
    CALL WRRRN ('AINV', AINV)
    END
    ```

\section*{Output}

AINV
\begin{tabular}{rrrrr} 
& \multicolumn{4}{c}{ AINV } \\
& 1 & 2 & 3 & 4 \\
1 & -1.000 & -1.000 & 0.400 & -0.800 \\
2 & -3.000 & -2.000 & 0.800 & -1.600 \\
3 & 0.000 & 0.000 & -0.200 & 0.400 \\
4 & 0.000 & 0.000 & 0.400 & 0.200
\end{tabular}

\section*{LFIRB}

Uses iterative refinement to improve the solution of a real system of linear equations in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{A}-(\mathrm{NUCA}+\mathrm{NLCA}+1)\) by N array containing the N by N banded coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-(2 * \mathrm{NLCA}+\mathrm{NUCA}+1)\) by N array containing the \(L U\) factorization of the matrix A as output from routines LFCRB/DLFCRB or LFTRB/DLFTRB. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine LFCRB/DLFCRB or LFTRB/DLFTRB. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)
\(\boldsymbol{R E S}\) - Vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{T} X=B\) is solved.
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIRB (A, NLCA, NUCA, FACT, IPVT, B, X, RES [,...])
Specific: The specific interface names are \(S_{\text {_LFIRB }}\) and \(D_{-}\)LFIRB.

\section*{FORTRAN 77 Interface}

Single: CALL LFIRB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, B, IPATH, X, RES)

Double: The double precision name is DLFIRB.

\section*{Description}

Routine LFIRB computes the solution of a system of linear algebraic equations having a real banded coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCRB or LFTRB.

Iterative refinement fails only if the matrix is very ill-conditioned.
LFIRB and LFSRB both solve a linear system given its \(L U\) factorization. LFIRB generally takes more time and produces a more accurate answer than LFSRB. Each iteration of the iterative refinement algorithm used by LFIRB calls LFSRB.

\section*{Comments}

Informational error
\begin{tabular}{cc} 
Type & Code \\
3 & 2
\end{tabular} \begin{tabular}{l} 
The input matrix is too ill-conditioned for iterative refinement to be \\
effective
\end{tabular}

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding 0.5 to the second element.
```

USE LFIRB INT
USE LFCRB INT
USE UMACH INT
USE WRRRN -INT
! Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
REAL A(LDA,N), B(N), FACT(LDFACT,N), RCOND, RES(N), X(N)
Set values for A in band form, and B
A =( }$$
\begin{array}{lllll}{0.0}&{-1.0}&{-2.0}&{2.0}\end{array}
$$
( 2.0 1.0 -1.0 1.0)
( -3.0 0.0 2.0 0.0)
B = ( 3.0 5.0 7.0 -9.0)
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0,\&
2.0, 1.0, 0.0/
DATA B/3.0, 5.0, 7.0, -9.0/
CALL LFCRB (A, NLCA, NUCA, FACT, IPVT, RCOND)
Print the reciprocal condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.OEO/RCOND
Solve the three systems

```
```

    DO 10 J=1, 3
        CALL LFIRB (A, NLCA, NUCA, FACT, IPVT, B, X, RES)
    ! Print results
CALL WRRRN ('X', X, 1, N, 1)
! Perturb B by adding 0.5 to B(2)
B(2) = B(2) + 0.5E0
1 0 ~ C O N T I N U E ~
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND < . 07
L1 Condition number \(=25.0\)
\begin{tabular}{rrrr}
1 & 2 & 3 & 4 \\
2.000 & 1.000 & -5.000 & 1.000
\end{tabular}
\begin{tabular}{rrrrr} 
& & X & & \\
1 & 2 & & 3 & 4 \\
1.500 & 0.000 & -5.000 & 1.000
\end{tabular}
\begin{tabular}{rrrr} 
& & X & \\
1 & 2 & 3 & 4 \\
1.000 & -1.000 & -5.000 & 1.000
\end{tabular}

\section*{LFDRB}

Computes the determinant of a real matrix in band storage mode given the \(L U\) factorization of the matrix.

\section*{Required Arguments}
\(\boldsymbol{F A C T}-(2 * N L C A+N U C A+1)\) by N array containing the \(L U\) factorization of the matrix A as output from routine LFTRB/DLFTRB or LFCRB/DLFCRB. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization as output from routine LFTRB/DLFTRB or LFCRB/DLFCRB. (Input)

DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET} 1 * 10^{\mathrm{DET} 2}\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDRB (FACT, NLCA, NUCA, IPVT, DET1, DET2 [,...])

Specific: The specific interface names are S_LFDRB and D_LFDRB.

\section*{FORTRAN 77 Interface}

Single: CALL LFDRB (N, FACT, LDFACT, NLCA, NUCA, IPVT, DET1, DET2)
Double: The double precision name is DLFDRB.

\section*{Description}

Routine LFDRB computes the determinant of a real banded coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LfCRB or Lftrb. The formula \(\operatorname{det} A=\operatorname{det} L \operatorname{det} U\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} U=\prod_{i=1}^{N} U_{i i}
\]
(The matrix \(U\) is stored in the upper NUCA + NLCA +1 rows of \(\operatorname{FACT}\) as a banded matrix.) Since \(L\) is the product of triangular matrices with unit diagonals and of permutation matrices, \(\operatorname{det} L=(-1)^{k}\), where \(k\) is the number of pivoting interchanges.

LFDRB is based on the LINPACK routine CGBDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a real banded \(4 \times 4\) matrix with one upper and one lower codiagonal.
```

USE LFDRB INT
USE LFTRB INT
USE UMACH_INT
N, NTCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
REAL A(LDA,N), DET1, DET2, FACT(LDFACT,N)
Set values for A in band form

```
```

!
!
!
DATA A/0.0, 2.0, -3.0, -1.0, 1.0, 0.0, -2.0, -1.0, 2.0.\&
2.0, 1.0, 0.0/
CALL LFTRB (A, NLCA, NUCA, FACT, IPVT)
Compute the determinant
CALL LFDRB (FACT, NLCA, NUCA, IPVT, DET1, DET2)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
99999 FORMAT (' The determinant of A is ', F6.3, ' * 10**', F2.0)
END

```

\section*{Output}

The determinant of \(A\) is \(5.000 * 10 * * 0\).

\section*{LSAQS}

Solves a real symmetric positive definite system of linear equations in band symmetric storage mode with iterative refinement.

\section*{Required Arguments}

A-NCODA +1 by N array containing the N by N positive definite band coefficient matrix in band symmetric storage mode. (Input)

NCODA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSAQS (A, NCODA, B, X [, ...])

Specific: The specific interface names are S_LSAQS and D_LSAQS.

\section*{FORTRAN 77 Interface}

Single: CALL LSAQS (N, A, LDA, NCODA, B, X)
Double: The double precision name is DLSAQS.

\section*{Description}

Routine LSAQS solves a system of linear algebraic equations having a real symmetric positive definite band coefficient matrix. It first uses the routine LFCQS to compute an \(R^{T} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. \(R\) is an upper triangular band matrix. The solution of the linear system is then found using the iterative refinement routine LFIQS.

LSAQS fails if any submatrix of \(R\) is not positive definite, if \(R\) has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSAQS solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2AQS/DL2AQS. The reference is:

CALL L2AQS (N, A, LDA, NCODA, B, X, FACT, WK)
The additional arguments are as follows:
FACT - Work vector of length NCODA +1 by n containing the \(R^{T} R\) factorization of A in band symmetric storage form on output.

WK - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is not positive definite.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2AQS the leading dimension of FACT is increased by
\(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and \(\operatorname{IVAL}(4)\) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSAQS. Additional memory allocation for FACT and option value restoration are done automatically in LSAQS.

Users directly calling L2AQS can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSAQS or L2AQS. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).
17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSAQS temporarily replaces IVAL(2) by IVAL(1). The routine L2CQS computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CQS skips this computation. LSAQS restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has real positive definite band form, and the right-hand-side vector \(b\) has four elements.
```

USE LSAQS INT
USE WRRRN_INT
INTEGER LDA, N, NCODA
PARAMETER (LDA=3, N=4, NCODA=2)
REAL A(LDA,N), B(N), X(N)
Set values for A in band symmetric form, and B
A =( (0.0 0.0 -1.0 1.0 )
( 0.0 0.0 2.0 -1.0 )
( 2.0 4.0 7.0 3.0)
B =( 6.0 -11.0 -11.0 19.0 )
DATA A/2*0.0, 2.0, 2*0.0, 4.0, -1.0, 2.0, 7.0, 1.0, -1.0, 3.0/
DATA B/6.0, -11.0, -11.0, 19.0/
Solve A*X = B
CALL LSAQS (A, NCODA, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```
!
\(!\)

\section*{Output}
\begin{tabular}{cccc} 
& \multicolumn{3}{c}{X} \\
1 & 2 & 3 & 4 \\
4.000 & -6.000 & 2.000 & 9.000
\end{tabular}

\section*{LSLQS}

Solves a real symmetric positive definite system of linear equations in band symmetric storage mode without iterative refinement.

\section*{Required Arguments}

A-NCODA +1 by N array containing the N by N positive definite band symmetric coefficient matrix in band symmetric storage mode. (Input)

NCODA - Number of upper codiagonals of A. (Input)
B - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N-\) Number of equations. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLQS (A, NCODA, B, X [, ...])
Specific: The specific interface names are S_LSLQS and D_LSLQS.

\section*{FORTRAN 77 Interface}

Single: CALL LSLQS (N, A, LDA, NCODA, B, X)
Double: The double precision name is DLSLQS.

\section*{Description}

Routine LSLQS solves a system of linear algebraic equations having a real symmetric positive definite band coefficient matrix. It first uses the routine LFCQS to compute an \(\mathrm{R}^{T} \mathrm{R}\) Cholesky
factorization of the coefficient matrix and to estimate the condition number of the matrix. \(R\) is an upper triangular band matrix. The solution of the linear system is then found using the routine LFSQS.

LSLQS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSAQS be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LQS/DL2LQS. The reference is:

CALL L2LQS (N, A, LDA, NCODA, B, X, FACT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}-\mathrm{NCODA}+1\) by N work array containing the \(R^{T} R\) factorization of A in band symmetric form on output. If A is not needed, A and FACT can share the same storage locations.

WK - Work vector of length N .
2. Informational errors

Type Code
31 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is not positive definite.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LQS the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by \(\operatorname{IVAL}(1)\) and IVAL(2), respectively, in LSLQS. Additional memory allocation for FACT and option value restoration are done automatically in LSLQS. Users directly calling L2LQS can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLQS or L2LQS. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLQS temporarily replaces IVAL(2) by IVAL(1). The routine L2CQS computes the condition number if IVAL \((2)=2\). Otherwise L2CQS
skips this computation. LSLQS restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has real positive definite band form and the right-hand-side vector \(b\) has four elements.
```

USE LSLQS_INT
USE WRRRN_INT
INTEGER LDA, N, NCODA
PARAMETER (LDA=3, N=4, NCODA=2)
REAL A(LDA,N), B(N), X(N)
Set values for A in band symmetric form, and B
A=($$
\begin{array}{lllll}{0.0}&{0.0}&{-1.0}&{1.0}\end{array}
$$)
( 0.0 0.0 2.0 -1.0 )
(2.0 4.0 7.0 3.0)
B=( 6.0 -11.0 -11.0 19.0 )
DATA A/2*0.0, 2.0, 2*0.0, 4.0, -1.0, 2.0, 7.0, 1.0, -1.0, 3.0/
DATA B/6.0, -11.0, -11.0, 19.0/
Solve A*X = B
CALL LSLQS (A, NCODA, B, X)

```
!
! Print results
CALL WRRRN (' \(X^{\prime}, \mathrm{X}, \mathrm{I}, \mathrm{N}, 1\) )
END

\section*{Output}
\begin{tabular}{cccc}
\multicolumn{4}{c}{X} \\
1 & 2 & 3 & 4 \\
4.000 & -6.000 & 2.000 & 9.000
\end{tabular}

\section*{LSLPB}

Computes the \(R^{T} D R\) Cholesky factorization of a real symmetric positive definite matrix \(A\) in codiagonal band symmetric storage mode. Solve a system \(A x=b\).

\section*{Required Arguments}

A-Array containing the N by N positive definite band coefficient matrix and right hand side in codiagonal band symmetric storage mode. (Input/Output)
The number of array columns must be at least NCODA +2 . The number of column is not an input to this subprogram.

On output, A contains the solution and factors. See Comments section for details.

NCODA - Number of upper codiagonals of matrix A. (Input)
Must satisfy NCODA \(\geq 0\) and NCODA \(<\mathrm{N}\).
\(\boldsymbol{U}\) - Array of flags that indicate any singularities of A, namely loss of positive-definiteness of a leading minor. (Output)
A value \(U(I)=0\). means that the leading minor of dimension \(I\) is not positive-definite.
Otherwise, \(\mathrm{U}(\mathrm{I})=1\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Must satisfy \(\mathrm{N}>0\).
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Must satisfy LDA \(\geq \mathrm{N}+\mathrm{NCODA}\).
Default: LDA \(=\operatorname{size}(A, 1)\).
IJOB - Flag to direct the desired factorization or solving step. (Input)
Default: \(\mathrm{IJOB}=1\).

\section*{IJOB Meaning}

1 factor the matrix \(A\) and solve the system \(A x=b\), where \(b\) is stored in column NCODA +2 of array A. The vector \(x\) overwrites \(b\) in storage.

2 solve step only. Use \(b\) as column NCODA +2 of A. (The factorization step has already been done.) The vector \(x\) overwrites \(b\) in storage.

3 factor the matrix A but do not solve a system.
\(4,5,6\) same meaning as with the value IJOB - 3 . For efficiency, no error checking is done on values LDA, N , NCODA, and \(\mathrm{U}(*)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLPB (A, NCODA, U [, ...])
Specific: The specific interface names are S_LSLPB and D_LSLPB.

\section*{FORTRAN 77 Interface}

Single: CALL LSLPB (N, A, LDA, NCODA, IJOB, U)
Double: The double precision name is DLSLPB.

\section*{Description}

Routine LSLPB factors and solves the symmetric positive definite banded linear system \(A x=b\). The matrix is factored so that \(A=R^{T} D R\), where \(R\) is unit upper triangular and \(D\) is diagonal. The reciprocals of the diagonal entries of \(D\) are computed and saved to make the solving step more efficient. Errors will occur if \(D\) has a non-positive diagonal element. Such events occur only if \(A\) is very close to a singular matrix or is not positive definite.

LSLPB is efficient for problems with a small band width. The particular cases NCODA \(=0,1,2\) are done with special loops within the code. These cases will give good performance. See Hanson (1989) for details. When solving tridiagonal systems, NCODA \(=1\), the cyclic reduction code LSLCR should be considered as an alternative. The expectation is that LSLCR will outperform LSLPB on vector or parallel computers. It may be inferior on scalar computers or even parallel computers with non-optimizing compilers.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LPB/DL2LPB. The reference is:

CALL L2LPB (N, A, LDA, NCODA, IJOB, U, WK)
The additional argument is:
WK - Work vector of length NCODA.
2. If \(I J O B=1,3,4\), or 6 , A contains the factors \(R\) and \(D\) on output. These are stored in codiagonal band symmetric storage mode. Column 1 of A contains the reciprocal of diagonal matrix \(D\). Columns 2 through NCODA +1 contain the upper diagonal values for upper unit diagonal matrix \(R\). If \(\operatorname{IJOB}=1,2,4\), or 5 , the last column of A contains the solution on output, replacing b.
3. Informational error

Type Code
42 The input matrix is not positive definite.

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has real positive definite codiagonal band form and the right-hand-side vector \(b\) has four elements.
```

USE LSLPB_INT
USE WRRRN_INT
INTEGER LDA, N, NCODA
PARAMETER ( }\textrm{N}=4, NCODA=2, LDA=N+NCODA
INTEGER IJOB
REAL A(LDA,NCODA+2), U(N)
REAL R(N,N), RT(N,N), D(N,N), WK(N,N), AA (N,N)

```
\(!\)
!

```

DATA ((A (I+NCODA,J),I=1,N),J=1,NCODA+2)/2.0, 4.0, 7.0, 3.0, 0.0,\&
0.0, 2.0, -1.0, 0.0, 0.0, -1.0, 1.0, 6.0, -11.0, -11.0,\&
19.0/
DATA R/16*0.0/, D/16*0.0/, RT/16*0.0/
Factor and solve A*x = b.
CALL LSLPB(A, NCODA, U)
CALL WRRRN ('X', A((NCODA+1):,(NCODA+2):), NRA=1, NCA=N, LDA=1)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|}
\hline 1 & 2 & 3 & 4 \\
\hline 4.000 & -6.000 & 2.000 & 9.000 \\
\hline
\end{tabular}

\section*{LFCQS}

Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{NCODA}+1\) by N array containing the N by N positive definite band coefficient matrix in band symmetric storage mode to be factored. (Input)

NCODA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-\mathrm{NCODA}+1\) by N array containing the \(R^{T} R\) factorization of the matrix A in band symmetric form. (Output) If \(A\) is not needed, A and FACT can share the same storage locations.
\(\boldsymbol{R C O N D}\) - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\mathrm{FACT}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCQS (A, NCODA, FACT, RCOND [, ...])

Specific: The specific interface names are S_LFCQS and D_LFCQS.

\section*{FORTRAN 77 Interface}

Single: CALL LFCQS (N, A, LDA, NCODA, FACT, LDFACT, RCOND)
Double: The double precision name is DLFCQS.

\section*{Description}

Routine LFCQS computes an \(R^{T} R\) Cholesky factorization and estimates the condition number of a real symmetric positive definite band coefficient matrix. \(R\) is an upper triangular band matrix.

The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCQS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

The \(R^{T} R\) factors are returned in a form that is compatible with routines LFIQS, LFSQS and LFDQS. To solve systems of equations with multiple right-hand-side vectors, use LFCQS followed by either LFIQS or LFSQS called once for each right-hand side. The routine LFDQS can be called to compute the determinant of the coefficient matrix after LFCQS has performed the factorization.
LFCQS is based on the LINPACK routine SPBCO; see Dongarra et al. (1979).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CQS/DL2CQS. The reference is:
```

CALL L2CQS (N, A, LDA, NCODA, FACT, LDFACT, RCOND, WK)

```

The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Work vector of length N .
2. Informational errors

Type Code
33 The input matrix is algorithmically singular.
42 The input matrix is not positive definite.

\section*{Example}

The inverse of a \(4 \times 4\) symmetric positive definite band matrix with one codiagonal is computed. LFCQS is called to factor the matrix and to check for nonpositive definiteness or ill-conditioning. LFIQS is called to determine the columns of the inverse.
```

    USE LFCQS INT
    USE LFIQS INT
    USE UMACH_INT
    USE WRRRN INT
    INTEGER LDA, LDFACT, N, NCODA, NOUT
    PARAMETER (LDA=2, LDFACT=2, N=4, NCODA=1)
    REAL A(LDA,N), AINV (N,N), RCOND, FACT (LDFACT,N), &
        RES(N), RJ(N)
            Set values for A in band symmetric form
            A =( lllll
            (2.0 2.5 2.5 2.0 )
    DATA A/0.0, 2.0, 1.0, 2.5, 1.0, 2.5, 1.0, 2.0/
        Factor the matrix A
    CALL LFCQS (A, NCODA, FACT, RCOND)
                        Set up the columns of the identity
                        matrix one at a time in RJ
    RJ = 0.0E0
    DO 10 J=1, N
        RJ(J) = 1.0E0
    RJ is the J-th column of the identity
    matrix so the following LFIQS
    reference places the J-th column of
    the inverse of A in the J-th column
    of AINV
    CALL LFIQS (A, NCODA, FACT, RJ, AINV(:,J), RES)
    RJ(J) = 0.0E0
    10 CONTINUE

```
```

! Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
CALL WRRRN ('AINV', AINV)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}
```

RCOND = 0.160
L1 Condition number = 6.239

|  | AINV |  |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 |  | 3 |
| 4 |  |  |  |  |
| 1 | 0.6667 | -0.3333 | 0.1667 | -0.0833 |
| 2 | -0.3333 | 0.6667 | -0.3333 | 0.1667 |
| 3 | 0.1667 | -0.3333 | 0.6667 | -0.3333 |
| 4 | -0.0833 | 0.1667 | -0.3333 | 0.6667 |

```

\section*{LFTQS}

Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode.

\section*{Required Arguments}
\(\boldsymbol{A}-\mathrm{NCODA}+1\) by N array containing the N by N positive definite band coefficient matrix in band symmetric storage mode to be factored. (Input)

NCODA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-\mathrm{NCODA}+1\) by N array containing the \(R^{T} R\) factorization of the matrix A. (Output) If A s not needed, A and FACT can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTQS (A, NCODA, FACT [, ...])

Specific: The specific interface names are S_LFTQS and D_LFTQS.

\section*{FORTRAN 77 Interface}

Single: CALL LFTQS (N, A, LDA, NCODA, FACT, LDFACT)
Double: The double precision name is DLFTQS.

\section*{Description}

Routine LFTQS computes an \(R^{T} R\) Cholesky factorization of a real symmetric positive definite band coefficient matrix. \(R\) is an upper triangular band matrix.

LFTQS fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) is very close to a singular matrix or to a matrix which is not positive definite.

The \(R^{T} R\) factors are returned in a form that is compatible with routines LFIQS, LFSQS and LFDQS. To solve systems of equations with multiple right hand-side vectors, use LFTQS followed by either LFIQS or LFSQS called once for each right-hand side. The routine LFDQS can be called to compute the determinant of the coefficient matrix after LFTQS has performed the factorization.

LFTQS is based on the LINPACK routine CPBFA; see Dongarra et al. (1979).

\section*{Comments}

Informational error
\(\begin{array}{ccc}\text { Type } & \text { Code } & \\ 4 & 2 & \text { The input matrix is not positive definite. }\end{array}\)

\section*{Example}

The inverse of a \(3 \times 3\) matrix is computed. LFTQS is called to factor the matrix and to check for nonpositive definiteness. LFSQS is called to determine the columns of the inverse.
```

USE LFTQS INT
USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
USE LFSQS_INT
INTEGER LDA, LDFACT, N, NCODA
PARAMETER (LDA=2, LDFACT=2, N=4, NCODA=1)
REAL A(LDA,N), AINV (N,N), FACT(LDFACT,N), RJ (N)
Set values for A in band symmetric form
A =( lllll}$$
\begin{array}{llll}{0.0}&{1.0}&{1.0}&{1.0}\end{array}
$$
(2.0 2.5 2.5 2.0 )
DATA A/0.0, 2.0, 1.0, 2.5, 1.0, 2.5, 1.0, 2.0/
Factor the matrix A
CALL LFTQS (A, NCODA, FACT)

```
```

Set up the columns of the identity

```
matrix one at a time in RJ
    \(R J=0.0 \mathrm{E} 0\)
    DO \(10 \mathrm{~J}=1\), N
        RJ(J) = 1.0E0
    RJ is the J-th column of the identity
    matrix so the following LFSQS
    reference places the J-th column of
    the inverse of \(A\) in the \(J\)-th column
    of AINV
    CALL LFSQS (FACT, NCODA, RJ, AINV(:, J))
        RJ (J) \(=0.0 \mathrm{E} 0\)
    10 CONTINUE
! Print the results
    CALL WRRRN ('AINV', AINV, ITRING=1)
    END

\section*{Output}
\begin{tabular}{rrrr}
\multicolumn{4}{c}{ AINV } \\
1 & 2 & 3 & 4 \\
0.6667 & -0.3333 & 0.1667 & -0.0833 \\
& 0.6667 & -0.3333 & 0.1667 \\
& & 0.6667 & -0.3333 \\
& & & 0.6667
\end{tabular}

\section*{LFSQS}

Solves a real symmetric positive definite system of linear equations given the factorization of the coefficient matrix in band symmetric storage mode.

\section*{Required Arguments}

FACT — NCODA +1 by N array containing the \(R^{T} R\) factorization of the positive definite band matrix A in band symmetric storage mode as output from subroutine LFCQS/DLFCQS or LFTQS/DLFTQS. (Input)

NCODA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) an share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSQS (FACT, NCODA, B, X [, ...])
Specific: The specific interface names are S_LFSQS and D_LFSQS.

\section*{FORTRAN 77 Interface}

Single: CALL LFSQS (N, FACT, LDFACT, NCODA, B, X)
Double: The double precision name is DLFSQS.

\section*{Description}

This routine computes the solution for a system of linear algebraic equations having a real symmetric positive definite band coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either LFCQS or LFTQS. \(R\) is an upper triangular band matrix.
The solution to \(A x=b\) is found by solving the triangular systems \(R^{T} y=b\) and \(R x=y\).
LFSQS and LFIQS both solve a linear system given its \(R^{T} R\) factorization. LFIQS generally takes more time and produces a more accurate answer than LFSQS. Each iteration of the iterative refinement algorithm used by LFIQS calls LFSQS.

LFSQS is based on the LINPACK routine SPBSL; see Dongarra et al. (1979).

\section*{Comments}

Informational error
Type Code
41 The factored matrix is singular.

\section*{Example}

A set of linear systems is solved successively. LFTQS is called to factor the coefficient matrix. LFSQS is called to compute the four solutions for the four right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCQS to perform the factorization, and LFIQS to compute the solutions.
```

USE LFSQS_INT
USE LFTQS_INT
USE WRRRN_INT
! Declare variables
INTEGER LDA, LDFACT, N, NCODA
PARAMETER (LDA=3, LDFACT=3, N=4, NCODA=2)

```
```

    REAL A(LDA,N), B(N,4), FACT (LDFACT,N), X(N,4)
        Set values for A in band symmetric form, and B
    A=( lllll
        ( 0.0 0.0 2.0 2.0 -1.0 )
    B = ( 4.0 -3.0 9.0 -1.0 )
        ( 6.0 10.0 29.0 3.0 )
        ( 15.0 12.0 11.0 6.0 )
        (\begin{array}{clll}{-7.0}&{1.0}&{14.0}&{2.0}\end{array})
    DATA A/2*0.0, 2.0, 2*0.0, 4.0, -1.0, 2.0, 7.0, 1.0, -1.0, 3.0/
    DATA B/4.0, 6.0, 15.0, -7.0, -3.0, 10.0, 12.0, 1.0, 9.0, 29.0,&
        11.0, 14.0, -1.0, 3.0, 6.0, 2.0/
                            Factor the matrix A
    CALL LFTQS (A, NCODA, FACT)
    DO 10 I=1, 4
        CALL LFSQS (FACT, NCODA, B(:,I), X(:,I))
    1 0
    CONTINUE
    END
    ```

\section*{Output}
\begin{tabular}{rrrrr} 
& & X & & \\
& 1 & 2 & 3 & 4 \\
1 & 3.000 & -1.000 & 5.000 & 0.000 \\
2 & 1.000 & 2.000 & 6.000 & 0.000 \\
3 & 2.000 & 1.000 & 1.000 & 1.000 \\
4 & -2.000 & 0.000 & 3.000 & 1.000
\end{tabular}

\section*{LFIQS}

Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations in band symmetric storage mode.

\section*{Required Arguments}

A- NCODA +1 by n array containing the n by n positive definite band coefficient matrix in band symmetric storage mode. (Input)

NCODA - Number of upper codiagonals of A. (Input)
FACT — NCODA +1 by N array containing the \(R^{T} R\) factorization of the matrix A from routine LFCQS/DLFCQS or LFTQS/DLFTQS. (Input)
\(\boldsymbol{B}-\) Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}-\) Vector of length N containing the solution to the system. (Output)
\(\boldsymbol{R E S}\) - Vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIQS (A, NCODA, FACT, B, X, RES [, ...])
Specific: The specific interface names are S_LFIQS and D_LFIQS.

\section*{FORTRAN 77 Interface}

Single: CALL LFIQS (N, A, LDA, NCODA, FACT, LDFACT, B, X, RES)
Double: The double precision name is DLFIQS.

\section*{Description}

Routine LFIQS computes the solution of a system of linear algebraic equations having a real symmetric positive-definite band coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.
To compute the solution, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either IMSL routine LFCQS or LFTQS.

Iterative refinement fails only if the matrix is very ill-conditioned.
LFIQS and LFSQS both solve a linear system given its \(R^{T} R\) factorization. LFIQS generally takes more time and produces a more accurate answer than LFSQS. Each iteration of the iterative refinement algorithm used by LFIQS calls LFSQS.

\section*{Comments}

Informational error
Type Code
34 The input matrix is too ill-conditioned for iterative refinement to be effective.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding 0.5 to the second element.
```

USE LFIQS_INT
USE UMACH INT
USE LFCQS_INT
USE WRRRN INT
INTEGER LDA, LDFACT, N, NCODA, NOUT
PARAMETER (LDA=2, LDFACT=2, N=4, NCODA=1)
REAL A (LDA,N), B(N), RCOND, FACT (LDFACT,N), RES (N, 3),\&
X(N, 3)
Set values for A in band symmetric form, and B
A =($$
\begin{array}{llll}{0.0}&{1.0}&{1.0}&{1.0}\end{array}
$$)
B = ( 3.0 5.0 7.0 4.0 )
DATA A/0.0, 2.0, 1.0, 2.5, 1.0, 2.5, 1.0, 2.0/
DATA B/3.0, 5.0, 7.0, 4.0/
Factor the matrix A
CALL LFCQS (A, NCODA, FACT, RCOND)
Print the estimated condition number
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) RCOND, 1.0EO/RCOND
DO 10 I=1, 3
CALL LFIQS (A, NCODA, FACT, B, X(:,I), RES (:,I))
B(2) = B(2) + 0.5E0
1 0 ~ C O N T I N U E ~
CALL WRRRN ('X', X)
CALL WRRRN ('RES', RES)
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```
\(!\)
!

\section*{Output}
```

RCOND = 0.160
L1 Condition number = 6.239
X
rrrr

```
```

2 0.667 1.000 1.333
3 2.167 2.000 1.833
40.917 1.000 1.083

|  | RES |  |  |
| ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 |
| 1 | $7.947 \mathrm{E}-08$ | $0.000 \mathrm{E}+00$ | $9.934 \mathrm{E}-08$ |
| 2 | $7.947 \mathrm{E}-08$ | $0.000 \mathrm{E}+00$ | $3.974 \mathrm{E}-08$ |
| 3 | $7.947 \mathrm{E}-08$ | $0.000 \mathrm{E}+00$ | $1.589 \mathrm{E}-07$ |
| 4 | $-3.974 \mathrm{E}-08$ | $0.000 \mathrm{E}+00$ | $-7.947 \mathrm{E}-08$ |

```

\section*{LFDQS}

Computes the determinant of a real symmetric positive definite matrix given the \(R^{T} R\) Cholesky factorization of the band symmetric storage mode.

\section*{Required Arguments}
\(\boldsymbol{F A C T}-\mathrm{NCODA}+1\) by N array containing the \(R^{T} R\) factorization of the positive definite band matrix, A, in band symmetric storage mode as output from subroutine LFCQS/DLFCQS or LFTQS/DLFTQS. (Input)

NCODA - Number of upper codiagonals of A. (Input)
DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\mathrm{DET}}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT — Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDQS (FACT, NCODA, DET1, DET2 [,...])
Specific: The specific interface names are \(S_{\text {_LFDQ }}\) and \(D_{-}\)LFDQS.

\section*{FORTRAN 77 Interface}

Single: CALL LFDQS (N, FACT, LDFACT, NCODA, DET1, DET2)

Double: \(\quad\) The double precision name is DLFDQS.

\section*{Description}

Routine LFDQS computes the determinant of a real symmetric positive-definite band coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(R^{T} R\) factorization. This may be done by calling either IMSL routine LFCQS or LFTQS. The formula \(\operatorname{det} A=\operatorname{det} R^{T} \operatorname{det} R=\left(\operatorname{det} R^{22}\right.\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} R=\prod_{i=1}^{v} R_{i}
\]

LFDQS is based on the LINPACK routine SPBDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a real positive definite \(4 \times 4\) matrix with 2 codiagonals.
```

    USE LFDQS_INT
    USE LFTQS_INT
USE UMACH_
! - Declare variables
INTEGER LDA, LDFACT, N, NCODA, NOUT
PARAMETER (LDA=3, N=4, LDFACT=3, NCODA=2)
REAL A(LDA,N), DET1, DET2, FACT(LDFACT,N)
Set values for A in band symmetric form
A=( lllll
( 0.0 2.0 1.0 3.0)
( 7.0 6.0 6.0 8.0)
DATA A/2*0.0, 7.0, 0.0, 2.0, 6.0, 1.0, 1.0, 6.0, -2.0, 3.0, 8.0/
CALL LFTQS (A, NCODA, FACT)
Compute the determinant
CALL LFDQS (FACT, NCODA, DET1, DET2)
Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is ',F6.3,' * 10**', F2.0)
END

```

\section*{Output}

The determinant of \(A\) is \(1.186 * 10 * * 3\).

\section*{LSLTQ}

Solves a complex tridiagonal system of linear equations.

\section*{Required Arguments}
\(\boldsymbol{C}\) - Complex vector of length N containing the subdiagonal of the tridiagonal matrix in \(\mathrm{C}(2)\)
through \(\mathrm{C}(\mathrm{N})\). (Input/Output)
On output C is destroyed.
\(\boldsymbol{D}\) - Complex vector of length N containing the diagonal of the tridiagonal matrix.
(Input/Output)
On output \(D\) is destroyed.
\(\boldsymbol{E}\) - Complex vector of length N containing the superdiagonal of the tridiagonal matrix in \(\mathrm{E}(1)\) through \(\mathrm{E}(\mathrm{N}-1)\). (Input/Output)
On output E is destroyed.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system on entry and the solution vector on return. (Input/Output)

\section*{Optional Arguments}
\(N\) - Order of the tridiagonal matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{C}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLTQ (C, D, E, B [,...])

Specific: The specific interface names are S_LSLTQ and D_LSLTQ.

\section*{FORTRAN 77 Interface}

Single: CALL LSLTQ (N, C, D, E, B)
Double: The double precision name is DLSLTQ.

\section*{Description}

Routine LSLTQ factors and solves the complex tridiagonal linear system \(A x=b\). LSLTQ is intended just for tridiagonal systems. The coefficient matrix does not have to be symmetric. The algorithm is Gaussian elimination with pivoting for numerical stability. See Dongarra et al. (1979), LINPACK subprograms CGTSL/ZGTSL, for details. When computing on vector or parallel computers the cyclic reduction algorithm, LSLCQ, should be considered as an alternative method to solve the system.

\section*{Comments}

Informational error
\(\begin{array}{ccc}\text { Type } & \text { Code } \\ 4 & 2 & \text { An element along the diagonal became exactly zero during execution. }\end{array}\)

\section*{Example}

A system of \(n=4\) linear equations is solved.
```

USE LSLTQ_INT
USE WRCRL_INT
INTEGER N
PARAMETER (N=4)
COMPLEX B(N), C(N), D(N), E(N)
CHARACTER CLABEL(1)*6, FMT*8, RLABEL(1)*4
DATA FMT/'(E13.6)'/
DATA CLABEL/'NUMBER'/
DATA RLABEL/'NONE'/
C(*), D(*), E(*) and B(*)
contain the subdiagonal,
diagonal, superdiagonal and
right hand side.
DATA C/(0.0,0.0), (-9.0,3.0), (2.0,7.0), (7.0,-4.0)/
DATA D/ (3.0,-5.0), (4.0,-9.0), (-5.0,-7.0), (-2.0,-3.0)/
DATA E/ (-9.0,8.0), (1.0,8.0), (8.0,3.0), (0.0,0.0)/
DATA B/(-16.0,-93.0), (128.0,179.0), (-60.0,-12.0), (9.0,-108.0)/
CALL LSLTQ (C, D, E, B)
CALL WRCRL ('Solution:', B, RLABEL, CLABEL, 1, N, 1, FMT=FMT)
END

```
!
\(!\)

\section*{Output}

Solution:
\begin{tabular}{ccc}
\((-0.400000 \mathrm{E}+01,-0.700000 \mathrm{E}+01)\) & \((-0.700000 \mathrm{E}+01,0.400000 \mathrm{E}+01)\) \\
3 & 4 \\
\((0.700000 \mathrm{E}+01,-0.700000 \mathrm{E}+01)\) & \((0.900000 \mathrm{E}+01,0.200000 \mathrm{E}+01)\)
\end{tabular}

\section*{LSLCQ}

Computes the \(L D U\) factorization of a complex tridiagonal matrix \(A\) using a cyclic reduction algorithm.

\section*{Required Arguments}
\(\boldsymbol{C}\) - Complex array of size 2 N containing the upper codiagonal of the N by N tridiagonal matrix in the entries \(\mathrm{C}(1), \ldots, \mathrm{C}(\mathrm{N}-1)\). (Input/Output)
\(\boldsymbol{A}\) - Complex array of size 2 N containing the diagonal of the N by N tridiagonal matrix in the entries \(A(1), \ldots, A(N)\). (Input/Output)
\(\boldsymbol{B}\) - Complex array of size 2 N containing the lower codiagonal of the N by N tridiagonal matrix in the entries \(B(1), \ldots, B(N-1)\). (Input/Output)
\(\boldsymbol{Y}\) - Complex array of size 2 N containing the right-hand side of the system \(A x=y\) in the order \(\mathrm{Y}(1), \ldots, \mathrm{Y}(\mathrm{N})\). (Input/Output)
The vector \(x\) overwrites \(Y\) in storage.
\(\boldsymbol{U}\) - Real array of size 2 N of flags that indicate any singularities of A. (Output) A value \(U(I)=1\). means that a divide by zero would have occurred during the factoring. Otherwise \(U(I)=0\).
\(\boldsymbol{I R}\) - Array of integers that determine the sizes of loops performed in the cyclic reduction algorithm. (Output)

IS - Array of integers that determine the sizes of loops performed in the cyclic reduction algorithm. (Output)
The sizes of these arrays must be at least \(\log _{2}(\mathrm{~N})+3\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input) N must be greater than zero. Default: \(\mathrm{N}=\operatorname{size}(\mathrm{C}, 1)\).

IJOB - Flag to direct the desired factoring or solving step. (Input) Default: I JOB \(=1\).

\section*{IJOB Action}

1 Factor the matrix \(A\) and solve the system \(A x=y\), where \(y\) is stored in array \(Y\).
2 Do the solve step only. Use \(y\) from array \(Y\). (The factoring step has already been done.)
3 Factor the matrix \(A\) but do not solve a system.
4 Same meaning as with the value \(I J O B=3\). For efficiency, no error checking is done on the validity of any input value.

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCQ (C, A, B, Y, U, IR, IS [, ...])

Specific: The specific interface names are S_LSLCQ and D_LSLCQ.

\section*{FORTRAN 77 Interface}

Single: CALL LSLCQ (N, C, A, B, IJOB, Y, U, IR, IS )
Double: \(\quad\) The double precision name is DLSLCQ.

\section*{Description}

Routine LSLCQ factors and solves the complex tridiagonal linear system \(A x=y\). The matrix is decomposed in the form \(A=L D U\), where \(L\) is unit lower triangular, \(U\) is unit upper triangular, and \(D\) is diagonal. The algorithm used for the factorization is effectively that described in Kershaw (1982). More details, tests and experiments are reported in Hanson (1990).

LSLCQ is intended just for tridiagonal systems. The coefficient matrix does not have to be Hermitian. The algorithm amounts to Gaussian elimination, with no pivoting for numerical stability, on the matrix whose rows and columns are permuted to a new order. See Hanson (1990) for details. The expectation is that LSLCQ will outperform either LSLTQ or LSLQB on vector or parallel computers. Its performance may be inferior for small values of \(n\), on scalar computers, or high-performance computers with non-optimizing compilers.

\section*{Example}

A real skew-symmetric tridiagonal matrix, \(A\), of dimension \(n=1000\) is given by \(\mathrm{c}_{k}=-k, a_{k}=0\), and \(b_{k}=k, k=1, \ldots, n-1, a_{n}=0\). This matrix will have eigenvalues that are purely imaginary. The eigenvalue closest to the imaginary unit is required. This number is obtained by using inverse iteration to approximate a complex eigenvector \(y\). The eigenvalue is approximated by \(\lambda=y^{H} \mathrm{Ay} / y^{H} y\). (This example is contrived in the sense that the given tridiagonal skew-symmetric matrix eigenvalue problem is essentially equivalent to the tridiagonal symmetic eigenvalue problem where the \(c_{k}=k\) and the other data are unchanged.)
```

USE LSLCQ_INT
USE UMACH_INT
! Declare variables
INTEGER LP, N, N2
PARAMETER (LP=12, N=1000, N2=2*N)
INTEGER I, IJOB, IR(LP), IS(LP), K, NOUT
REAL AIMAG, U(N2)
COMPLEX A(N2), B(N2), C(N2), CMPLX, CONJG, S, T, Y(N2)
INTRINSIC AIMAG, CMPLX, CONJG
Define entries of skew-symmetric
matrix, A:
DO 10 I=1, N - 1
C(I) = -I
This amounts to subtracting the
positive imaginary unit from the
diagonal. (The eigenvalue closest
to this value is desired.)
A(I) = CMPLX(0.E0,-1.0E0)

```
!
```

        B(I) = I
    ! This initializes the approximate
! Y(I) = 1.E0
10 CONTINUE
A(N) = CMPLX(0.E0,-1.0E0)
Y(N) = 1.EO
First step of inverse iteration
CALL LSLCQ (C, A, B, Y, U, IR, IS, N=N, IJOB=IJOB)
Next steps of inverse iteration
follow. Solve the system again with
the decomposition ready:
IJOB = 2
DO 20 K=1, 3
CALL LSLCQ (C, A, B, Y, U, IR, IS, N=N, IJOB=IJOB)
20 CONTINUE
!
Compute the Raleigh quotient to
estimate the eigenvalue closest to
the positive imaginary unit. After
the approximate eigenvector, y, is
computed, the estimate of the
eigenvalue is ctrans(y)*A*y/t,
where t = ctrans(y)*y.
S = -CONJG(Y(1))*Y(2)
T = CONJG(Y(1))*Y(1)
DO 30 I=2, N - 1
S = S + CONJG(Y(I))*((I-1)*Y(I-1)-I*Y(I+1))
T = T + CONJG(Y(I))*Y(I)
30 CONTINUE
S = S + CONJG (Y (N))* (N-1)*Y (N-1)
T = T + CONJG(Y(N))*Y(N)
S = S/T
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The value of n is: ', N
WRITE (NOUT,*) ' Value of approximate imaginary eigenvalue:',\&
AIMAG (S)
STOP
END

```

\section*{Output}

The value of \(n\) is: 1000
Value of approximate imaginary eigenvalue: 1.03811

\section*{LSACB}

Solves a complex system of linear equations in band storage mode with iterative refinement.

\section*{Required Arguments}

A-Complex NLCA + NUCA +1 by N array containing the N by N banded coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved.
Default: \(\operatorname{IPATH}=1\)

\section*{FORTRAN 90 Interface}

Generic: CALL LSACB (A, NLCA, NUCA, B, X [,...])

Specific: The specific interface names are S_LSACB and D_LSACB.

\section*{FORTRAN 77 Interface}

Single: CALL LSACB (N, A, LDA, NLCA, NUCA, B, IPATH, X)

Double: The double precision name is DLSACB.

\section*{Description}

Routine LSACB solves a system of linear algebraic equations having a complex banded coefficient matrix. It first uses the routine LFCCB to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using the iterative refinement routine LFICB.

LSACB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element or if the iterative refinement algorithm fails to converge. These errors occur only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSACB solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 A C B / D L 2 A C B\). The reference is:

CALL L2ACB (N, A, LDA, NLCA, NUCA, B, IPATH, X, FACT, IPVT, WK)
The additional arguments are as follows:
\(\boldsymbol{F A C T}\) - Complex work vector of length \((2 * \mathrm{NLCA}+\mathrm{NUCA}+1) * \mathrm{~N}\) containing the \(L U\) factorization of A on output.

IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.
\(\boldsymbol{W} \boldsymbol{K}\) - Complex work vector of length N .
2. Informational errors

Type Code
33 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2ACB the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and \(\operatorname{IVAL}(4)\) are temporarily replaced by \(\operatorname{IVAL}(1)\) and IVAL(2), respectively, in LSACB. Additional memory allocation for FACT and option value restoration are done automatically in LSACB. Users directly calling L2ACB can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSACB or L2ACB. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSACB temporarily replaces IVAL(2) by IVAL(1). The routine L2CCB computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise

L2CCB skips this computation. LSACB restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has complex banded form with one upper and one lower codiagonal. The right-hand-side vector \(b\) has four elements.
```

USE LSACB_INT
USE WRCRN_INT
INTEGER Declare variables
PARAMETER (LDA=3, N=4, NLCA=1, NUCA=1)
COMPLEX A(LDA,N), B(N), X(N)
Set values for A in band form, and B
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( -2.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 1.0+1.0i 0.0+2.0i 0.0+0.0i )
B = ( -10.0-5.0i 9.5+5.5i 12.0-12.0i 0.0+8.0i )
DATA A/ (0.0,0.0), (-2.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0),\&
(1.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),\&
(1.0,-1.0), (0.0,0.0)/
DATA B/(-10.0,-5.0), (9.5.5.5), (12.0,-12.0), (0.0,8.0)/
Solve A*X = B
CALL LSACB (A, NLCA, NUCA, B, X)
Print results
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}

X
\((3.000,0.000)(-1.000,1.000) \quad(3.000,0.000) \quad(-1.000,1.000)\)

\section*{LSLCB}

Solves a complex system of linear equations in band storage mode without iterative refinement.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex NLCA + NUCA +1 by N array containing the N by N banded coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, then \(B\) and \(X\) may share the same storage locations)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCB (A, NLCA, NUCA, B, X [, ...])

Specific: The specific interface names are \(S_{\text {_LSLCB }}\) and \(D_{-}\)LSLCB.

\section*{FORTRAN 77 Interface}

Single: CALL LSLCB (N, A, LDA, NLCA, NUCA, B, IPATH, X)
Double: The double precision name is DLSLCB.

\section*{Description}

Routine LSLCB solves a system of linear algebraic equations having a complex banded coefficient matrix. It first uses the routine LFCCB to compute an \(L U\) factorization of the coefficient matrix and to estimate the condition number of the matrix. The solution of the linear system is then found using LFSCB.

LSLCB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This occurs only if \(A\) is singular or very close to a singular matrix.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly scaled, it is recommended that LSACB be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LCB/DL2 LCB The reference is:
```

CALL L2LCB (N, A, LDA, NLCA, NUCA, B, IPATH, X, FACT, IPVT, WK)

```

The additional arguments are as follows:
\(\boldsymbol{F A C T}-(2 *\) NLCA \(+\mathrm{NUCA}+1) \times \mathrm{N}\) complex work array containing the \(L U\) factorization of \(A\) on output. If \(A\) is not needed, A can share the first \((N L C A+N U C A+1) * N\) locations with FACT.

IPVT - Integer work vector of length N containing the pivoting information for the \(L U\) factorization of A on output.

WK - Complex work vector of length \(N\).
2. Informational errors

Type Code
33 The input matrix is too ill-conditioned. The solution might not be accurate.
42 The input matrix is singular
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2LCB the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSLCB. Additional memory allocation for FACT and option value restoration are done automatically in LSLCB. Users directly calling L2LCB can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLCB or L2LCB. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLCB temporarily replaces IVAL(2) by IVAL(1). The routine L2CCB computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L2CCB skips this computation. LSLCB restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of four linear equations is solved. The coefficient matrix has complex banded form with one upper and one lower codiagonal. The right-hand-side vector \(b\) has four elements.
```

USE LSLCB_INT
USE WRCRN_-INT

```
```

! Declare variables
INTEGER LDA, N, NLCA, NUCA
PARAMETER (LDA=3, N=4, NLCA=1, NUCA=1)
COMPLEX A(LDA,N), B(N), X(N)
Set values for A in band form, and B
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( -2.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 1.0+1.0i 0.0+2.0i 0.0+0.0i )
B = ( -10.0-5.0i 9.5+5.5i 12.0-12.0i 0.0+8.0i )
DATA A/ (0.0,0.0), (-2.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0),\&
(1.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),\&
(1.0,-1.0), (0.0,0.0)/
DATA B/(-10.0,-5.0), (9.5,5.5), (12.0,-12.0), (0.0.8.0)/
CALL LSLCB (A, NLCA, NUCA, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}

X
\(\left(\begin{array}{cccc}1 & { }^{2} & \left.\begin{array}{c}3 \\ 4\end{array}\right)\end{array}\right.\)

\section*{LFCCB}

Computes the \(L U\) factorization of a complex matrix in band storage mode and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex NLCA + NUCA +1 by N array containing the N by N matrix in band storage mode to be factored. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-\) Complex 2 * NLCA \(+\mathrm{NUCA}+1\) by N array containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, \(A\) can share the first (NLCA \(+N U C A+1) * N\) locations with FACT .
IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization. (Output)

RCOND - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A. (Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCCB (A, NLCA, NUCA, FACT, IPVT, RCOND [, ...])
Specific: The specific interface names are S_LFCCB and D_LFCCB.

\section*{FORTRAN 77 Interface}

Single: CALL LFCCB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, RCOND)
Double: The double precision name is DLFCCB.

\section*{Description}

Routine LFCCB performs an \(L U\) factorization of a complex banded coefficient matrix. It also estimates the condition number of the matrix. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm.

The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCCB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) is singular or very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with IMSL routines LFICB, LFSCB and LFDCB. To solve systems of equations with multiple right-hand-side vectors, use LFCCB followed
by either LFICB or LFSCB called once for each right-hand side. The routine LFDCB can be called to compute the determinant of the coefficient matrix after LFCCB has performed the factorization.

Let \(F\) be the matrix FACT , let \(m_{l}=\mathrm{NLCA}\) and let \(m_{u}=\) NUCA. The first \(m_{l}+m_{u}+1\) rows of \(F\) contain the triangular matrix \(U\) in band storage form. The lower \(m_{l}\) rows of \(F\) contain the multipliers needed to reconstruct \(L\).

LFCCB is based on the LINPACK routine CGBCO; see Dongarra et al. (1979). CGBCO uses unscaled partial pivoting.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 C C B / D L 2 C C B\). The reference is:

CALL L2CCB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, RCOND, WK)

The additional argument is
\(W K\) - Complex work vector of length N.
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
42 The input matrix is singular.

\section*{Example}

The inverse of a \(4 \times 4\) band matrix with one upper and one lower codiagonal is computed. LFCCB is called to factor the matrix and to check for singularity or ill-conditioning. LFICB is called to determine the columns of the inverse.
```

USE LFCCB INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE LFICB_INT
USE WRCRN_INT
Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT (N)
REAL RCOND
COMPLEX A (LDA,N), AINV (N,N), FACT(LDFACT,N), RJ (N), RES (N)
Set values for A in band form
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( 0.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 4.0+1.0i 0.0+2.0i 0.0+0.0i )
DATA A/ (0.0,0.0), (0.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0),\&

```
```

        (4.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),&
        (1.0,-1.0), (0.0,0.0)/
    !
CALL LFCCB (A, NLCA, NUCA, FACT, IPVT, RCOND)
Print the reciprocal condition number
and the L1 condition number
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) RCOND, 1.0E0/RCOND
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0E0,0.0EO)
DO 10 J=1, N
RJ(J) = (1.0E0,0.0EO)
RJ is the J-th column of the identity
matrix so the following LFICB
reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFICB (A, NLCA, NUCA, FACT, IPVT, RJ, AINV(:,J), RES)
RJ(J) = (0.0E0,0.0EO)
1 0 ~ C O N T I N U E ~
! Print results
CALL WRCRN ('AINV', AINV)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 condition number = ',F6.3)
END

```

\section*{Output}

RCOND \(=0.022\)
L1 condition number \(=45.933\)
\begin{tabular}{ccccccc} 
\\
& & \multicolumn{4}{c}{ AINV } \\
1 & \((0.562\), & \(0.170)\) & \((0.125,0.260)\) & \((-0.385,-0.135)\) & \((-0.239,-1.165)\) \\
2 & \((0.122\), & \(0.421)\) & \((-0.195,0.094)\) & \((0.101,-0.289)\) & \((0.874,-0.179)\) \\
3 & \((0.034\), & \(0.904)\) & \((-0.437,0.090)\) & \((-0.153,-0.527)\) & \((1.087,-1.172)\) \\
4 & \((0.938\), & \(0.870)\) & \((-0.347\), & \(0.527)\) & \((-0.679,-0.374)\) & \((0.415,-1.759)\)
\end{tabular}

\section*{LFTCB}

Computes the \(L U\) factorization of a complex matrix in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{A}-\) Complex NLCA + NUCA +1 by N array containing the N by N matrix in band storage mode to be factored. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-\) Complex \(2 *\) NLCA \(+\mathrm{NUCA}+1\) by N array containing the \(L U\) factorization of the matrix A. (Output)
If \(A\) is not needed, \(A\) can share the first \((N L C A+N U C A+1) * N\) locations with FACT .
IPVT - Integer vector of length N containing the pivoting information for the \(L U\)
factorization. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(N=\operatorname{size}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size \((\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFTCB (A, NLCA, NUCA, FACT, IPVT [, ...])

Specific: The specific interface names are S_LFTCB and D_LFTCB.

\section*{FORTRAN 77 Interface}

Single: CALL LFTCB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT)
Double: The double precision name is DLFTCB.

\section*{Description}

Routine LFTCB performs an \(L U\) factorization of a complex banded coefficient matrix. The \(L U\) factorization is done using scaled partial pivoting. Scaled partial pivoting differs from partial pivoting in that the pivoting strategy is the same as if each row were scaled to have the same \(\infty\)-norm.

LFTCB fails if \(U\), the upper triangular part of the factorization, has a zero diagonal element. This can occur only if \(A\) is singular or very close to a singular matrix.

The \(L U\) factors are returned in a form that is compatible with routines LFICB, LFSCB and LFDCB. To solve systems of equations with multiple right-hand-side vectors, use LFTCB followed by either LFICB or LFSCB called once for each right-hand side. The routine LFDCB can be called to compute the determinant of the coefficient matrix after LFTCB has performed the factorization.

Let \(F\) be the matrix FACT, let \(m_{l}=\) NLCA and let \(m_{u}=\) NUCA. The first \(m_{l}+m_{u}+1\) rows of \(F\) contain the triangular matrix \(U\) in band storage form. The lower \(m_{l}\) rows of \(F\) contain the multipliers needed to reconstruct \(L^{-1}\). LFTCB is based on the LINPACK routine CGBFA; see Dongarra et al. (1979). CGBFA uses unscaled partial pivoting.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2TCB/DL2TCB The reference is:

CALL L2TCB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, WK)
The additional argument is:
WK - Complex work vector of length N used for scaling.
2. Informational error
\begin{tabular}{ccc} 
Type & Code & \\
4 & 2 & The input matrix is singular.
\end{tabular}

\section*{Example}

A linear system with multiple right-hand sides is solved. LFTCB is called to factor the coefficient matrix. LFSCB is called to compute the two solutions for the two right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call LFCCB to perform the factorization, and LFICB to compute the solutions.
```

USE LFTCB_INT
USE LFSCB INT
USE WRCRN_INT
INTEGER LDA, LDFACT, N, NLCA, NUCA
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT (N)
COMPLEX A(LDA,N), B(N,2), FACT (LDFACT,N), X(N,2)
Set values for A in band form, and B
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( 0.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 4.0+1.0i 0.0+2.0i 0.0+0.0i )
B = ( -4.0-5.0i 16.0-4.0i )
( 9.5+5.5i -9.5+19.5i )
( 9.0-9.0i 12.0+12.0i )
( 0.0+8.0i -8.0-2.0i )
DATA A/ (0.0,0.0), (0.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0),\&
(4.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),\&
(1.0,-1.0), (0.0,0.0)/

```
```

    DATA B/ (-4.0,-5.0), (9.5,5.5), (9.0,-9.0), (0.0,8.0),&
                (16.0,-4.0), (-9.5.19.5), (12.0,12.0), (-8.0,-2.0)/
    !
CALL LFTCB (A, NLCA, NUCA, FACT, IPVT)
DO 10 J=1, 2
CALL LFSCB (FACT, NLCA, NUCA, IPVT, B(:,J), X(:,J))
1 0 CONTINUE
! Print results
CALL WRCRN ('X', X)
!
END

```

\section*{Output}
X
\begin{tabular}{cc}
1 & 2 \\
\((3.000,0.000)\) & \((0.000,4.000)\) \\
\((-1.000,1.000)\) & \((1.000,-1.000)\) \\
\((3.000,0.000)\) & \((0.000,4.000)\) \\
\((-1.000,1.000)\) & \((1.000,-1.000)\)
\end{tabular}

\section*{LFSCB}

Solves a complex system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{F A C T}\) - Complex 2 * NLCA \(+\mathrm{NUCA}+1\) by N array containing the \(L U\) factorization of the coefficient matrix A as output from subroutine LFCCB/DLFCCB or LFTCB/DLFTCB. (Input)

NLCA — Number of lower codiagonals of A. (Input)
\(N U C A\) - Number of upper codiagonals of A. (Input)
\(\boldsymbol{I P V T}\) - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from subroutine \(\mathrm{LFCCB} / \mathrm{DLFCCB}\) or \(\mathrm{LFTCB} / \mathrm{DLFTCB}\). (Input)

B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output) If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSCB (FACT, NLCA, NUCA, IPVT, B, X [,...])
Specific: The specific interface names are S_LFSCB and D_LFSCB.

\section*{FORTRAN 77 Interface}

Single: CALL LFSCB (N, FACT, LDFACT, NLCA, NUCA, IPVT, B, IPATH, X)
Double: The double precision name is DLFSCB.

\section*{Description}

Routine LFSCB computes the solution of a system of linear algebraic equations having a complex banded coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCB or LFTCB. The solution to \(A x=b\) is found by solving the banded triangular systems \(L y=b\) and \(U x=y\). The forward elimination step consists of solving the system \(L y=b\) by applying the same permutations and elimination operations to \(b\) that were applied to the columns of \(A\) in the factorization routine. The backward substitution step consists of solving the banded triangular system \(U x=y\) for \(x\).

LFSCB and LFICB both solve a linear system given its \(L U\) factorization. LFICB generally takes more time and produces a more accurate answer than LFSCB. Each iteration of the iterative refinement algorithm used by LFICB calls LFSCB.
LFSCB is based on the LINPACK routine CGBSL; see Dongarra et al. (1979).

\section*{Example}

The inverse is computed for a real banded \(4 \times 4\) matrix with one upper and one lower codiagonal. The input matrix is assumed to be well-conditioned; hence LFTCB is used rather than LFCCB.
```

USE LFSCB_INT
USE LFTCB INT
USE WRCRN_INT
! Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
COMPLEX A(LDA,N), AINV (N,N), FACT(LDFACT,N), RJ(N)

```
```

        Set values for A in band form
        A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
        ( -2.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
        ( 6.0+1.0i 1.0+1.0i 0.0+2.0i 0.0+0.0i )
    DATA A/ (0.0,0.0), (-2.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0), &
        (1.0,1.0),(-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),&
        (1.0,-1.0), (0.0,0.0)/
    CALL LFTCB (A, NLCA, NUCA, FACT, IPVT)
                Set up the columns of the identity
                matrix one at a time in RJ
    RJ = (0.0EO,0.0EO)
    DO 10 J=1, N
        RJ(J) = (1.0E0,0.0E0)
        RJ is the J-th column of the identity
                matrix so the following LFSCB
                reference places the J-th column of
                the inverse of A in the J-th column
                of AINV
        CALL LFSCB (FACT, NLCA, NUCA, IPVT, RJ, AINV (:,J))
        RJ(J) = (0.0E0,0.0E0)
    1 0 ~ C O N T I N U E
    CALL WRCRN ('AINV', AINV)
    END
    ```

\section*{Output}
\begin{tabular}{ccccc} 
& \multicolumn{3}{c}{2} \\
\((0.165,-0.341)\) & \((0.376,-0.094)\) & \((-0.282,0.471)\) & \((-1.600,0.000)\) \\
\((0.588,-0.047)\) & \((0.259,0.235)\) & \((-0.494,0.024)\) & \((-0.800,-1.200)\) \\
\((0.318,-0.271)\) & \((0.012,0.247)\) & \((-0.759,-0.235)\) & \((-0.550,-2.250)\) \\
\((0.588,-0.047)\) & \((0.259,0.235)\) & \((-0.994,0.524)\) & \((-2.300,-1.200)\)
\end{tabular}

\section*{LFICB}

Uses iterative refinement to improve the solution of a complex system of linear equations in band storage mode.

\section*{Required Arguments}

A-Complex NLCA + NUCA +1 by N array containing the N by N coefficient matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)

NUCA - Number of upper codiagonals of A. (Input)
\(\boldsymbol{F A C T}-\) Complex 2 * NLCA + NUCA +1 by N array containing the \(L U\) factorization of the matrix A as output from routine LFCCB/DLFCCB or LFTCB/DLFTCB. (Input)

IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization of A as output from routine \(\operatorname{LFCCB} / D L F C C B\) or LFTCB/DLFTCB. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution. (Output)
\(\boldsymbol{R E S}\) - Complex vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input) Default: \(N=\operatorname{size}(A, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input) Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A X=B\) is solved.
IPATH \(=2\) means the system \(A^{H} X=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFICB (A, NLCA, NUCA, FACT, IPVT, B, X, RES [,...])
Specific: The specific interface names are S_LFICB and D_LFICB.

\section*{FORTRAN 77 Interface}

Single: CALL LFICB (N, A, LDA, NLCA, NUCA, FACT, LDFACT, IPVT, B, IPATH, X, RES)

Double: The double precision name is DLFICB.

\section*{Description}

Routine LFICB computes the solution of a system of linear algebraic equations having a complex banded coefficient matrix. Iterative refinement is performed on the solution vector to improve the accuracy. Usually almost all of the digits in the solution are accurate, even if the matrix is somewhat ill-conditioned.

To compute the solution, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCB or LFTCB.
Iterative refinement fails only if the matrix is very ill-conditioned.
LFICB and LFSCB both solve a linear system given its \(L U\) factorization. LFICB generally takes more time and produces a more accurate answer than LFSCB. Each iteration of the iterative refinement algorithm used by LFICB calls LFSCB.

\section*{Comments}

Informational error
Type Code
33 The input matrix is too ill-conditioned for iterative refinement be effective.

\section*{Example}

A set of linear systems is solved successively. The right-hand-side vector is perturbed after solving the system each of the first two times by adding \((1+i) / 2\) to the second element.
```

USE LFICB INT
USE LFCCB -INT
USE WRCRN_INT
USE UMACH_INT
Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT (N)
REAL RCOND
COMPLEX A(LDA,N), B(N), FACT(LDFACT,N), RES(N), X(N)
Set values for A in band form, and B
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( -2.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 1.0+1.0i 0.0+2.0i 0.0+0.0i )
B = ( -10.0-5.0i 9.5+5.5i 12.0-12.0i 0.0+8.0i )
DATA A/ (0.0,0.0), (-2.0,-3.0), (6.0.1.0), (4.0,0.0), (-0.5,3.0),\&
(1.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),\&
(1.0,-1.0), (0.0,0.0)/
DATA B/(-10.0,-5.0), (9.5,5.5), (12.0,-12.0), (0.0.8.0)/
CALL LFCCB (A, NLCA, NUCA, FACT, IPVT, RCOND)

```
```

! Print the reciprocal condition number
CALL UMACH (2, NOUT)
WRITE (NOUT,99998) RCOND, 1.0E0/RCOND
!
DO 10 J=1, 3
CALL LFICB (A, NLCA, NUCA, FACT, IPVT, B, X, RES)
Print results
WRITE (NOUT, 99999) J
CALL WRCRN ('X', X, 1, N, 1)
CALL WRCRN ('RES', RES, 1, N, 1)
Perturb B by adding 0.5+0.5i to B(2)
B(2) = B(2) + (0.5E0,0.5E0)
10 CONTINUE
!
99998 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
99999 FORMAT (//,' For system ',I1)
END

```

\section*{Output}
```

RCOND = 0.014
L1 Condition number = 72.414
For system 1
X
(3.000,0.000) (-1.000, 1.000)
$(3.000,0.000)(-1.000,1.000)$
RES
$(0.000 \mathrm{E}+00,0.000 \mathrm{E}+00)^{1} \quad(0.000 \mathrm{E}+00,0.000 \mathrm{E}+00)^{2} \quad(0.000 \mathrm{E}+00,5.684 \mathrm{E}-14)^{3}$
( 3.494E-22,-6.698E-22)
For system 2
X
(3.235, 0.141) (-0.988, 1.247) ( 2 882,0.129)
RES

| $(-1.402 \mathrm{E}-08$, | $6.486 \mathrm{E}-09)^{1}$ |  |
| ---: | ---: | ---: |
| $(-7.012 \mathrm{E}-10,4.488 \mathrm{E}-08)^{4}$ | $(-7.012 \mathrm{E}-10,4.488 \mathrm{E}-08)^{2}$ | $\left.(-1.122 \mathrm{E}-07,7.188 \mathrm{E}-09)^{3}\right)$ |

For system 3
X
$(3.471,0.28)^{1} \quad(-0.976,1.494)^{2} \quad(2.765,0.259)^{3} \quad(-0.976,1.494)^{4}$

```
RES
\((-2.805 \mathrm{E}-08,1.297 \mathrm{E}-08)^{1}\)
\((-1.402 \mathrm{E}-09,-2.945 \mathrm{E}-08)^{2}\)\(\quad(-1.402 \mathrm{E}-09,-2.945 \mathrm{E}-08)^{2} \quad(1.402 \mathrm{E}-08,1.438 \mathrm{E}-08)^{3}\)

\section*{LFDCB}

Computes the determinant of a complex matrix given the \(L U\) factorization of the matrix in band storage mode.

\section*{Required Arguments}
\(\boldsymbol{F A C T}\) - Complex \((2 *\) NLCA + NUCA +1\()\) by N array containing the \(L U\) factorization of the matrix A as output from routine LFTCB/DLFTCB or LFCCB/DLFCCB. (Input)

NLCA - Number of lower codiagonals in matrix A. (Input)

NUCA - Number of upper codiagonals in matrix A. (Input)
IPVT - Vector of length N containing the pivoting information for the \(L U\) factorization as output from routine \(L F T C B / D L F T C B\) or \(\operatorname{LFCCB} / D L F C C B\). (Input)

DET1 - Complex scalar containing the mantissa of the determinant. (Output) The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).

DET2 - Scalar containing the exponent of the determinant. (Output) The determinant is returned in the form \(\operatorname{det}(A)=\) DET1 \(* 10^{\text {DET2 }}\).

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\) size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDCB (FACT, NLCA, NUCA, IPVT, DET1, DET2 [,...])

Specific: The specific interface names are \(S_{\text {_LFDCB }}\) and \(D_{\text {_LFDCB }}\).

\section*{FORTRAN 77 Interface}

Single: CALL LFDCB (N, FACT, LDFACT, NLCA, NUCA, IPVT, DET1, DET2)

Double: \(\quad\) The double precision name is DLFDCB.

\section*{Description}

Routine LFDCB computes the determinant of a complex banded coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(L U\) factorization. This may be done by calling either LFCCB or LFTCB. The formula \(\operatorname{det} A=\operatorname{det} L \operatorname{det} U\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} U=\prod_{i=1}^{N} U_{i i}
\]
(The matrix \(U\) is stored in the upper NUCA + NLCA +1 rows of FACT as a banded matrix.) Since \(L\) is the product of triangular matrices with unit diagonals and of permutation matrices, \(\operatorname{det} L=(-1)^{k}\), where \(k\) is the number of pivoting interchanges.

LFDCB is based on the LINPACK routine CGBDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a complex banded \(4 \times 4\) matrix with one upper and one lower codiagonal.
```

USE LFDCB_INT
USE LFTCB_INT
USE UMACH_INT
D Declare variables
INTEGER LDA, LDFACT, N, NLCA, NUCA, NOUT
PARAMETER (LDA=3, LDFACT=4, N=4, NLCA=1, NUCA=1)
INTEGER IPVT(N)
REAL DET2
COMPLEX A(LDA,N), DET1, FACT (LDFACT,N)
Set values for A in band form
A = ( 0.0+0.0i 4.0+0.0i -2.0+2.0i -4.0-1.0i )
( -2.0-3.0i -0.5+3.0i 3.0-3.0i 1.0-1.0i )
( 6.0+1.0i 1.0+1.0i 0.0+2.0i 0.0+0.0i )
DATA A/ (0.0,0.0), (-2.0,-3.0), (6.0,1.0), (4.0,0.0), (-0.5,3.0),\&
(1.0,1.0), (-2.0,2.0), (3.0,-3.0), (0.0,2.0), (-4.0,-1.0),\&
(1.0,-1.0), (0.0,0.0)/
CALL LFTCB (A, NLCA, NUCA, FACT, IPVT)
Compute the determinant
CALL LFDCB (FACT, NLCA, NUCA, IPVT, DET1, DET2)
Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
99999 FORMAT (' The determinant of A is (', F6.3, ',', F6.3, ') * 10**',\&
F2.0)

```
\(!\)

\section*{Output}

The determinant of \(A\) is ( \(2.500,-1.500) * 10 * * 1\).

\section*{LSAQH}

Solves a complex Hermitian positive definite system of linear equations in band Hermitian storage mode with iterative refinement.

\section*{Required Arguments}

A-Complex NCODA +1 by N array containing the N by N positive definite band Hermitian coefficient matrix in band Hermitian storage mode. (Input)

NCODA - Number of upper or lower codiagonals of A. (Input)
B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size \((A, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSAQH (A, NCODA, B, X [, ...])

Specific: The specific interface names are S_LSAQH and D_LSAQH.

\section*{FORTRAN 77 Interface}

Single: CALL LSAQH (N, A, LDA, NCODA, B, X)
Double: The double precision name is DLSAQH.

\section*{Description}

Routine LSAQH solves a system of linear algebraic equations having a complex Hermitian positive definite band coefficient matrix. It first uses the IMSL routine LFCQH to compute an \(R^{H} R\)

Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. \(R\) is an upper triangular band matrix. The solution of the linear system is then found using the iterative refinement IMSL routine LFIQH.

LSAQH fails if any submatrix of \(R\) is not positive definite, if \(R\) has a zero diagonal element, or if the iterative refinement agorithm fails to converge. These errors occur only if the matrix \(A\) either is very close to a singular matrix or is a matrix that is not positive definite.
If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system. LSAQH solves the problem that is represented in the computer; however, this problem may differ from the problem whose solution is desired.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2AQH/DL2AQH The reference is:
```

CALL L2AQH (N, A, LDA, NCODA, B, X, FACT, WK)

```

The additional arguments are as follows:
FACT - Complex work vector of length (NCODA +1 ) * N containing the \(R^{H} R\) factorization of A in band Hermitian storage form on output.
\(\boldsymbol{W K}\) - Complex work vector of length N .
2. Informational errors

\section*{Type Code}

33 The input matrix is too ill-conditioned. The solution might not be accurate.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix is not positive definite.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine \(L 2 A Q H\) the leading dimension of \(F A C T\) is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of \(\operatorname{IVAL}(4)\). The values \(\operatorname{IVAL}(3)\) and \(\operatorname{IVAL}(4)\) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSAQH. Additional memory allocation for FACT and option value restoration are done automatically in LSAQH. Users directly calling L2AQH can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSAQH or L2AQH. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSAQH temporarily replaces IVAL(2) by IVAL(1). The routine L 2 CQH computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L 2 CQH skips this computation. LSAQH restores the option. Default values for the option \(\operatorname{are} \operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of five linear equations is solved. The coefficient matrix has complex Hermitian positive definite band form with one codiagonal and the right-hand-side vector \(b\) has five elements.
```

USE LSAQH INT
USE WRCRN_INT
INTEGER LDA, N, NCODA
PARAMETER (LDA=2, N=5, NCODA=1)
COMPLEX A(LDA,N), B(N), X(N)
Set values for A in band Hermitian form, and B
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )
B=( 1.0+5.0i 12.0-6.0i 1.0-16.0i -3.0-3.0i 25.0+16.0i )
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0.1.0), (4.0, 0.0), (1.0, 2.0),\&
(10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
DATA B/(1.0,5.0), (12.0,-6.0), (1.0,-16.0), (-3.0,-3.0),\&
(25.0,16.0)/
CALL LSAQH (A, NCODA, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}

\section*{X}
\((2.000,1.000)^{1} \quad(3.000,0.000)^{2}(-1.000,-1.000)^{3}(0.000,-2.000)^{4}\)
( 3.000, 2.000)

\section*{LSLQH}

Solves a complex Hermitian positive definite system of linear equations in band Hermitian storage mode without iterative refinement.

\section*{Required Arguments}

A-Complex NCODA +1 by N array containing the N by N positive definite band Hermitian coefficient matrix in band Hermitian storage mode. (Input)

NCODA - Number of upper or lower codiagonals of A. (Input)
B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLQH (A, NCODA, B, X [,...])

Specific: The specific interface names are S_LSLQH and D_LSLQH.

\section*{FORTRAN 77 Interface}

Single: CALL LSLQH (N, A, LDA, NCODA, B, X)
Double: The double precision name is DLSLQH.

\section*{Description}

Routine LSLQH solves a system of linear algebraic equations having a complex Hermitian positive definite band coefficient matrix. It first uses the routine \(\angle F C Q H\) to compute an \(R^{H} R\) Cholesky factorization of the coefficient matrix and to estimate the condition number of the matrix. \(R\) is an upper triangular band matrix. The solution of the linear system is then found using the routine LFSQH.

LSLQH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) either is very close to a singular matrix or is a matrix that is not positive definite.

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in A can cause very large changes in the solution \(x\). If the coefficient matrix is ill-conditioned or poorly sealed, it is recommended that LSAQH be used.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 L Q H / D L 2 L Q H\) The reference is:
```

CALL L2LQH (N, A, LDA, NCODA, B, X, FACT, WK)

```

The additional arguments are as follows:
\(\boldsymbol{F A C T}-(\mathrm{NCODA}+1) \times \mathrm{N}\) complex work array containing the \(R^{H} R\) factorization of A in band Hermitian storage form on output. If A is not needed, A and FACT can share the same storage locations.

WK - Complex work vector of length N.
2. Informational errors

Type Code
33 The input matrix is too ill-conditioned. The solution might not be accurate.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix is not positive definite.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine \(L 2 L Q H\) the leading dimension of FACT is increased by \(\operatorname{IVAL}(3)\) when \(N\) is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSLQH.
Additional memory allocation for FACT and option value restoration are done automatically in LSLQH. Users directly calling L2LQH can allocate additional space for FACT and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSLQH or L2LQH. Default values for the option are \(\operatorname{IVAL}(*)=1,16,0,1\).

17 This option has two values that determine if the \(L_{1}\) condition number is to be computed. Routine LSLQH temporarily replaces IVAL(2) by IVAL(1). The routine \(L 2 \mathrm{CQH}\) computes the condition number if \(\operatorname{IVAL}(2)=2\). Otherwise L 2 CQH skips this computation. LSLQH restores the option. Default values for the option are \(\operatorname{IVAL}(*)=1,2\).

\section*{Example}

A system of five linear equations is solved. The coefficient matrix has complex Hermitian positive definite band form with one codiagonal and the right-hand-side vector \(b\) has five elements.
```

    USE LSLQH_INT
    USE WRCRN_INT
    ! Declare variables
INTEGER N, NCODA, LDA
PARAMETER (N=5, NCODA=1, LDA=NCODA+1)
COMPLEX A(LDA,N), B(N), X(N)
Set values for A in band Hermitian form, and B
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )
B = ( 1.0+5.0i 12.0-6.0i 1.0-16.0i -3.0-3.0i 25.0+16.0i )
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0),\&
(10.0,0.0), (0.0,4.0), (6.0.0.0), (1.0,1.0), (9.0,0.0)/
DATA B/(1.0,5.0), (12.0,-6.0), (1.0,-16.0), (-3.0,-3.0),\&
(25.0,16.0)/
CALL LSLQH (A, NCODA, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\((2.000,1.000)^{1}(3.000,0.000)^{2}(-1.000,-1.000)^{3} \quad(0.000,-2.000)^{4}\)

5
\((3.000,2.000)\)

\section*{LSLQB}

Computes the \(R^{H} D R\) Cholesky factorization of a complex Hermitian positive-definite matrix \(A\) in codiagonal band Hermitian storage mode. Solve a system \(A x=b\).

\section*{Required Arguments}

A - Array containing the N by n positive-definite band coefficient matrix and the right hand side in codiagonal band Hermitian storage mode. (Input/Output)
The number of array columns must be at least 2 * NCODA +3 . The number of columns is not an input to this subprogram.

NCODA - Number of upper codiagonals of matrix A. (Input)
Must satisfy NCODA \(\geq 0\) and NCODA \(<\mathrm{N}\).
\(\boldsymbol{U}\) - Array of flags that indicate any singularities of A, namely loss of positive-definiteness of a leading minor. (Output)

A value \(U(I)=0\). means that the leading minor of dimension \(I\) is not positive-definite. Otherwise, \(\mathrm{U}(\mathrm{I})=1\).

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Must satisfy \(\mathrm{N}>0\).
Default: N = size (A,2).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Must satisfy LDA \(\geq \mathrm{N}+\mathrm{NCODA}\).
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
IJOB - flag to direct the desired factorization or solving step. (Input)
Default: \(\mathrm{IJOB}=1\).
IJOB Meaning
1 factor the matrix \(A\) and solve the system \(\mathrm{A} x=b\); where the real part of \(b\) is stored in column \(2 *\) NCODA +2 and the imaginary part of \(b\) is stored in column \(2 *\) NCODA +3 of array A. The real and imaginary parts of \(b\) are overwritten by the real and imaginary parts of \(x\).

2 solve step only. Use the real part of \(b\) as column 2 * NCODA +2 and the imaginary part of \(b\) as column \(2 *\) NCODA +3 of A. (The factorization step has already been done.) The real and imaginary parts of \(b\) are overwritten by the real and imaginary parts of \(x\).

3 factor the matrix A but do not solve a system.
\(4,5,6\) same meaning as with the value \(I J O B=3\). For efficiency, no error checking is done on values LDA, \(N\), NCODA, and \(U(*)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLQB (A, NCODA, U [, ...])
Specific: The specific interface names are S_LSLQB and D_LSLQB.

\section*{FORTRAN 77 Interface}

Single: CALL LSLQB (N, A, LDA, NCODA, IJOB, U)
Double: The double precision name is DLSLQB.

\section*{Description}

Routine LSLQB factors and solves the Hermitian positive definite banded linear system \(A x=b\). The matrix is factored so that \(A=R^{H} D R\), where \(R\) is unit upper triangular and \(D\) is diagonal and real. The reciprocals of the diagonal entries of \(D\) are computed and saved to make the solving step more efficient. Errors will occur if \(D\) has a nonpositive diagonal element. Such events occur only if \(A\) is very close to a singular matrix or is not positive definite.

LSLQB is efficient for problems with a small band width. The particular cases NCODA \(=0,1\) are done with special loops within the code. These cases will give good performance. See Hanson (1989) for more on the algorithm. When solving tridiagonal systems, \(\mathrm{NCODA}=1\), the cyclic reduction code LSLCQ should be considered as an alternative. The expectation is that LSLCQ will outperform LSLQB on vector or parallel computers. It may be inferior on scalar computers or even parallel computers with non-optimizing compilers.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 L Q B / D L 2 L Q B\) The reference is:

CALL L2LQB (N, A, LDA, NCODA, IJOB, U, WK1, WK2)
The additional arguments are as follows:

WK1 - Work vector of length NCODA.

WK2 - Work vector of length NCODA.
2. Informational error

Type Code
42 The input matrix is not positive definite.

\section*{Example}

A system of five linear equations is solved. The coefficient matrix has real positive definite codiagonal Hermitian band form and the right-hand-side vector \(b\) has five elements.
```

USE LSLQB_INT
USE WRRRN_INT
INTEGER LDA, N, NCODA
PARAMETER (N=5, NCODA=1, LDA=N+NCODA)
INTEGER I, IJOB, J
REAL A(LDA,2*NCODA+3),U(N)
Set values for A and right hand side
in codiagonal band Hermitian form:
$\left.A=\begin{array}{rcccrr}(* & * & * & * & *\end{array}\right)$

```
\(!\)
```

! ( 6.0 0.0 0. 4.0 -3.0 - -3.0)
DATA ((A (I+NCODA,J),I=1,N),J=1,2*NCODA+3)/2.0, 4.0, 10.0, 6.0,\&
9.0, 0.0, -1.0, 1.0, 0.0, 1.0, 0.0, 1.0, 2.0, 4.0, 1.0,\&
1.0, 12.0, 1.0, -3.0, 25.0, 5.0, -6.0, -16.0, -3.0, 16.0/
Factor and solve A*x = b.
IJOB = 1
CALL LSLQB (A, NCODA, U)
!
Print results
CALL WRRRN ('REAL(X)', A((NCODA+1):,(2*NCODA+2):), 1, N, 1)
CALL WRRRN ('IMAG(X)', A((NCODA+1):,(2*NCODA+3):), 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrrrr} 
& & \multicolumn{4}{c}{ REAL (X) } \\
2.000 & 3.000 & -1.000 & 0.000 & 3.000
\end{tabular}
\begin{tabular}{rrrrr} 
& \multicolumn{4}{c}{ IMAG (X) } \\
1 & 2 & 3 & 4 & 5 \\
1.000 & 0.000 & -1.000 & -2.000 & 2.000
\end{tabular}

\section*{LFCQH}

Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode and estimate its \(L_{1}\) condition number.

\section*{Required Arguments}

A-Complex NCODA +1 by N array containing the N by N positive definite band Hermitian matrix to be factored in band Hermitian storage mode. (Input)

NCODA - Number of upper or lower codiagonals of A. (Input)
FACT - Complex NCODA +1 by N array containing the \(R^{H} R\) factorization of the matrix A. (Output)
If A is not needed, A and FACT can share the same storage locations.
RCOND - Scalar containing an estimate of the reciprocal of the \(L_{1}\) condition number of A.
(Output)

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(F A C T, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFCQH (A, NCODA, FACT, RCOND [, ...])

Specific: The specific interface names are S_LFCQH and D_LFCQH.

\section*{FORTRAN 77 Interface}

Single: CALL LFCQH (N, A, LDA, NCODA, FACT, LDFACT, RCOND)
Double: The double precision name is DLFCQH.

\section*{Description}

Routine LFCQH computes an \(R^{H} R\) Cholesky factorization and estimates the condition number of a complex Hermitian positive definite band coefficient matrix. \(R\) is an upper triangular band matrix.
The \(L_{1}\) condition number of the matrix \(A\) is defined to be \(\kappa(A)=\|A\|_{1}\left\|A^{-1}\right\|_{1}\). Since it is expensive to compute \(\left\|A^{-1}\right\|_{1}\), the condition number is only estimated. The estimation algorithm is the same as used by LINPACK and is described by Cline et al. (1979).

If the estimated condition number is greater than \(1 / \varepsilon\) (where \(\varepsilon\) is machine precision), a warning error is issued. This indicates that very small changes in \(A\) can cause very large changes in the solution \(x\). Iterative refinement can sometimes find the solution to such a system.

LFCQH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) either is very close to a singular matrix or is a matrix which is not positive definite.
The \(R^{H} R\) factors are returned in a form that is compatible with routines LFIQH, LFSQH and LFDQH. To solve systems of equations with multiple right-hand-side vectors, use LFCQH followed by either \(L F I Q H\) or \(L F S Q H\) called once for each right-hand side. The routine \(L F D Q H\) can be called to compute the determinant of the coefficient matrix after LFCQH has performed the factorization.

LFCQH is based on the LINPACK routine CPBCO; see Dongarra et al. (1979).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(\mathrm{L} 2 \mathrm{CQH} / \mathrm{DL} 2 \mathrm{CQH}\). The reference is:

CALL L2CQH (N, A, LDA, NCODA, FACT, LDFACT, RCOND, WK)
The additional argument is:
\(\boldsymbol{W} \boldsymbol{K}\) - Complex work vector of length N .
2. Informational errors

Type Code
31 The input matrix is algorithmically singular.
34 The input matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The input matrix is not positive definite.
44 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part

\section*{Example}

The inverse of a \(5 \times 5\) band Hermitian matrix with one codiagonal is computed. LFCQH is called to factor the matrix and to check for nonpositive definiteness or ill-conditioning. LFIQH is called to determine the columns of the inverse.
```

USE LFCQH_INT
USE LFIQH_INT
USE UMACH_INT
USE WRCRN_INT
INTEGER N, NCODA, LDA, LDFACT, NOUT
PARAMETER (N=5, NCODA=1, LDA=NCODA+1, LDFACT=LDA)
REAL RCOND
COMPLEX A(LDA,N), AINV (N,N), FACT(LDFACT,N), RES (N), RJ (N)
Set values for A in band Hermitian form
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0), \&
(10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
Factor the matrix A
CALL LFCQH (A, NCODA, FACT, RCOND)
Set up the columns of the identity
matrix one at a time in RJ
RJ = (0.0E0,0.0E0)
DO 10 J=1, N
RJ(J) = (1.0E0,0.0E0)
RJ is the J-th column of the identity
matrix so the following LFIQH

```
```

! reference places the J-th column of
the inverse of A in the J-th column
of AINV
CALL LFIQH (A, NCODA, FACT, RJ, AINV(:,J), RES)
RJ(J) = (0.0E0,0.0E0)
1 0 CONTINUE
! Print the results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) RCOND, 1.0E0/RCOND
CALL WRCRN ('AINV', AINV)
!
99999 FORMAT (' RCOND = ',F5.3,/,' L1 Condition number = ',F6.3)
END

```

\section*{Output}

RCOND \(=0.067\)
L1 Condition number \(=14.961\)
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{4}{|c|}{AINV} \\
\hline 1 & 2 & 3 & 4 \\
\hline \(1(0.7166,0.0000)\) & ( 0.2166,-0.2166) & (-0.0899, -0.0300) & \((-0.0207,0.0622)\) \\
\hline \(2(0.2166,0.2166)\) & ( 0.4332, 0.0000) & (-0.0599,-0.1198) & (-0.0829, 0.0415) \\
\hline \(3(-0.0899,0.0300)\) & (-0.0599, 0.1198) & ( 0.1797, 0.0000) & ( 0.0000, -0.1244) \\
\hline 4 (-0.0207,-0.0622) & (-0.0829,-0.0415) & ( 0.0000, 0.1244) & ( 0.2592, 0.0000) \\
\hline \[
5(0.0092,0.0046)_{5}
\] & ( 0.0138,-0.0046) & (-0.0138, -0.0138) & (-0.0288, 0.0288) \\
\hline \(1(0.0092,-0.0046)\) & & & \\
\hline \(2(0.0138,0.0046)\) & & & \\
\hline 3 (-0.0138, 0.0138) & & & \\
\hline \(4(-0.0288,-0.0288)\) & & & \\
\hline \(5 \quad(0.1175,0.0000)\) & & & \\
\hline
\end{tabular}

\section*{LFTQH}

Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode.

\section*{Required Arguments}

A-Complex NCODA +1 by N array containing the N by N positive definite band Hermitian matrix to be factored in band Hermitian storage mode. (Input)

NCODA - Number of upper or lower codiagonals of A. (Input)
FACT - Complex NCODA +1 by N array containing the \(R^{H} R\) factorization of the matrix A. (Output)
If A is not needed, A and FACT can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Order of the matrix. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: \(\operatorname{IDFACT}=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LfTQH (A, NCODA, FACT [, ...])
Specific: The specific interface names are S_LFTQH and D_LFTQH.

\section*{FORTRAN 77 Interface}

Single: CALL LFTQH (N, A, LDA, NCODA, FACT, LDFACT)
Double: \(\quad\) The double precision name is DLFTQH.

\section*{Description}

Routine LFTQH computes an \(R^{H} R\) Cholesky factorization of a complex Hermitian positive definite band coefficient matrix. \(R\) is an upper triangular band matrix.

LFTQH fails if any submatrix of \(R\) is not positive definite or if \(R\) has a zero diagonal element. These errors occur only if \(A\) either is very close to a singular matrix or is a matrix which is not positive definite.
The \(R^{H} R\) factors are returned in a form that is compatible with routines LFIQH, LFSQH and LFDQH. To solve systems of equations with multiple right-hand-side vectors, use \(\operatorname{LFTQH}\) followed by either LFIQH or LFSQH called once for each right-hand side. The routine LFDQH can be called to compute the determinant of the coefficient matrix after LFTQH has performed the factorization.

LfTQH is based on the LINPACK routine SPBFA; see Dongarra et al. (1979).

\section*{Comments}

Informational errors
\begin{tabular}{ccl} 
Type & Code & \\
3 & 4 & The input matrix is not Hermitian. It has a diagonal entry with a small \\
4 & 2 & \begin{tabular}{l} 
imaginary part. \\
The input matrix is not positive definite.
\end{tabular}
\end{tabular}

4 The input matrix is not Hermitian. It has a diagonal entry with an imaginary part.

\section*{Example}

The inverse of a \(5 \times 5\) band Hermitian matrix with one codiagonal is computed. LFTQH is called to factor the matrix and to check for nonpositive definiteness. LFSQH is called to determine the columns of the inverse.
```

    USE LFTQH_INT
    USE LFSQH_INT
    USE WRCRN_INT
    INTEGER LDA, LDFACT, N, NCODA
    PARAMETER (LDA=2, LDFACT=2, N=5, NCODA=1)
    COMPLEX A(LDA,N), AINV (N,N), FACT(LDFACT,N), RJ(N)
        Set values for A in band Hermitian form
        A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
            ( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )
    DATA A/(0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0),&
        (10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
                Factor the matrix A
    CALL LFTQH (A, NCODA, FACT)
        Set up the columns of the identity
                matrix one at a time in RJ
    RJ = (0.0EO,0.0E0)
    DO 10 J=1, N
        RJ(J) = (1.0E0,0.0EO)
                                RJ is the J-th column of the identity
                matrix so the following LFSQH
                reference places the J-th column of
                the inverse of A in the J-th column
                of AINV
        CALL LFSQH (FACT, NCODA, RJ, AINV(:,J))
        RJ(J) = (0.0E0,0.0E0)
    1 0 ~ C O N T I N U E ~
    CALL WRCRN ('AINV', AINV)
    END
    ```

\section*{Output}


\section*{LFSQH}

Solves a complex Hermitian positive definite system of linear equations given the factorization of the coefficient matrix in band Hermitian storage mode.

\section*{Required Arguments}

FACT - Complex NCODA +1 by N array containing the \(R^{H} R\) factorization of the Hermitian positive definite band matrix A. (Input)
FACT is obtained as output from routine LFCQH/DLFCQH or LFTQH/DLFTQH.
NCODA - Number of upper or lower codiagonals of A. (Input)
B - Complex vector of length N containing the right-hand-side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
If \(B\) is not needed, \(B\) and \(X\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).
LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT \(=\operatorname{size}(\) FACT, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSQH (FACT, NCODA, B, X [, ...])
Specific: The specific interface names are S_LFSQH and D_LFSQH.

\section*{FORTRAN 77 Interface}

Single: CALL LFSQH (N, FACT, LDFACT, NCODA, B, X)
Double: The double precision name is DLFSQH.

\section*{Description}

This routine computes the solution for a system of linear algebraic equations having a complex Hermitian positive definite band coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either IMSL routine LFCQH or LFTQH. \(R\) is an upper triangular band matrix.
The solution to \(A x=b\) is found by solving the triangular systems \(R^{H} y=b\) and \(R x=y\).
LFSQH and LFIQH both solve a linear system given its \(R^{H} R\) factorization. LFIQH generally takes more time and produces a more accurate answer than LFSQH. Each iteration of the iterative refinement algorithm used by LFIQH calls LFSQH.

LFSQH is based on the LINPACK routine CPBSL; see Dongarra et al. (1979).

\section*{Comments}

Informational error
Type Code
41 The factored matrix has a diagonal element close to zero.

\section*{Example}

A set of linear systems is solved successively. LFTQH is called to factor the coefficient matrix. LFSQH is called to compute the three solutions for the three right-hand sides. In this case the coefficient matrix is assumed to be well-conditioned and correctly scaled. Otherwise, it would be better to call \(\angle F C Q H\) to perform the factorization, and \(\operatorname{LFIQH}\) to compute the solutions.
```

USE LFSQH INT
USE LFTQH_INT
USE WRCRN_INT
INTEGER LDA, LDFACT, N, NCODA
PARAMETER (LDA=2, LDFACT=2, N=5, NCODA=1)
COMPLEX A(LDA,N), B(N,3), FACT(LDFACT,N), X(N,3)
Set values for A in band Hermitian form, and B
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )
B = ( 3.0+3.0i 4.0+0.0i 29.0-9.0i )
( 5.0-5.0i 15.0-10.0i -36.0-17.0i )
( 5.0+4.0i -12.0-56.0i -15.0-24.0i )
( 9.0+7.0i -12.0+10.0i -23.0-15.0i )
(-22.0+1.0i 3.0-1.0i -23.0-28.0i )
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0),\&
(10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
DATA B/(3.0,3.0), (5.0,-5.0), (5.0.4.0), (9.0,7.0), (-22.0,1.0),\&
(4.0,0.0), (15.0,-10.0), (-12.0,-56.0), (-12.0,10.0),\&
(3.0,-1.0), (29.0,-9.0), (-36.0,-17.0), (-15.0,-24.0),\&
(-23.0,-15.0), (-23.0,-28.0)/
Factor the matrix A

```
```

    CALL LFTQH (A, NCODA, FACT)
    ! Compute the solutions
DO 10 I=1, 3
CALL LFSQH (FACT, NCODA, B(:,I), X(:,I))
10 CONTINUE
! Print solutions
CALL WRCRN ('X', X)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|c|}{X} \\
\hline & 1 & & & 2 & & ) \\
\hline 1.00, & \(0.00)\) & & 3.00 , & -1.00) & ( 11.00, & -1.00) \\
\hline ( 1.00, & -2.00) & ( & 2.00 , & \(0.00)\) & ( -7.00, & 0.00 ) \\
\hline ( 2.00, & \(0.00)\) & & -1.00, & -6.00) & ( -2.00, & -3.00) \\
\hline ( 2.00, & \(3.00)\) & & 2.00, & 1.00) & ( -2.00 , & -3.00) \\
\hline ( -3.00, & \(0.00)\) & ( & 0.00 , & \(0.00)\) & ( -2.00 , & -3.00) \\
\hline
\end{tabular}

\section*{LFIQH}

Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations in band Hermitian storage mode.

\section*{Required Arguments}

A-Complex NCODA +1 by N array containing the N by N positive definite band Hermitian coefficient matrix in band Hermitian storage mode. (Input)

NCODA - Number of upper or lower codiagonals of A. (Input)
FACT - Complex NCODA +1 by N array containing the \(R^{H} R\) factorization of the matrix A as output from routine \(\mathrm{LFCQH} / \mathrm{DLFCQH}\) or \(\mathrm{LFTQH} / \mathrm{DLFTQH}\). (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)
\(\boldsymbol{R E S}\) - Complex vector of length N containing the residual vector at the improved solution. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input) Default: \(N=\operatorname{size}(A, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input) Default: LDA \(=\operatorname{size}(\mathrm{A}, 1)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT, 1 ).

\section*{FORTRAN 90 Interface}

Generic: CALL LFIQH (A, NCODA, FACT, B, X, RES [,...])
Specific: The specific interface names are S_LFIQH and D_LFIQH.

\section*{FORTRAN 77 Interface}

Single: CALL LFIQH (N, A, LDA, NCODA, FACT, LDFACT, B, X, RES)
Double: The double precision name is DLFIQH.

\section*{Description}

This routine computes the solution for a system of linear algebraic equations having a complex Hermitian positive definite band coefficient matrix. To compute the solution, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either IMSL routine LFCQH or LFTQH. \(R\) is an upper triangular band matrix.

The solution to \(A x=b\) is found by solving the triangular systems \(R^{H} y=b\) and \(R x=y\).
LFSQH and LFIQH both solve a linear system given its \(R^{H} R\) factorization. LFIQH generally takes more time and produces a more accurate answer than LFSQH. Each iteration of the iterative refinement algorithm used by LFIQH calls LFSQH.

\section*{Comments}

Informational error
Type Code
41 The factored matrix has a diagonal element close to zero.

\section*{Example}

A set of linear systems is solved successively. The right-hand side vector is perturbed after solving the system each of the fisrt two times by adding \((1+i) / 2\) to the second element.
```

USE IMSL_LIBRARIES
! - Declare variables
INTEGER LDA, LDFACT, N, NCODA
PARAMETER (LDA=2, LDFACT=2, N=5, NCODA=1)
REAL RCOND
COMPLEX A(LDA,N), B(N), FACT(LDFACT,N), RES(N,3), X(N,3)
Set values for A in band Hermitian form, and B
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )

```
```

!
!
B = ( 3.0+3.0i 5.0-5.0i 5.0+4.0i 9.0+7.0i -22.0+1.0i )
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0),\&
(10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
DATA B/(3.0,3.0), (5.0,-5.0), (5.0,4.0), (9.0,7.0), (-22.0,1.0)/
Factor the matrix A
CALL LFCQH (A, NCODA, FACT, RCOND=RCOND)
Print the estimated condition number
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) RCOND, 1.0E0/RCOND
DO 10 I=1, 3
CALL LFIQH (A, NCODA, FACT, B, X(:,I), RES(:,I))
B(2) = B (2) + (0.5E0, 0.5E0)
O CONTINUE
!
CALL WRCRN ('X', X)
CALL WRCRN ('RES', RES)
99999 FORMAT (' RCOND = ', F5.3, /, ' L1 Condition number = ', F6.3)
END

```

\section*{Output}

X


\section*{LFDQH}

Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\) Cholesky factorization in band Hermitian storage mode.

\section*{Required Arguments}
\(\boldsymbol{F A C T}\) - Complex NCODA +1 by N array containing the \(R^{H} R\) factorization of the Hermitian positive definite band matrix A. (Input)
FACT is obtained as output from routine LFCQH/DLFCQH or LFTQH/DLFTQH.
NCODA - Number of upper or lower codiagonals of A. (Input)
DET1 - Scalar containing the mantissa of the determinant. (Output)
The value DET1 is normalized so that \(1.0 \leq \mid\) DET1 \(\mid<10.0\) or DET1 \(=0.0\).
DET2 - Scalar containing the exponent of the determinant. (Output)
The determinant is returned in the form \(\operatorname{det}(A)=\operatorname{DET1} * 10^{\mathrm{DET}}\).

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{FACT}, 2)\).

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT = size (FACT,1).

\section*{FORTRAN 90 Interface}

Generic: CALL LFDQH (FACT, NCODA, DET1, DET2 [,...])
Specific: The specific interface names are S_LFDQH and D_LFDQH.

\section*{FORTRAN 77 Interface}

Single: CALL LFDQH (N, FACT, LDFACT, NCODA, DET1, DET2)
Double: The double precision name is DLFDQH.

\section*{Description}

Routine LFDQH computes the determinant of a complex Hermitian positive definite band coefficient matrix. To compute the determinant, the coefficient matrix must first undergo an \(R^{H} R\) factorization. This may be done by calling either LFCQH or LFTQH. The formula \(\operatorname{det} A=\operatorname{det} R^{H} \operatorname{det} R=(\operatorname{det} R)^{2}\) is used to compute the determinant. Since the determinant of a triangular matrix is the product of the diagonal elements,
\[
\operatorname{det} R=\prod_{i=1}^{N} R_{i i}
\]

LFDQH is based on the LINPACK routine CPBDI; see Dongarra et al. (1979).

\section*{Example}

The determinant is computed for a \(5 \times 5\) complex Hermitian positive definite band matrix with one codiagonal.
```

USE LFDQH_INT
USE LFTQH-INT
USE UMACH_INT
! Declare variables
INTEGER LDA, LDFACT, N, NCODA, NOUT
PARAMETER (LDA=2, N=5, LDFACT=2, NCODA=1)
REAL DET1, DET2
COMPLEX A(LDA,N), FACT(LDFACT,N)
Set values for A in band Hermitian form
A = ( 0.0+0.0i -1.0+1.0i 1.0+2.0i 0.0+4.0i 1.0+1.0i )
( 2.0+0.0i 4.0+0.0i 10.0+0.0i 6.0+0.0i 9.0+0.0i )

```
```

!
DATA A/ (0.0,0.0), (2.0,0.0), (-1.0,1.0), (4.0, 0.0), (1.0,2.0),\&
(10.0,0.0), (0.0,4.0), (6.0,0.0), (1.0,1.0), (9.0,0.0)/
Factor the matrix
CALL LFTQH (A, NCODA, FACT)
Compute the determinant
CALL LFDQH (FACT, NCODA, DET1, DET2)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) DET1, DET2
!
99999 FORMAT (' The determinant of A is ',F6.3,' * 10**',F2.0)
END

```

\section*{Output}

The determinant of A is 1.736 * \(10 * * 3\).

\section*{LSLXG}

Solves a sparse system of linear algebraic equations by Gaussian elimination.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Vector of length NZ containing the nonzero coefficients of the linear system. (Input)
IROW - Vector of length NZ containing the row numbers of the corresponding elements in A. (Input)

JCOL - Vector of length NZ containing the column numbers of the corresponding elements in A. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N-\) Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input)
Default: NZ \(=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A x=b\) is solved.
IPATH \(=2\) means the system \(A^{T} x=b\) is solved.
Default: IPATH \(=1\).

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM.
Default: \(\operatorname{IPARAM}(1)=0\).
See Comment 3.

RPARAM - Parameter vector of length 5. (Input/Output)
See Comment 3.

\section*{FORTRAN 90 Interface}

Generic: CALL LSLXG (A, IROW, JCOL, B, X [, ...])
Specific: The specific interface names are S_LSLXG and D_LSLXG.

\section*{FORTRAN 77 Interface}

Single: CALL LSLXG (N, NZ, A, IROW, JCOL, B, IPATH, IPARAM, RPARAM, X)
Double: The double precision name is DLSLXG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a \(n \times n\) sparse matrix. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column numbers for these entries in \(A\). That is
\[
A_{\text {irow }(i), \operatorname{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero.
The routine LSLXG solves a system of linear algebraic equations having a real sparse coefficient matrix. It first uses the routine LFTXG to perform an \(L U\) factorization of the coefficient matrix. The solution of the linear system is then found using LFSXG.
The routine LFTXG by default uses a symmetric Markowitz strategy (Crowe et al. 1990) to choose pivots that most likely would reduce fill-ins while maintaining numerical stability. Different strategies are also provided as options for row oriented or column oriented problems. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively.
Finally, the solution \(x\) is obtained by the following calculations:
\[
\text { 1) } L z=P b
\]
\[
\begin{aligned}
& \text { 2) } U y=z \\
& \text { 3) } x=Q y
\end{aligned}
\]

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LXG/DL2LXG. The reference is:

CALL L2LXG (N, NZ, A, IROW, JCOL, B, IPATH, IPARAM, RPARAM, X, WK, LWK, IWK, LIWK)

The additional arguments are as follows:

WK - Real work vector of length LWK.
\(\boldsymbol{L W K}\) - The length of wK, LWK should be at least \(2 \mathrm{~N}+\) MAXNz.
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(17 \mathrm{~N}+4\) * MAXNZ.
The workspace limit is determined by MAXNZ, where
MAXNZ \(=\) MINO(LWK-2N, INT (0.25(LIWK-17N)))
2. Informational errors

Type Code
31 The coefficient matrix is numerically singular.
32 The growth factor is too large to continue.
33 The matrix is too ill-conditioned for iterative refinement.
3. If the default parameters are desired for LSLXG, then set IPARAM(1) to zero and call the routine LSLXG. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM. then the following steps should be taken before calling LSLXG.

CALL L4LXG (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to L4LXG will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above.

IPARAM - Integer vector of length 6.
\(\operatorname{IPARAM}(1)=\) Initialization flag.
\(\operatorname{IPARAM}(2)=\) The pivoting strategy
IPARAM(2) Action

1 Markowitz row search
2 Markowitz column search
3 Symmetric Markowitz search
Default: 3 .
\(\operatorname{IPARAM}(3)=\) The number of rows which have least numbers of nonzero elements that will be searched for a pivotal element.
Default: 3.
\(\operatorname{IPARAM}(4)=\) The maximal number of nonzero elements in \(A\) at any stage of the Gaussian elimination. (Output)
\(\operatorname{IPARAM}(5)=\) The workspace limit.
IPARAM(5) Action
0
integer This integer value replaces the default workspace limit.
When L2LXG is called, the values of LWK and LIWK are used instead of I PARAM(5).

Default: 0.
\(\operatorname{IPARAM}(6)=\) Iterative refinement is done when this is nonzero.
Default: 0.
\(\boldsymbol{R P A R A M}\) - Real vector of length 5.
\(\operatorname{RPARAM}(1)=\) The upper limit on the growth factor. The computation stops when the growth factor exceeds the limit.

Default: \(10^{16}\).
\(\operatorname{RPARAM}(2)=\) The stability factor. The absolute value of the pivotal element must be bigger than the largest element in absolute value in its row divided by RPARAM(2).

Default: 10.0
\(\operatorname{RPARAM}(3)=\) Drop-tolerance. Any element in the lower triangular factor \(L\) will be removed if its absolute value becomes smaller than the drop-tolerance at any stage of the Gaussian elimination.

Default: 0.0.
\(\operatorname{RPARAM}(4)=\) The growth factor. It is calculated as the largest element in absolute value in A at any stage of the Gaussian elimination divided by the largest element in absolute value in the original A matrix. (Output) Large value of the growth factor indicates that an appreciable error in the computed solution is possible.
\(\operatorname{RPARAM}(5)=\) The value of the smallest pivotal element in absolute value. (Output)

If double precision is required, then DL4LXG is called and RPARAM is declared double precision.

\section*{Example}

As an example consider the \(6 \times 6\) linear system:
\[
A=\left[\begin{array}{rrrrrr}
10 & 0 & 0 & 0 & 0 & 0 \\
0 & 10 & -3 & -1 & 0 & 0 \\
0 & 0 & 15 & 0 & 0 & 0 \\
-2 & 0 & 0 & 10 & -1 & 0 \\
-1 & 0 & 0 & -5 & 1 & -3 \\
-1 & -2 & 0 & 0 & 0 & 6
\end{array}\right]
\]

Let \(x^{T}=(1,2,3,4,5,6)\) so that \(A x=(10,7,45,33,-34,31)^{T}\). The number of nonzeros in \(A\) is \(\mathrm{nz}=15\). The sparse coordinate form for \(A\) is given by:
\begin{tabular}{rrrrrrrrrrrrrrrr} 
irow & 6 & 2 & 3 & 2 & 4 & 4 & 5 & 5 & 5 & 5 & 1 & 6 & 6 & 2 & 4 \\
jcol & 6 & 2 & 3 & 3 & 4 & 5 & 1 & 6 & 4 & 5 & 1 & 1 & 2 & 4 & 1 \\
a & 6 & 10 & 15 & -3 & 10 & -1 & -1 & -3 & -5 & 1 & 10 & -1 & -2 & -1 & -2
\end{tabular}
```

USE LSLXG INT
USE WRRRN INT
USE L4LXG_INT
INTEGER - N, NZ
PARAMETER (N=6, NZ=15)
INTEGER IPARAM(6), IROW(NZ), JCOL(NZ)
REAL A(NZ), B(N), RPARAM(5), X(N)
DATA A/6., 10., 15., -3., 10., -1., -1., -3., -5., 1., 10., -1.,\&
-2., -1., -2./
DATA B/10., 7., 45., 33., -34., 31./
DATA IROW/6, 2, 3, 2, 4, 4, 5, 5, 5, 5, 1, 6, 6, 2, 4/
DATA JCOL/6, 2, 3, 3, 4, 5, 1, 6, 4, 5, 1, 1, 2, 4, 1/
Change a default parameter
CALL L4LXG (IPARAM, RPARAM)
IPARAM(5) = 203
Solve for X

```
```

CALL LSLXG (A, IROW, JCOL, B, X, IPARAM=IPARAM)

```
\(!\)
CALL WRRRN (' x ', \(\mathrm{X}, \mathrm{I}, \mathrm{N}, 1\) )
END

\section*{Output}
\begin{tabular}{rrrrrrr}
1 & & \multicolumn{5}{c}{\(x\)} \\
1.000 & 2.000 & 3.000 & 4.000 & 5.000 & 6.000
\end{tabular}

\section*{LFTXG}

Computes the \(L U\) factorization of a real general sparse matrix..

\section*{Required Arguments}
\(\boldsymbol{A}\) - Vector of length NZ containing the nonzero coefficients of the linear system. (Input)
IROW - Vector of length NZ containing the row numbers of the corresponding elements in A. (Input)

JCOL - Vector of length NZ containing the column numbers of the corresponding elements in A. (Input)
\(N L\) - The number of nonzero coefficients in the triangular matrix \(L\) excluding the diagonal elements. (Output)

NFAC - On input, the dimension of vector FACT. (Input/Output)
On output, the number of nonzero coefficients in the triangular matrix \(L\) and \(U\).
FACT - Vector of length NFAC containing the nonzero elements of \(L\) (excluding the diagonals) in the first NL locations and the nonzero elements of \(U\) in NL +1 to NFAC locations. (Output)

IRFAC - Vector of length NFAC containing the row numbers of the corresponding elements in FACT. (Output)

JCFAC - Vector of length NFAC containing the column numbers of the corresponding elements in FACT. (Output)

IPVT - Vector of length N containing the row pivoting information for the \(L U\) factorization. (Output)
\(\boldsymbol{J P V T}\) - Vector of length N containing the column pivoting information for the \(L U\) factorization. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{IPVT}, 1)\).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input)
Default: NZ \(=\operatorname{size}(A, 1)\).

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM.
Default: \(\operatorname{IPARAM}(1)=0\).
See Comment 3.
RPARAM - Parameter vector of length 5. (Input/Output)
See Comment 3.

\section*{FORTRAN 90 Interface}

Generic: CALL LFTXG (A, IROW, JCOL, NL, NFAC, FACT, IRFAC, JCFAC, IPVT, JPVT [,...])

Specific: The specific interface names are S_LFTXG and D_LFTXG.

\section*{FORTRAN 77 Interface}

Single: CALL LFTXG (N, NZ, A, IROW, JCOL, IPARAM, RPARAM, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT)

Double: The double precision name is DLFTXG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a \(n \times n\) sparse matrix. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column numbers for these entries in \(A\). That is
\[
A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero.
The routine LFTXG performs an \(L U\) factorization of the coefficient matrix \(A\). It by default uses a symmetric Markowitz strategy (Crowe et al. 1990) to choose pivots that most likely would reduce fillins while maintaining numerical stability. Different strategies are also provided as options for row oriented or column oriented problems. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively.

Finally, the solution \(x\) is obtained using LFSXG by the following calculations:
1) \(L z=P b\)
2) \(U y=z\)
3) \(x=Q y\)

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2TXG/DL2TXG. The reference is:

CALL L2TXG (N, NZ, A, IROW, JCOL, IPARAM, RPARAM, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, WK, LWK, IWK, LIWK)

The additional arguments are as follows:

WK - Real work vector of length LWK.
\(\mathbf{L W K}\) - The length of wK, Lwk should be at least MAXNz.

IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(15 \mathrm{~N}+4 *\) MAXNZ.
The workspace limit is determined by MAXNZ, where

MAXNZ \(=\) MINO (LWK, \(\operatorname{INT}(0.25(L I W K-15 N)))\)
2. Informational errors

Type Code
31 The coefficient matrix is numerically singular.
32 The growth factor is too large to continue.
3. If the default parameters are desired for LFTXG, then set IPARAM(1) to zero and call the routine LFTXG. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling LFTXG.

CALL L4LXG (IPARAM, RPARAM)

Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to L4LXG will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above.

The arguments are as follows:

IPARAM - Integer vector of length 6.
\[
\begin{aligned}
& \operatorname{IPARAM}(1)=\text { Initialization flag. } \\
& \operatorname{IPARAM}(2)=\text { The pivoting strategy. }
\end{aligned}
\]

\section*{IPARAM(2) Action}

1 Markowitz row search
2 Markowitz column search
3 Symmetric Markowitz search
Default: 3 .
\(\operatorname{IPARAM}(3)=\) The number of rows which have least numbers of nonzero
elements that will be searched for a pivotal element.
Default: 3 .
\(\operatorname{IPARAM}(4)=\) The maximal number of nonzero elements in \(A\) at any stage of the Gaussian elimination. (Output)
\(\operatorname{IPARAM}(5)=\) The workspace limit.

\section*{IPARAM(5) Action}

0
integer \(\quad\) This integer value replaces the default workspace limit.

When L2TXG is called, the values of LWK and LIWK are used instead of IPARAM(5).
\(\operatorname{IPARAM}(6)=\) Not used in LFTXG.
RPARAM - Real vector of length 5.
\(\operatorname{RPARAM}(1)=\) The upper limit on the growth factor. The computation stops when the growth factor exceeds the limit.

Default: 10 .
\(\operatorname{RPARAM}(2)=\) The stability factor. The absolute value of the pivotal element must be bigger than the largest element in absolute value in its row divided by RPARAM(2).

Default: 10.0.
\(\operatorname{RPARAM}(3)=\) Drop-tolerance. Any element in the lower triangular factor L will be removed if its absolute value becomes smaller than the drop-tolerance at any stage of the Gaussian elimination.

Default: 0.0.
\(\operatorname{RPARAM}(4)=\) The growth factor. It is calculated as the largest element in absolute value in A at any stage of the Gaussian elimination divided by the largest element in absolute value in the original A matrix. (Output) Large value of the growth factor indicates that an appreciable error in the computed solution is possible.
\(\operatorname{RPARAM}(5)=\) The value of the smallest pivotal element in absolute value. (Output)
If double precision is required, then DL4LXG is called and RPARAM is declared double precision.

\section*{Example}

As an example, consider the \(6 \times 6\) matrix of a linear system:
\[
A=\left[\begin{array}{cccccc}
10 & 0 & 0 & 0 & 0 & 0 \\
0 & 10 & -3 & -1 & 0 & 0 \\
0 & 0 & 15 & 0 & 0 & 0 \\
-2 & 0 & 0 & 10 & -1 & 0 \\
-1 & 0 & 0 & -5 & 1 & -3 \\
-1 & -2 & 0 & 0 & 0 & 6
\end{array}\right]
\]

The sparse coordinate form for \(A\) is given by:
\begin{tabular}{crrrrrrrrrrrrrrr} 
irow & 6 & 2 & 3 & 2 & 4 & 4 & 5 & 5 & 5 & 5 & 1 & 6 & 6 & 2 & 4 \\
jcol & 6 & 2 & 3 & 3 & 4 & 5 & 1 & 6 & 4 & 5 & 1 & 1 & 2 & 4 & 1 \\
a & 6 & 10 & 15 & -3 & 10 & -1 & -1 & -3 & -5 & 1 & 10 & -1 & -2 & -1 & -2
\end{tabular}
```

USE LFTXG_INT
USE WRRRN_INT
USE WRIRN_INT
INTEGER - N, NZ
PARAMETER (N=6, NZ=15)
INTEGER IROW(NZ), JCOL(NZ), NFAC, NL,\&
IRFAC(3*NZ), JCFAC(3*NZ), IPVT(N), JPVT(N)
REAL A(NZ), FACT(3*NZ)
DATA A/6., 10., 15., -3., 10., -1., -1., -3., -5., 1., 10., -1.,\&
-2., -1., -2./
DATA IROW/6, 2, 3, 2, 4, 4, 5, 5, 5, 5, 1, 6, 6, 2, 4/
DATA JCOL/6, 2, 3, 3, 4, 5, 1, 6, 4, 5, 1, 1, 2, 4, 1/
NFAC = 3*NZ
CALL LFTXG (A, IROW, JCOL, NL, NFAC, FACT, IRFAC, JCFAC, IPVT, JPVT)
CALL WRRRN (' fact ', FACT, 1, NFAC, 1)
CALL WRIRN (' irfac ', IRFAC, 1, NFAC, 1)

```
```

CALL WRIRN (' jcfac ', JCFAC, 1, NFAC, 1)
CALL WRIRN (' p ', IPVT, 1, N, 1)
CALL WRIRN (' q ', JPVT, 1, N, 1)
!
END

```

\section*{Output}
\begin{tabular}{rrrrrrrrr}
1 & & \multicolumn{7}{c}{ fact } \\
-0.10 & -5.00 & -0.20 & -0.10 & -0.10 & -1.00 & -0.20 & 4.90 & -5.10 \\
11 & 12 & 13 & 14 & 15 & 16 & & 10 \\
-1.00 & 30.00 & 6.00 & -2.00 & 10.00 & 15.00 & & & \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline & \multicolumn{15}{|c|}{irfac} \\
\hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 \\
\hline 3 & 4 & 4 & 5 & 5 & 6 & 6 & 6 & 5 & 5 & 4 & 4 & 3 & 3 & 2 & 1 \\
\hline & \multicolumn{15}{|c|}{jcfac} \\
\hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 \\
\hline 2 & 3 & 1 & 4 & 2 & 5 & 2 & 6 & 6 & 5 & 6 & 4 & 4 & 3 & 2 & 1 \\
\hline
\end{tabular}
\begin{tabular}{llllll} 
& & & \(P\) & & \\
1 & 2 & 3 & 4 & 5 & 6 \\
3 & 1 & 6 & 2 & 5 & 4
\end{tabular}
\begin{tabular}{llllll}
1 & 2 & 3 & 4 & 5 & 6 \\
3 & 1 & 2 & 6 & 5 & 4
\end{tabular}

\section*{LFSXG}

Solves a sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.

\section*{Required Arguments}

NFAC - The number of nonzero coefficients in FACT as output from subroutine LFTXG/DLFTXG. (Input)
\(N L\) - The number of nonzero coefficients in the triangular matrix \(L\) excluding the diagonal elements as output from subroutine LFTXG/DLFTXG. (Input)
\(\boldsymbol{F A C T}\) - Vector of length NFAC containing the nonzero elements of \(L\) (excluding the diagonals) in the first NL locations and the nonzero elements of \(U\) in NL +1 to NFAC locations as output from subroutine LFTXG/DLFTXG. (Input)

IRFAC - Vector of length NFAC containing the row numbers of the corresponding elements in FACT as output from subroutine LFTXG/DLFTXG. (Input)

JCFAC - Vector of length NFAC containing the column numbers of the corresponding elements in FACT as output from subroutine LFTXG/DLFTXG. (Input)

IPVT - Vector of length N containing the row pivoting information for the \(L U\) factorization as output from subroutine LFTXG/DLFTXG. (Input)
\(\boldsymbol{J P V T}\) - Vector of length N containing the column pivoting information for the \(L U\) factorization as output from subroutine LFTXG/DLFTXG. (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) — Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) — Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
IPATH - Path indicator. (Input)
IPATH \(=1\) means the system \(A x=B\) is solved.
IPATH \(=2\) means the system \(A^{T} x=B\) is solved.
Default: IPATH \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSXG (NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, B, X [,...])
Specific: The specific interface names are S_LFSXG and D_LFSXG.

\section*{FORTRAN 77 Interface}

Single: CALL LFSXG (N, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, B, IPATH, X)

Double: The double precision name is DLFSXG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a \(n \times n\) sparse matrix. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column numbers for these entries in \(A\). That is
\[
A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero. The routine LFSXG computes the solution of the linear equation given its \(L U\) factorization. The factorization is performed by calling LFTXG. The solution of the linear system is then found by the forward and backward substitution. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively. Finally, the solution \(x\) is obtained by the following calculations:
\[
\begin{aligned}
& \text { 1) } L z=P b \\
& \text { 2) } U y=z \\
& \text { 3) } x=Q y
\end{aligned}
\]

For more details, see Crowe et al. (1990).

\section*{Example}

As an example, consider the \(6 \times 6\) linear system:
\[
A=\left[\begin{array}{cccccc}
10 & 0 & 0 & 0 & 0 & 0 \\
0 & 10 & -3 & -1 & 0 & 0 \\
0 & 0 & 15 & 0 & 0 & 0 \\
-2 & 0 & 0 & 10 & -1 & 0 \\
-1 & 0 & 0 & -5 & 1 & -3 \\
-1 & -2 & 0 & 0 & 0 & 6
\end{array}\right]
\]

Let
\[
x_{1}^{T}=(1,2,3,4,5,6)
\]
so that \(A x_{1}=(10,7,45,33,-34,31)^{T}\), and
\[
x_{2}^{T}=(6,5,4,3,2,1)
\]
so that \(A x_{2}=(60,35,60,16,-22,10)^{T}\). The sparse coordinate form for \(A\) is given by:
\begin{tabular}{crrrrrrrrrrrrrrr} 
irow & 6 & 2 & 3 & 2 & 4 & 4 & 5 & 5 & 5 & 5 & 1 & 6 & 6 & 2 & 4 \\
jcol & 6 & 2 & 3 & 3 & 4 & 5 & 1 & 6 & 4 & 5 & 1 & 1 & 2 & 4 & 1 \\
a & 6 & 10 & 15 & -3 & 10 & -1 & -1 & -3 & -5 & 1 & 10 & -1 & -2 & -1 & -2
\end{tabular}
```

USE LFSXG INT
USE WRRRL-}\mp@subsup{}{}{-
USE LFTXG_INT
INTEGER - N, NZ
PARAMETER (N=6, NZ=15)
INTEGER IPATH, IROW(NZ), JCOL(NZ), NFAC, \&
NL, IRFAC(3*NZ), JCFAC (3*NZ), IPVT(N), JPVT(N)
REAL X(N), A(NZ), B (N,2), FACT(3*NZ)
CHARACTER TITLE(2)*2, RLABEL(1)*4, CLABEL(1)*6

```
```

    DATA RLABEL(1)/'NONE'/, CLABEL(1)/'NUMBER'/
    !
DATA A/6., 10., 15., -3., 10., -1., -1., -3., -5., 1., 10., -1.,\&
-2., -1., -2./
DATA B/10., 7., 45., 33., -34., 31.,\&
60., 35., 60., 16., -22., -10./
DATA IROW/6, 2, 3, 2, 4, 4, 5, 5, 5, 5, 1, 6, 6, 2, 4/
DATA JCOL/6, 2, 3, 3, 4, 5, 1, 6, 4, 5, 1, 1, 2, 4, 1/
DATA TITLE/'x1', 'x2'/
NFAC = 3*NZ
CALL LFTXG (A, IROW, JCOL, NL, NFAC, FACT, IRFAC, JCFAC, IPVT, JPVT)
DO 10 I = 1, 2
CALL LFSXG (NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, B(:,I), X)
CALL WRRRL (TITLE(I), X, RLABEL, CLABEL, 1, N, 1)
1 0 ~ C O N T I N U E
END

```

\section*{Output}
\begin{tabular}{rrrrrr}
1 & & & x1 & & \\
1.0 & 2.0 & 3.0 & 4 & 5 & 6 \\
& & & & 5.0 & 6.0 \\
1 & 2 & 3 & 4 & 5 & 6 \\
6.0 & 5.0 & 4.0 & 3.0 & 2.0 & 1.0
\end{tabular}

\section*{LSLZG}

Solves a complex sparse system of linear equations by Gaussian elimination.

\section*{Required Arguments}

A - Complex vector of length NZ containing the nonzero coefficients of the linear system. (Input)

IROW - Vector of length NZ containing the row numbers of the corresponding elements in A. (Input)

JCOL - Vector of length nZ containing the column numbers of the corresponding elements in A. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input)
Default: NZ \(=\operatorname{size}(\mathrm{A}, 1)\).

IPATH - Path indicator. (Input) IPATH \(=1\) means the system \(\mathrm{Ax}=\mathrm{b}\) is solved. IPATH \(=2\) means the system \(\mathrm{A}^{H} \mathrm{x}=\mathrm{b}\) is solved. Default: \(\operatorname{IPATH}=1\).

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 3. Default: IPARAM \(=0\).

RPARAM - Parameter vector of length 5. (Input/Output) See Comment 3

\section*{FORTRAN 90 Interface}

Generic: CALL LSLZG (A, IROW, JCOL, B, X [,...])
Specific: The specific interface names are S_LSLZG and D_LSLZG.

\section*{FORTRAN 77 Interface}

Single: CALL LSLZG (N, NZ, A, IROW, JCOL, B, IPATH, IPARAM, RPARAM, X)
Double: The double precision name is DLSLZG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a \(n \times n\) complex sparse matrix. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column numbers for these entries in \(A\). That is
\[
A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero.
The subroutine LSLZG solves a system of linear algebraic equations having a complex sparse coefficient matrix. It first uses the routine LFTZG to perform an \(L U\) factorization of the coefficient
matrix. The solution of the linear system is then found using LFSZG. The routine LFTZG by default uses a symmetric Markowitz strategy (Crowe et al. 1990) to choose pivots that most likely would reduce fill-ins while maintaining numerical stability. Different strategies are also provided as options for row oriented or column oriented problems. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively. Finally, the solution \(x\) is obtained by the following calculations:
\[
\begin{aligned}
& \text { 1) } L z=P b \\
& \text { 2) } U y=z \\
& \text { 3) } x=Q y
\end{aligned}
\]

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LZG/DL2LZG. The reference is:
```

CALL L2LZG (N,NZ, A, IROW, JCOL, B, IPATH, IPARAM, RPARAM, X, WK, LWK, IWK,
LIWK)

```

The additional arguments are as follows:

WK - Complex work vector of length LWK.
LWK - The length of wK, LWK should be at least 2N+ MAXNZ.
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(17 \mathrm{~N}+4\) * MAXNZ.
The workspace limit is determined by MAXNZ, where
```

MAXNZ = MINO(LWK-2N, INT (0.25(LIWK-17N)))

```
2. Informational errors

Type Code
31 The coefficient matrix is numerically singular.
32 The growth factor is too large to continue.
33 The matrix is too ill-conditioned for iterative refinement.
3. If the default parameters are desired for LSLZG, then set IPARAM(1) to zero and call the routine LSLZG. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM. then the following steps should be taken before calling LSLZG.
```

CALL L4LZG (IPARAM, RPARAM)

```

Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to L4LZG will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above. The arguments are as follows:

IPARAM - Integer vector of length 6.
\(\operatorname{IPARAM}(1)=\) Initialization flag.
\(\operatorname{IPARAM}(2)=\) The pivoting strategy.

IPARAM(2) Action
1 Markowitz row search
2 Markowitz column search
3 Symmetric Markowitz search
Default: 3 .
\(\operatorname{IPARAM}(3)=\) The number of rows which have least numbers of nonzero
elements that will be searched for a pivotal element.
Default: 3 .
\(\operatorname{IPARAM}(4)=\) The maximal number of nonzero elements in \(A\) at any stage of the Gaussian elimination. (Output)
\(\operatorname{IPARAM}(5)=\) The workspace limit.

IPARAM(5)
0 integer

Action Default limit, see Comment 1.

This integer value replaces the default workspace limit.

When L2LZG is called, the values of LWK and LIWK are used instead of IPARAM(5).

Default: 0 .
\(\operatorname{IPARAM}(6)=\) Iterative refinement is done when this is nonzero.
Default: 0.
RPARAM - Real vector of length 5.
\(\operatorname{RPARAM}(1)=\) The upper limit on the growth factor. The computation stops when the growth factor exceeds the limit.
Default: 10 .
\(\operatorname{RPARAM}(2)=\) The stability factor. The absolute value of the pivotal element must be bigger than the largest element in absolute value in its row divided by RPARAM(2).

Default: 10.0.
\(\operatorname{RPARAM}(3)=\) Drop-tolerance. Any element in A will be removed if its absolute value becomes smaller than the drop-tolerance at any stage of the Gaussian elimination.

Default: 0.0.
\(\operatorname{RPARAM}(4)=\) The growth factor. It is calculated as the largest element in absolute value in A at any stage of the Gaussian elimination divided by the largest element in absolute value in the original A matrix. (Output) Large value of the growth factor indicates that an appreciable error in the computed solution is possible.
\(\operatorname{RPARAM}(5)=\) The value of the smallest pivotal element in absolute value.
(Output)
If double precision is required, then DL4LZG is called and RPARAM is declared double precision.

\section*{Example}

As an example, consider the \(6 \times 6\) linear system:
\[
A=\left[\begin{array}{rrrrrr}
10+7 i & 0 & 0 & 0 & 0 & 0 \\
0 & 3+2 i & -3+0 i & -1+2 i & 0 & 0 \\
0 & 0 & 4+2 i & 0 & 0 & 0 \\
-2-4 i & 0 & 0 & 1+6 i & -1+3 i & 0 \\
-5+4 i & 0 & 0 & -5+0 i & 12+2 i & -7+7 i \\
-1+12 i & -2+8 i & 0 & 0 & 0 & 3+7 i
\end{array}\right]
\]

Let
\[
x^{T}=(1+i, 2+2 i, 3+3 i, 4+4 i, 5+5 i, 6+6 i)
\]
so that
\[
A x=(3+17 i,-19+5 i, 6+18 i,-38+32 i,-63+49 i,-57+83 i)^{T}
\]

The number of nonzeros in \(A\) is \(\mathrm{nz}=15\). The sparse coordinate form for \(A\) is given by:
\[
\begin{array}{clllllllllllllll}
\text { irow } & 6 & 2 & 2 & 4 & 3 & 1 & 5 & 4 & 6 & 5 & 5 & 6 & 4 & 2 & 5 \\
\text { jcol } & 6 & 2 & 3 & 5 & 3 & 1 & 1 & 4 & 1 & 4 & 5 & 2 & 1 & 4 & 6
\end{array}
\]
```

USE LSLZG_INT
USE WRCRN -
INTEGER N, NZ
PARAMETER (N=6, NZ=15)
INTEGER IROW (NZ), JCOL (NZ)
COMPLEX A(NZ), B(N), X(N)
DATA A/ (3.0,7.0), (3.0,2.0), (-3.0,0.0), (-1.0,3.0), (4.0,2.0),\&
(10.0,7.0), (-5.0,4.0), (1.0,6.0), (-1.0,12.0), (-5.0,0.0),\&
(12.0,2.0), (-2.0,8.0), (-2.0,-4.0), (-1.0,2.0), (-7.0,7.0)/

```
```

DATA B/ (3.0,17.0), (-19.0,5.0), (6.0,18.0), (-38.0,32.0),\&
(-63.0,49.0), (-57.0,83.0)/
DATA IROW/6, 2, 2, 4, 3, 1, 5, 4, 6, 5, 5, 6, 4, 2, 5/
DATA JCOL/6, 2, 3, 5, 3, 1, 1, 4, 1, 4, 5, 2, 1, 4, 6/
Use default options
CALL LSLZG (A, IROW, JCOL, B, X)
CALL WRCRN ('X', X)
END

```

\section*{Output}

\section*{X}
```

( $1.000,1.000$ )
( $2.000,2.000$ )
( $3.000,3.000$ )
( 4.000, 4.000)
(5.000, 5.000)
( $6.000,6.000$ )

```

\section*{LFTZG}

Computes the \(L U\) factorization of a complex general sparse matrix.

\section*{Required Arguments}

A-Complex vector of length NZ containing the nonzero coefficients of the linear system. (Input)

IROW - Vector of length NZ containing the row numbers of the corresponding elements in A. (Input)

JCOL - Vector of length nZ containing the column numbers of the corresponding elements in A. (Input)
\(N F A C\) - On input, the dimension of vector FACT. (Input/Output)
On output, the number of nonzero coefficients in the triangular matrix \(L\) and \(U\).
\(N L\) - The number of nonzero coefficients in the triangular matrix \(L\) excluding the diagonal elements. (Output)

FACT - Complex vector of length NFAC containing the nonzero elements of \(L\) (excluding the diagonals) in the first NL locations and the nonzero elements of \(U\) in NL +1 to NFAC locations. (Output)

IRFAC - Vector of length NFAC containing the row numbers of the corresponding elements in FACT. (Output)

JCFAC - Vector of length NFAC containing the column numbers of the corresponding elements in FACT. (Output)

IPVT - Vector of length N containing the row pivoting information for the \(L U\) factorization. (Output)
\(\boldsymbol{J P V T}\) - Vector of length N containing the column pivoting information for the \(L U\) factorization. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input) Default: \(\mathrm{N}=\operatorname{size}(\operatorname{IPVT}, 1)\).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input) Default: NZ \(=\operatorname{size}(A, 1)\).

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 3. Default: \(\operatorname{IPARAM}=0\).

RPARAM - Parameter vector of length 5. (Input/Output) See Comment 3.

\section*{FORTRAN 90 Interface}

Generic: CALL LFTZG (A, IROW, JCOL, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT [,...])

Specific: The specific interface names are S_LFTZG and D_LFTZG.

\section*{FORTRAN 77 Interface}

Single: CALL LFTZG (N, NZ, A, IROW, JCOL, IPARAM, RPARAM, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT)

Double: The double precision name is DLFTZG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a complex \(n \times n\) sparse matrix. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column indices for these entries in \(A\). That is
\[
A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero.
The routine LFTZG performs an \(L U\) factorization of the coefficient matrix \(A\). It uses by default a symmetric Markowitz strategy (Crowe et al. 1990) to choose pivots that most likely would reduce fill-ins while maintaining numerical stability. Different strategies are also provided as options for row oriented or column oriented problems. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively.

Finally, the solution \(x\) is obtained using LFSZG by the following calculations:
1) \(L z=P b\)
2) \(U y=z\)
3) \(x=Q y\)

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(L 2 T Z G / D L 2 T Z G\). The reference is:

CALL L2TZG (N, NZ, A, IROW, JCOL, IPARAM, RPARAM, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, WK, LWK, IWK, LIWK)

The additional arguments are as follows:
WK - Complex work vector of length LWK.
\(\mathbf{L W K}\) - The length of wK, LWK should be at least MAXNZ.
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(15 \mathrm{~N}+4\) * MAXNZ.
The workspace limit is determined by MAXNZ, where
```

MAXNZ = MINO(LWK,INT(0.25(LIWK-15N)))

```
2. Informational errors

Type Code
31 The coefficient matrix is numerically singular.
32 The growth factor is too large to continue.
3. If the default parameters are desired for LFTZG, then set IPARAM(1) to zero and call the routine LFTZG. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM. then the following steps should be taken before calling LFTZG:

CALL L4LZG (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to L4LZG will set IPARAM and RPARAM to their default values so only nondefault values need to be set above. The arguments are as follows:

IPARAM - Integer vector of length 6.
\(\operatorname{IPARAM}(1)=\) Initialization flag.
\(\operatorname{IPARAM}(2)=\) The pivoting strategy.

\section*{IPARAM(2) \\ Action}

1
Markowitz row search
2 Markowitz column search
3 Symmetric Markowitz search
\(\operatorname{IPARAM}(3)=\) The number of rows which have least numbers of nonzero elements that will be searched for a pivotal element.

Default: 3 .
\(\operatorname{IPARAM}(4)=\) The maximal number of nonzero elements in A at any stage of the Gaussian elimination. (Output)
\(\operatorname{IPARAM}(5)=\) The workspace limit.

IPARAM(5) Action
0
integer \(\quad\) This integer value replaces the default workspace limit. When L2tzG is called, the values of LWK and LIWK are used instead of IPARAM(5).

Default: 0 .
\(\operatorname{IPARAM}(6)=\) Not used in LFTZG.
RPARAM - Real vector of length 5.
\(\operatorname{RPARAM}(1)=\) The upper limit on the growth factor. The computation stops when the growth factor exceeds the limit.

Default: 10 .
\(\operatorname{RPARAM}(2)=\) The stability factor. The absolute value of the pivotal element must be bigger than the largest element in absolute value in its row divided by RPARAM(2).
Default: 10.0.
\(\operatorname{RPARAM}(3)=\) Drop-tolerance. Any element in the lower triangular factor \(L\) will be removed if its absolute value becomes smaller than the drop-tolerance at any stage of the Gaussian elimination.

Default: 0.0.
\(\operatorname{RPARAM}(4)=\) The growth factor. It is calculated as the largest element in absolute value in A at any stage of the Gaussian elimination divided by the largest element in absolute value in the original A matrix. (Output) Large value of the growth factor indicates that an appreciable error in the computed solution is possible.
\(\operatorname{RPARAM}(5)=\) The value of the smallest pivotal element in absolute value.
(Output)
If double precision is required, then DL4LZG is called and RPARAM is declared double precision.

\section*{Example}

As an example, the following \(6 \times 6\) matrix is factorized, and the outcome is printed:
\[
A=\left[\begin{array}{rrrrrr}
10+7 i & 0 & 0 & 0 & 0 & 0 \\
0 & 3+2 i & -3+0 i & -1+2 i & 0 & 0 \\
0 & 0 & 4+2 i & 0 & 0 & 0 \\
-2-4 i & 0 & 0 & 1+6 i & -1+3 i & 0 \\
-5+4 i & 0 & 0 & -5+0 i & 12+2 i & -7+7 i \\
-1+12 i & -2+8 i & 0 & 0 & 0 & 3+7 i
\end{array}\right]
\]

The sparse coordinate form for \(A\) is given by:
\begin{tabular}{clllllllllllllll} 
irow & 6 & 2 & 2 & 4 & 3 & 1 & 5 & 4 & 6 & 5 & 5 & 6 & 4 & 2 & 5 \\
jcol & 6 & 2 & 3 & 5 & 3 & 1 & 1 & 4 & 1 & 4 & 5 & 2 & 1 & 4 & 6
\end{tabular}
USE LFTZG INT
USE WRCRN \({ }^{-}\)INT
USE WRIRN INT
INTEGER - N, NFAC, NZ
PARAMETER \(\quad(\mathrm{N}=6, \mathrm{NZ}=15)\)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER IPVT (N), IRFAC(45), IROW(NZ), JCFAC(45), \&
    JCOL (NZ), JPVT (N), NL
COMPLEX A(NZ), FAC(45)
!
DATA A/ (3.0,7.0), (3.0,2.0), (-3.0,0.0), (-1.0,3.0), (4.0,2.0), \&
    \((10.0,7.0),(-5.0,4.0),(1.0,6.0),(-1.0,12.0),(-5.0,0.0), \&\)
    \((12.0,2.0),(-2.0,8.0),(-2.0,-4.0),(-1.0,2.0),(-7.0,7.0) /\)
DATA IROW/6, 2, 2, 4, 3, 1, 5, 4, 6, 5, 5, 6, 4, 2, 5/
DATA JCOL/6, 2, 3, 5, 3, 1, 1, 4, 1, 4, 5, 2, 1, 4, 6/
DATA NFAC/45/
Use default options
CALL LFTZG (A, IROW, JCOL, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT)
!
```

CALL WRCRN ('fact',FACT, 1, NFAC, 1)
CALL WRIRN (' irfac ',IRFAC, 1, NFAC, 1)
CALL WRIRN (' jcfac ',JCFAC, 1, NFAC, 1)
CALL WRIRN (' p ',IPVT, 1, N, 1)
CALL WRIRN (' q ',JPVT, 1, N, 1)
END

```
!

\section*{Output}
\begin{tabular}{rrr} 
& \multicolumn{2}{c}{ fact } \\
1 & \(\left(\begin{array}{rrr}0.50, & 0.85) \\
2 & ( & 0.15,\end{array}\right.\) & \(-0.41)\) \\
3 & \((-0.60\), & \(0.30)\) \\
4 & \((\) & 2.23,
\end{tabular}\(\left.-1.97\right)\)
\begin{tabular}{rrrrrrrrrrrrrrrr} 
irfac \\
1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 \\
3 & 4 & 4 & 5 & 5 & 6 & 6 & 6 & 5 & 5 & 4 & 4 & 3 & 3 & 2 & 1
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline & \multicolumn{15}{|c|}{jcfac} \\
\hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 \\
\hline 2 & 3 & 1 & 4 & 2 & 5 & 2 & 6 & 6 & 5 & 6 & 4 & 4 & 3 & 2 & 1 \\
\hline
\end{tabular}
\begin{tabular}{llllll} 
& & & \(p\) & & \\
1 & 2 & 3 & 4 & 5 & 6 \\
3 & 1 & 6 & 2 & 5 & 4
\end{tabular}
\begin{tabular}{llllll} 
& & & \(q\) & & \\
1 & 2 & 3 & 4 & 5 & 6 \\
3 & 1 & 2 & 6 & 5 & 4
\end{tabular}

\section*{LFSZG}

Solves a complex sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.

\section*{Required Arguments}

NFAC - The number of nonzero coefficients in FACT as output from subroutine LFTZG/DLFTZG. (Input)
\(N L\) - The number of nonzero coefficients in the triangular matrix \(L\) excluding the diagonal elements as output from subroutine LFTZG/DLFTZG. (Input)

FACT - Complex vector of length NFAC containing the nonzero elements of \(L\) (excluding the diagonals) in the first NL locations and the nonzero elements of \(U\) in NL +1 to NFAC locations as output from subroutine LFTZG/DLFTZG. (Input)

IRFAC - Vector of length NFAC containing the row numbers of the corresponding elements in FACT as output from subroutine LFTZG/DLFTZG. (Input)

JCFAC - Vector of length NFAC containing the column numbers of the corresponding elements in FACT as output from subroutine LFTZG/DLFTZG. (Input)

IPVT - Vector of length N containing the row pivoting information for the \(L U\) factorization as output from subroutine LFTZG/DLFTZG. (Input)
\(\boldsymbol{J P V T}\) - Vector of length N containing the column pivoting information for the \(L U\) factorization as output from subroutine LFTZG/DLFTZG. (Input)

B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
IPATH - Path indicator. (Input) IPATH \(=1\) means the system \(\mathrm{Ax}=\mathrm{b}\) is solved. IPATH \(=2\) means the system \(\mathrm{A}^{H} \mathrm{x}=\mathrm{b}\) is solved. Default: ІРАТН \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LFSZG (NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, B, X [, ...])
Specific: The specific interface names are S_LFSZG and D_LFSZG.

\section*{FORTRAN 77 Interface}

Single: CALL LFSZG (N, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT, B, IPATH, X)

Double: The double precision name is DLFSZG.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is a complex \(n \times n\) sparse matrix. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in \(A\). Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column numbers for these entries in \(A\). That is
\[
A_{\text {irow }(i), \text { icol }(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in \(A\) zero.
The routine LFSZG computes the solution of the linear equation given its \(L U\) factorization. The factorization is performed by calling LFTZG. The solution of the linear system is then found by the forward and backward substitution. The algorithm can be expressed as
\[
P A Q=L U
\]
where \(P\) and \(Q\) are the row and column permutation matrices determined by the Markowitz strategy (Duff et al. 1986), and \(L\) and \(U\) are lower and upper triangular matrices, respectively.

Finally, the solution \(x\) is obtained by the following calculations:
1) \(L z=P b\)
2) \(U y=z\)
3) \(x=Q y\)

For more details, see Crowe et al. (1990).

\section*{Example}

As an example, consider the \(6 \times 6\) linear system:
\[
A=\left[\begin{array}{rrrrrr}
10+7 i & 0 & 0 & 0 & 0 & 0 \\
0 & 3+2 i & -3+0 i & -1+2 i & 0 & 0 \\
0 & 0 & 4+2 i & 0 & 0 & 0 \\
-2-4 i & 0 & 0 & 1+6 i & -1+3 i & 0 \\
-5+4 i & 0 & 0 & -5+0 i & 12+2 i & -7+7 i \\
-1+12 i & -2+8 i & 0 & 0 & 0 & 3+7 i
\end{array}\right]
\]

Let
\[
x_{1}^{T}=(1+i, 2+2 i, 3+3 i, 4+4 i, 5+5 i, 6+6 i)
\]
so that
\[
A x_{1}=(3+17 i,-19+5 i, 6+18 i,-38+32 i,-63+49 i,-57+83 i)^{T}
\]
and
\[
x_{2}^{T}=(6+6 i, 5+5 i, 4+4 i, 3+3 i, 2+2 i, 1+i)
\]
so that
\[
A x_{2}=(18+102 i,-16+16 i, 8+24 i,-11-11 i,-63+7 i,-132+106 i)^{T}
\]

The sparse coordinate form for \(A\) is given by:
\begin{tabular}{clllllllllllllll} 
irow & 6 & 2 & 2 & 4 & 3 & 1 & 5 & 4 & 6 & 5 & 5 & 6 & 4 & 2 & 5 \\
jcol & 6 & 2 & 3 & 5 & 3 & 1 & 1 & 4 & 1 & 4 & 5 & 2 & 1 & 4 & 6
\end{tabular}

USE LFSZG INT
USE WRCRN_INT
USE LFTZG_INT
INTEGER - N, NZ
PARAMETER ( \(\mathrm{N}=6, \mathrm{NZ}=15\) )
INTEGER IPATH, IPVT(N), IRFAC(3*NZ), IROW(NZ), \&
JCFAC(3*NZ), JCOL(NZ), JPVT(N), NFAC, NL
COMPLEX A(NZ), B(N,2), FACT(3*NZ), X(N)
CHARACTER TITLE (2)*2
!
DATA A/ (3.0,7.0), (3.0,2.0), (-3.0,0.0), (-1.0,3.0), (4.0.2.0), \&
\((10.0,7.0),(-5.0,4.0),(1.0,6.0),(-1.0,12.0),(-5.0,0.0), \&\)
\((12.0,2.0),(-2.0,8.0),(-2.0,-4.0),(-1.0,2.0),(-7.0,7.0) /\)
DATA \(\mathrm{B} /(3.0,17.0),(-19.0,5.0),(6.0,18.0),(-38.0,32.0), \&\) \((-63.0,49.0),(-57.0,83.0),(18.0,102.0),(-16.0,16.0), \&\) \((8.0,24.0),(-11.0,-11.0),(-63.0,7.0),(-132.0,106.0) /\)
DATA IROW/6, 2, 2, 4, 3, 1, 5, 4, 6, 5, 5, 6, 4, 2, 5/
DATA JCOL/6, 2, 3, 5, 3, 1, 1, 4, 1, 4, 5, 2, 1, 4, 6/
DATA TITLE/'x1','x2'/
NFAC \(=3 * N Z\)
Perform LU factorization
CALL LFTZG (A, IROW, JCOL, NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT)
IPATH \(=1\)
DO 10 I = 1,2
CALL LFSZG (NFAC, NL, FACT, IRFAC, JCFAC, IPVT, JPVT,\& B(:,I), X)
CALL WRCRN (TITLE(I), X)
10 CONTINUE
END

\section*{Output}

\section*{x1}
```

( 1.000, 1.000)
( 2.000, 2.000)
( 3.000, 3.000)
(4.000, 4.000)

```
```

5 ( 5.000, 5.000)
6 (6.000, 6.000)
x2
(6.000, 6.000)
(5.000, 5.000)
(4.000, 4.000)
( 3.000, 3.000)
(2.000, 2.000)
(1.000, 1.000)

```

\section*{LSLXD}

Solves a sparse system of symmetric positive definite linear algebraic equations by Gaussian elimination.

\section*{Required Arguments}

A - Vector of length NZ containing the nonzero coefficients in the lower triangle of the linear system. (Input)
The sparse matrix has nonzeroes only in entries (IROW \((i)\), JCOL \((i)\) ) for \(i=1\) to NZ , and at this location the sparse matrix has value \(A(i)\).

IROW - Vector of length NZ containing the row numbers of the corresponding elements in the lower triangle of A. (Input) Note \(\operatorname{IROW}(i) \geq \operatorname{JCOL}(i)\), since we are only indexing the lower triangle.

JCOL - Vector of length NZ containing the column numbers of the corresponding elements in the lower triangle of \(A\). (Input)
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}-\) Vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
\(N Z\) - The number of nonzero coefficients in the lower triangle of the linear system. (Input) Default: NZ \(=\operatorname{size}(\mathrm{A}, 1)\).

ITWKSP — The total workspace needed. (Input) If the default is desired, set ITWKSP to zero. Default: \(I T W K S P=0\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLXD (A, IROW, JCOL, B, X [, ...])

Specific: The specific interface names are S_LSLXD and D_LSLXD.

\section*{FORTRAN 77 Interface}

Single: CALL LSLXD (N, NZ, A, IROW, JCOL, B, ITWKSP, X)
Double: \(\quad\) The double precision name is DLSLXD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and symmetric. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be nz . The two integer arrays irow and jcol, each of length nz , contain the row and column indices for these entries in \(A\). That is
\[
\begin{aligned}
& A_{\text {irow }(i), \operatorname{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
& \operatorname{irow}(i) \geq \operatorname{jcol}(i) \quad i=1, \ldots, \mathrm{nz}
\end{aligned}
\]
with all other entries in the lower triangle of \(A\) zero.
The subroutine LSLXD solves a system of linear algebraic equations having a real, sparse and positive definite coefficient matrix. It first uses the routine LSCXD to compute a symbolic factorization of a permutation of the coefficient matrix. It then calls LNFXD to perform the numerical factorization. The solution of the linear system is then found using LFSXD.
The routine LSCXD computes a minimum degree ordering or uses a user-supplied ordering to set up the sparse data structure for the Cholesky factor, \(L\). Then the routine LNFXD produces the numerical entries in \(L\) so that we have
\[
P A P^{T}=L L^{T}
\]

Here \(P\) is the permutation matrix determined by the ordering.
The numerical computations can be carried out in one of two ways. The first method performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

Finally, the solution \(x\) is obtained by the following calculations:
1) \(L y_{1}=P b\)
2) \(L^{T} y_{2}=y_{1}\)
3) \(x=P^{T} y_{2}\)

The routine LFSXD accepts \(b\) and the permutation vector which determines \(P\). It then returns \(x\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LXD/DL2LXD. The reference is:

CALL L2LXD (N, NZ, A, IROW, JCOL, B, X, IPER, IPARAM, RPARAM, WK, LWK, IWK, LIWK)

The additional arguments are as follows:
IPER - Vector of length N containing the ordering.
IPARAM - Integer vector of length 4. See Comment 3.
RPARAM - Real vector of length 2. See Comment 3.

WK - Real work vector of length LWK.
\(\boldsymbol{L W K}\) - The length of wK, LWK should be at least \(2 \mathrm{~N}+6 \mathrm{Nz}\).
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(15 \mathrm{~N}+15 \mathrm{NZ}+9\).
Note that the parameter ITWKSP is not an argument to this routine.
2. Informational errors

Type Code
41 The coefficient matrix is not positive definite.
42 A column without nonzero elements has been found in the coefficient matrix.
3. If the default parameters are desired for L2LXD, then set IPARAM(1) to zero and call the routine L2LXD. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling L2LXD.

CALL L4LXD (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to L4LXD will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above. The arguments are as follows:

IPARAM - Integer vector of length 4.
\(\operatorname{IPARAM}(1)=\) Initialization flag.
\(\operatorname{IPARAM}(2)=\) The numerical factorization method.

IPARAM(2)
0
1

\section*{Action}

Multifrontal
Markowitz column search

Default: 0.
\(\operatorname{IPARAM}(3)=\) The ordering option.

IPARAM(3)
0
1

Action
Minimum degree ordering
User's ordering specified in IPER

Default: 0.
\(\operatorname{IPARAM}(4)=\) The total number of nonzeros in the factorization matrix.
RPARAM - Real vector of length 2.
\(\operatorname{RPARAM}(1)=\) The value of the largest diagonal element in the Cholesky factorization.
\(\operatorname{RPARAM}(2)=\) The value of the smallest diagonal element in the Cholesky factorization.

If double precision is required, then DL4LXD is called and RPARAM is declared double precision.

\section*{Example}

As an example consider the \(5 \times 5\) linear system:
\[
A=\left[\begin{array}{rrrrr}
10 & 0 & 1 & 0 & 2 \\
0 & 20 & 0 & 0 & 3 \\
1 & 0 & 30 & 4 & 0 \\
0 & 0 & 4 & 40 & 5 \\
2 & 3 & 0 & 5 & 50
\end{array}\right]
\]

Let \(x^{T}=(1,2,3,4,5)\) so that \(A x=(23,55,107,197,278)^{T}\). The number of nonzeros in the lower triangle of \(A\) is \(\mathrm{nz}=10\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{crrrrrrrrrr} 
irow & 1 & 2 & 3 & 3 & 4 & 4 & 5 & 5 & 5 & 5 \\
jcol & 1 & 2 & 1 & 3 & 3 & 4 & 1 & 2 & 4 & 5 \\
a & 10 & 20 & 1 & 30 & 4 & 40 & 2 & 3 & 5 & 50
\end{tabular}
or equivalently by
\begin{tabular}{crrrrrrrrrr} 
irow & 4 & 5 & 5 & 5 & 1 & 2 & 3 & 3 & 4 & 5 \\
jcol & 4 & 1 & 2 & 4 & 1 & 2 & 1 & 3 & 3 & 5 \\
a & 40 & 2 & 3 & 5 & 10 & 20 & 1 & 30 & 4 & 50
\end{tabular}
```

USE LSLXD INT
USE WRRRN }\mp@subsup{}{}{-}\mathrm{ INT
INTEGER N, NZ
PARAMETER (N=5, NZ=10)
INTEGER IROW (NZ), JCOL (NZ)
REAL A(NZ), B(N), X(N)
DATA A/10., 20., 1., 30., 4., 40., 2., 3., 5., 50./
DATA B/23., 55., 107., 197., 278./
DATA IROW/1, 2, 3, 3, 4, 4, 5, 5, 5, 5/
DATA JCOL/1, 2, 1, 3, 3, 4, 1, 2, 4, 5/
Solve A * X = B
CALL LSLXD (A, IROW, JCOL, B, X)
CALL WRRRN (' x ', X, 1, N, 1)
END

```

\section*{Output}
\begin{tabular}{rrrrr}
1 & \multicolumn{1}{l}{} \\
1.000 & \(2.000^{X}\) & \(3.000^{4}\) & 4.000 & 5.000
\end{tabular}

\section*{LSCXD}

Performs the symbolic Cholesky factorization for a sparse symmetric matrix using a minimum degree ordering or a user-specified ordering, and set up the data structure for the numerical Cholesky factorization

\section*{Required Arguments}

IROW - Vector of length NZ containing the row subscripts of the nonzeros in the lower triangular part of the matrix including the nonzeros on the diagonal. (Input)

JCOL - Vector of length NZ containing the column subscripts of the nonzeros in the lower triangular part of the matrix including the nonzeros on the diagonal. (Input) (IROW (K), JCOL(K)) gives the row and column indices of the \(k\)-th nonzero element of the matrix stored in coordinate form. Note, \(\operatorname{IROW}(\mathrm{K}) \geq \operatorname{JCOL}(\mathrm{K})\).

NZSUB - Vector of length MAXSUB containing the row subscripts for the off-diagonal nonzeros in the Cholesky factor in compressed format. (Output)

INZSUB - Vector of length \(N+1\) containing pointers for NZSUB. The row subscripts for the off-diagonal nonzeros in column J are stored in NZSUB from location INZSUB ( J ) to \(\operatorname{INZSUB}(J+(\operatorname{ILNZ}(J+1)-\operatorname{ILNZ}(J)-1) . \quad\) (Output)
\(M A X N Z\) - Total number of off-diagonal nonzeros in the Cholesky factor. (Output)
\(\boldsymbol{I L N Z}\) - Vector of length \(\mathrm{N}+1\) containing pointers to the Cholesky factor. The off-diagonal nonzeros in column \(J\) of the factor are stored from location ILNZ ( \(J\) ) to \(\operatorname{ILNZ}(J+1)-1\). (Output)
(ILNZ, NZSUB, INZSUB) sets up the data structure for the off-diagonal nonzeros of the Cholesky factor in column ordered form using compressed subscript format.

INVPER - Vector of length N containing the inverse permutation. (Output) \(\operatorname{INVPER}(\mathrm{K})=\) I indicates that the original row K is the new row I.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input) Default: \(\mathrm{N}=\) size (INVPER,1).
\(N Z\) - Total number of the nonzeros in the lower triangular part of the symmetric matrix, including the nonzeros on the diagonal. (Input)
Default: NZ = size (IROW,1).
IJOB - Integer parameter selecting an ordering to permute the matrix symmetrically. (Input) IJOB \(=0\) selects the user ordering specified in IPER and reorders it so that the multifrontal method can be used in the numerical factorization.
IJOB \(=1\) selects the user ordering specified in IPER.
\(I J O B=2\) selects a minimum degree ordering.
\(I J O B=3\) selects a minimum degree ordering suitable for the multifrontal method in the numerical factorization.
Default: \(\mathrm{IJOB}=3\).

ITWKSP - The total workspace needed. (Input) If the default is desired, set ITWKSP to zero. Default: \(I T W K S P=0\).

MAXSUB - Number of subscripts contained in array NZSUB. (Input/Output) On input, MAXSUB gives the size of the array NZSUB. Note that when default workspace ( \(\operatorname{ITWKSP}=0\) ) is used, set MAXSUB \(=3 * \mathrm{NZ}\). Otherwise ( \(\mathrm{ITWKSP}>0\) ), set MAXSUB \(=(\mathrm{ITWKSP}-10 * N-7) / 4\). On output, MAXSUB gives the number of subscripts used by the compressed subscript format.
Default: \(\mathrm{MAXSUB}=3^{*} \mathrm{NZ}\).
\(\boldsymbol{I P E R}\) - Vector of length \(N\) containing the ordering specified by IJOB. (Input/Output)
\(\operatorname{IPER}(I)=K\) indicates that the original row \(K\) is the new row \(I\).

ISPACE - The storage space needed for stack of frontal matrices. (Output)

\section*{FORTRAN 90 Interface}

Generic: Because the Fortran compiler cannot determine the precision desired from the required arguments, there is no generic Fortran 90 Interface for this routine. The specific Fortran 90 Interfaces are:

Single: CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER [,...])
Or

CALL S_LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER [, ...])

Double: CALL DLSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER [, ...])
Or

CALL D_LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER [, ...])

\section*{FORTRAN 77 Interface}

Single: CALL LSCXD (N, NZ, IROW, JCOL, IJOB, ITWKSP, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ, IPER, INVPER, ISPACE)

Double: \(\quad\) The double precision name is DLSCXD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and symmetric. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be nz . The two integer arrays irow and jcol, each of length nz , contain the row and column indices for these entries in \(A\). That is
\[
\begin{gathered}
A_{\text {irow }(i), \operatorname{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
\operatorname{irow}(i) \geq \operatorname{jcol}(i)
\end{gathered} \quad i=1, \ldots, \mathrm{nz} \text {. }
\]
with all other entries in the lower triangle of \(A\) zero.
The routine LSCXD computes a minimum degree ordering or uses a user-supplied ordering to set up the sparse data structure for the Cholesky factor, \(L\). Then the routine LNFXD produces the numerical entries in \(L\) so that we have
\[
P A P^{T}=L L^{T}
\]

Here, \(P\) is the permutation matrix determined by the ordering.
The numerical computations can be carried out in one of two ways. The first method performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2CXD. The reference is:
```

CALL L2CXD (N, NZ, IROW, JCOL, IJOB, MAXSUB, NZSUB, INZSUB,
MAXNZ, ILNZ, IPER, INVPER, ISPACE, LIWK,IWK)

```

The additional arguments are as follows:
LIWK - The length of IWK, LIWK should be at least \(10 \mathrm{~N}+12 \mathrm{NZ}+7\). Note that the argument MAXSUB should be set to (LIWK - 10N -7 )/4.

IWK - Integer work vector of length LIWK.
Note that the parameter ITWKSP is not an argument to this routine.
2. Informational errors

Type Code
\(4 \quad 1 \quad\) The matrix is structurally singular.

\section*{Example}

As an example, the following matrix is symbolically factorized, and the result is printed:
\[
A=\left[\begin{array}{rrrrr}
10 & 0 & 1 & 0 & 2 \\
0 & 20 & 0 & 0 & 3 \\
1 & 0 & 30 & 4 & 0 \\
0 & 0 & 4 & 40 & 5 \\
2 & 3 & 0 & 5 & 50
\end{array}\right]
\]

The number of nonzeros in the lower triangle of \(A\) is \(n z=10\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{lllllllllll} 
irow & 1 & 2 & 3 & 3 & 4 & 4 & 5 & 5 & 5 & 5 \\
jcol & 1 & 2 & 1 & 3 & 3 & 4 & 1 & 2 & 4 & 5
\end{tabular}
or equivalently by
\begin{tabular}{lllllllllll} 
irow & 4 & 5 & 5 & 5 & 1 & 2 & 3 & 3 & 4 & 5 \\
jcol & 4 & 1 & 2 & 4 & 1 & 2 & 1 & 3 & 3 & 5
\end{tabular}

USE LSCXD_INT
USE WRIRN_INT
INTEGER N, NZ
PARAMETER \(\quad(\mathrm{N}=5, \quad \mathrm{NZ}=10)\)

INTEGER ILNZ \((N+1), \operatorname{INVPER}(N), \operatorname{INZSUB}(N+1), \operatorname{IPER}(N), \&\) IROW (NZ), ISPACE, JCOL (NZ), MAXNZ, MAXSUB, \& NZSUB ( \(3 *\) NZ)

DATA IROW/1, 2, 3, 3, 4, 4, 5, 5, 5, 5/
DATA JCOL/1, 2, 1, 3, 3, 4, 1, 2, 4, 5/
MAXSUB \(=3 * \mathrm{NZ}\)
CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER, \& MAXSUB=MAXSUB, IPER=IPER) Print results
CALL WRIRN (' iper ', IPER, 1, N, 1)
CALL WRIRN (' invper ', INVPER, 1, N, 1)
CALL WRIRN (' nzsub ', NZSUB, 1, MAXSUB, 1)
CALL WRIRN (' inzsub ', INZSUB, 1, N+1, 1)
CALL WRIRN (' ilnz ', ILNZ, 1, N+1, 1)
END

\section*{Output}
\begin{tabular}{ccccc} 
& \multicolumn{5}{c}{ iper } \\
1 & 2 & 3 & 4 & 5 \\
2 & 1 & 5 & 4 & 3 \\
& & \multicolumn{4}{c}{ invper } \\
& & \\
1 & 2 & 3 & 4 & 5 \\
2 & 1 & 5 & 4 & 3
\end{tabular}
nzsub
\[
\begin{array}{lll}
2 & 3
\end{array}
\]
\[
\begin{array}{llll}
3 & 5 & 4 & 5
\end{array}
\]
\begin{tabular}{cccccc} 
& \multicolumn{5}{c}{ inzsub } \\
1 & 2 & 3 & 4 & 5 & 6 \\
1 & 1 & 3 & 4 & 4 & 4
\end{tabular}
\begin{tabular}{ccccc} 
& \multicolumn{5}{c}{ ilnz } \\
1 & 2 & 3 & 4 & 5 \\
1 & 2 & 4 & 6 & 7 \\
\hline
\end{tabular}

\section*{LNFXD}

Computes the numerical Cholesky factorization of a sparse symmetrical matrix \(A\).

\section*{Required Arguments}

A - Vector of length NZ containing the nonzero coefficients of the lower triangle of the linear system. (Input)

IROW - Vector of length NZ containing the row numbers of the corresponding elements in the lower triangle of A. (Input)

JCOL - Vector of length nZ containing the column numbers of the corresponding elements in the lower triangle of A. (Input)

MAXSUB - Number of subscripts contained in array NZSUB as output from subroutine LSCXD/DLSCXD. (Input)

NZSUB - Vector of length MAXSUB containing the row subscripts for the nonzeros in the Cholesky factor in compressed format as output from subroutine LSCXD/DLSCXD. (Input)

INZSUB - Vector of length \(\mathrm{N}+1\) containing pointers for NZSUB as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts for the nonzeros in column \(J\) are stored from location INZSUB ( J ) to \(\operatorname{INZSUB}(J+1)-1\).

MAXNZ — Length of RLNZ as output from subroutine LSCXD/DLSCXD. (Input)
\(\boldsymbol{I L N Z}\) - Vector of length \(\mathrm{N}+1\) containing pointers to the Cholesky factor as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts for the nonzeros in column \(J\) of the factor are stored from location \(\operatorname{ILNZ}(J)\) to \(\operatorname{ILNZ}(J+1)-1\). (ILNZ, NZSUB, INZSUB) sets up the compressed data structure in column ordered form for the Cholesky factor.

IPER - Vector of length N containing the permutation as output from subroutine LSCXD/DLSCXD. (Input)

INVPER - Vector of length N containing the inverse permutation as output from subroutine LSCXD/DLSCXD. (Input)

ISPACE - The storage space needed for the stack of frontal matrices as output from subroutine LSCXD/DLSCXD. (Input)

DIAGNL - Vector of length N containing the diagonal of the factor. (Output)

RLNZ - Vector of length MAXNZ containing the strictly lower triangle nonzeros of the Cholesky factor. (Output)

RPARAM - Parameter vector containing factorization information. (Output)
\(\operatorname{RPARAM}(1)=\) smallest diagonal element.
\(\operatorname{RPARAM}(2)=\) largest diagonal element.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{IPER}, 1)\).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input)
Default: NZ \(=\operatorname{size}(A, 1)\).

IJOB - Integer parameter selecting factorization method. (Input)
IJOB \(=1\) yields factorization in sparse column format.
\(I J O B=2\) yields factorization using multifrontal method.
Default: IJOB \(=1\).

ITWKSP - The total workspace needed. (Input) If the default is desired, set ITWKSP to zero. Default: \(I T W K S P=0\).

\section*{FORTRAN 90 Interface}

Generic: CALL LNFXD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ, IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM [,...])

Specific: The specific interface names are S_LNFXD and D_LNFXD.

\section*{FORTRAN 77 Interface}

Single: CALL LNFXD (N, NZ, A, IROW, JCOL, IJOB, ITWKSP, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ, IPER, INVPER, ISPACE, ITWKSP, DIAGNL, RLNZ, RPARAM)

Double: The double precision name is DLNFXD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and symmetric. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be nz . The two integer arrays irow and jcol, each of length nz , contain the row and column indices for these entries in \(A\). That is
\[
\begin{aligned}
& A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
& \text { irow }(i) \geq \operatorname{jcol}(i) \quad i=1, \ldots, n z
\end{aligned}
\]
with all other entries in the lower triangle of \(A\) zero. The routine LNFXD produces the Cholesky factorization of \(P A P^{T}\) given the symbolic factorization of \(A\) which is computed by LSCXD. That is, this routine computes \(L\) which satisfies
\[
P A P^{T}=L L^{T}
\]

The diagonal of \(L\) is stored in DIAGNL and the strictly lower triangular part of \(L\) is stored in compressed subscript form in \(R=\) RLNZ as follows. The nonzeros in the \(j\)-th column of \(L\) are stored in locations \(R(i), \ldots, R(i+k)\) where \(i=\operatorname{ILNZ}(j)\) and \(k=\operatorname{ILNZ}(j+1)-\operatorname{ILNZ}(j)-1\). The row subscripts are stored in the vector NZSUB from locations INZSUB \((j)\) to \(\operatorname{INZSUB}(j)+k\).

The numerical computations can be carried out in one of two ways. The first method (when \(I J O B=2\) ) performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method (when IJOB \(=1\) ) is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

\section*{Comments}
1. Workspace may be explicitly provided by use of L2FXD/DL2FXD. The reference is:
```

CALL L2FXD (N, NZ, A, IROW, JCOL, IJOB, MAXSUB, NZSUB, INZSUB,
MAXNZ, ILNZ, IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM, WK,
LWK, IWK,LIWK)

```

The additional arguments are as follows:
WK - Real work vector of length LWK.
\(\boldsymbol{L W K}\) - The length of \(W K, L W K\) should be at least \(\mathrm{N}+3 \mathrm{Nz}\).
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least 2 N .

Note that the parameter ITWKSP is not an argument to this routine.
2. Informational errors

Type Code
41 The coefficient matrix is not positive definite.
42 A column without nonzero elements has been found in the coefficient matrix.

\section*{Example}

As an example, consider the \(5 \times 5\) linear system:
\[
A=\left[\begin{array}{rrrrr}
10 & 0 & 1 & 0 & 2 \\
0 & 20 & 0 & 0 & 3 \\
1 & 0 & 30 & 4 & 0 \\
0 & 0 & 4 & 40 & 5 \\
2 & 3 & 0 & 5 & 50
\end{array}\right]
\]

The number of nonzeros in the lower triangle of \(A\) is \(n z=10\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{crrrrrrrrrr} 
irow & 1 & 2 & 3 & 3 & 4 & 4 & 5 & 5 & 5 & 5 \\
jcol & 1 & 2 & 1 & 3 & 3 & 4 & 1 & 2 & 4 & 5 \\
a & 10 & 20 & 1 & 30 & 4 & 40 & 2 & 3 & 5 & 50
\end{tabular}
or equivalently by
\begin{tabular}{crrrrrrrrrr} 
irow & 4 & 5 & 5 & 5 & 1 & 2 & 3 & 3 & 4 & 5 \\
jcol & 4 & 1 & 2 & 4 & 1 & 2 & 1 & 3 & 3 & 5 \\
a & 40 & 2 & 3 & 5 & 10 & 20 & 1 & 30 & 4 & 50
\end{tabular}

We first call LSCXD to produce the symbolic information needed to pass on to LNFXD. Then call LNFXD to factor this matrix. The results are displayed below.
```

    USE LNFXD INT
    USE LSCXD_INT
    USE WRRRN_INT
    INTEGER - N, NZ, NRLNZ
    PARAMETER ( }\textrm{N}=5,NZ=10, NRLNZ=10
    INTEGER IJOB, ILNZ(N+1), INVPER(N), INZSUB(N+1), IPER(N),&
        IROW(NZ), ISPACE, JCOL(NZ), MAXNZ, MAXSUB,&
        NZSUB(3*NZ)
    REAL A(NZ), DIAGNL(N), RLNZ(NRLNZ), RPARAM(2) , R(N,N)
    DATA A/10., 20., 1., 30., 4., 40., 2., 3., 5., 50./
    DATA IROW/1, 2, 3, 3, 4, 4, 5, 5, 5, 5/
    DATA JCOL/1, 2, 1, 3, 3, 4, 1, 2, 4, 5/
                Select minimum degree ordering
                for multifrontal method
    IJOB = 3
    MAXSUB = 3*NZ
    CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER, &
                MAXSUB=MAXSUB)
                        Check if NRLNZ is large enough
    IF (NRLNZ .GE. MAXNZ) THEN
                        Choose multifrontal method
        IJOB = 2
        CALL LNFXD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB, MAXNZ, &
                        ILNZ,IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM, &
                IJOB=IJOB)
                        Print results
    ```
```

            CALL WRRRN (' diagnl ', DIAGNL, NRA=1, NCA=N, LDA=1)
            CALL WRRRN (' rlnz ', RLNZ, NRA= 1, NCA= MAXNZ, LDA= 1)
        END IF
    !
Construct L matrix
DO I=1,N
R(I,I) = DIAG(I)
IF (ILNZ(I) .GT. MAXNZ) GO TO 50
Find elements of RLNZ for this column
ISTRT = ILNZ(I)
ISTOP = ILNZ(I+1) - 1
! Get starting index for NZSUB
K = INZSUB(I)
DO J=ISTRT, ISTOP
! NZSUB(K) is the row for this element of
RLNZ
R((NZSUB(K)),I) = RLNZ(J)
K = K + 1
END DO
END DO
50 CONTINUE
CALL WRRRN ('L', R, NRA=N, NCA=N)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|c|}{diagnl} \\
\hline 1 & 2 & 3 & 4 & 5 & & \\
\hline 4.472 & 3.162 & 7.011 & 6.284 & 5.430 & & \\
\hline \multicolumn{7}{|c|}{\(r \operatorname{lnz}\)} \\
\hline 1 & & 2 & 3 & 4 & 5 & 6 \\
\hline 0.6708 & 0.6 & 50 & 162 & 0.7132 & -0.0285 & 0.6398 \\
\hline
\end{tabular}
\begin{tabular}{rrrrrr} 
& & & \multicolumn{3}{c}{\(L\)} \\
& 1 & 2 & 3 & 4 & 5 \\
1 & 4.472 & 0.000 & 0.000 & 0.000 & 0.000 \\
2 & 0.000 & 3.162 & 0.000 & 0.000 & 0.000 \\
3 & 0.671 & 0.632 & 7.011 & 0.000 & 0.000 \\
4 & 0.000 & 0.000 & 0.713 & 6.284 & 0.000 \\
5 & 0.000 & 0.316 & -0.029 & 0.640 & 5.430
\end{tabular}

\section*{LFSXD}

Solves a real sparse symmetric positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.

\section*{Required Arguments}
\(N\) - Number of equations. (Input)

MAXSUB - Number of subscripts contained in array NZSUB as output from subroutine LSCXD/DLSCXD. (Input)

NZSUB - Vector of length MAXSUB containing the row subscripts for the off-diagonal nonzeros in the factor as output from subroutine LSCXD/DLSCXD. (Input)

INZSUB - Vector of length \(N+1\) containing pointers for NZSUB as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts of column \(J\) are stored from location \(\operatorname{INZSUB}(J)\) to \(\operatorname{INZSUB}(J+1)-1\).

MAXNZ - Total number of off-diagonal nonzeros in the Cholesky factor as output from subroutine LSCXD/DLSCXD. (Input)
\(\boldsymbol{R} \boldsymbol{L N Z}\) - Vector of length MAXNZ containing the off-diagonal nonzeros in the factor in column ordered format as output from subroutine LNFXD/DLNFXD. (Input)
\(\boldsymbol{I L N Z}\) - Vector of length \(N+1\) containing pointers to RLNZ as output from subroutine LSCXD/DLSCXD. The nonzeros in column \(J\) of the factor are stored from location \(\operatorname{ILNZ}(\mathrm{J})\) to \(\operatorname{ILNZ}(\mathrm{J}+1)-1\). (Input)
The values (RLNZ, ILNZ, NZSUB, INZSUB) give the off-diagonal nonzeros of the factor in a compressed subscript data format.

DIAGNL - Vector of length \(N\) containing the diagonals of the Cholesky factor as output from subroutine LNFXD/DLNFXD. (Input)

IPER - Vector of length \(N\) containing the ordering as output from subroutine LSCXD/DLSCXD. (Input) \(\operatorname{IPER}(I)=K\) indicates that the original row \(K\) is the new row \(I\).
\(\boldsymbol{B}\) - Vector of length N containing the right-hand side. (Input)
\(\boldsymbol{X}\) - Vector of length N containing the solution. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL LFSXD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL, IPER, B, X)

Specific: The specific interface names are S_LFSXD and D_LFSXD.

\section*{FORTRAN 77 Interface}

Single: CALL LFSXD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL, IPER, B, X)

Double: \(\quad\) The double precision name is DLFSXD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and symmetric. The sparse coordinate format for the matrix \(A\) requires one real and two integer vectors. The real array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column indices for these entries in \(A\). That is
\[
\begin{gathered}
A_{\text {irow }(i), \text { icol }(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
\operatorname{irow}(i) \geq \mathrm{jcol}(i)
\end{gathered}
\]
with all other entries in the lower triangle of \(A\) zero.
The routine LFSXD computes the solution of the linear system given its Cholesky factorization. The factorization is performed by calling LSCXD followed by LNFXD. The routine LSCXD computes a minimum degree ordering or uses a user-supplied ordering to set up the sparse data structure for the Cholesky factor, \(L\). Then the routine LNFXD produces the numerical entries in \(L\) so that we have
\[
P A P^{T}=L L^{T}
\]

Here \(P\) is the permutation matrix determined by the ordering.
The numerical computations can be carried out in one of two ways. The first method performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu \((1986,1989)\). The second method is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

Finally, the solution \(x\) is obtained by the following calculations:
\[
\begin{aligned}
& \text { 1) } L y_{1}=P b \\
& \text { 2) } L^{T} y_{2}=y_{1} \\
& \text { 3) } x=P^{T} y_{2}
\end{aligned}
\]

\section*{Comments}

Informational error
\begin{tabular}{cc} 
Type & Code \\
4 & 1 The input matrix is numerically singular.
\end{tabular}

\section*{Example}

As an example, consider the \(5 \times 5\) linear system:
\[
A=\left[\begin{array}{rrrrr}
10 & 0 & 1 & 0 & 2 \\
0 & 20 & 0 & 0 & 3 \\
1 & 0 & 30 & 4 & 0 \\
0 & 0 & 4 & 40 & 5 \\
2 & 3 & 0 & 5 & 50
\end{array}\right]
\]

Let
\[
x_{1}^{T}=(1,2,3,4,5)
\]
so that \(A x_{1}=(23,55,107,197,278)^{T}\), and
\[
x_{2}^{T}=(5,4,3,2,1)
\]
so that \(A x_{2}=(55,83,103,97,82)^{T}\). The number of nonzeros in the lower triangle of \(A\) is \(\mathrm{nz}=10\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{crrrrrrrrrr} 
irow & 1 & 2 & 3 & 3 & 4 & 4 & 5 & 5 & 5 & 5 \\
jcol & 1 & 2 & 1 & 3 & 3 & 4 & 1 & 2 & 4 & 5 \\
a & 10 & 20 & 1 & 30 & 4 & 40 & 2 & 3 & 5 & 50
\end{tabular}
or equivalently by
\begin{tabular}{crrrrrrrrrr} 
irow & 4 & 5 & 5 & 5 & 1 & 2 & 3 & 3 & 4 & 5 \\
jcol & 4 & 1 & 2 & 4 & 1 & 2 & 1 & 3 & 3 & 5 \\
a & 40 & 2 & 3 & 5 & 10 & 20 & 1 & 30 & 4 & 50
\end{tabular}
```

USE LFSXD_INT
USE LNFXD_INT
USE LSCXD -INT
USE WRRRN_INT
INTEGER N, NZ, NRLNZ
PARAMETER (N=5, NZ=10, NRLNZ=10)
INTEGER IJOB, ILNZ(N+1), INVPER(N), INZSUB(N+1), IPER(N),\&
IROW(NZ), ISPACE, ITWKSP, JCOL(NZ), MAXNZ, MAXSUB,\&
NZSUB(3*NZ)
REAL A(NZ), B1(N), B2(N), DIAGNL(N), RLNZ(NRLNZ), RPARAM(2),\&
X(N)
DATA A/10., 20., 1., 30., 4., 40., 2., 3., 5., 50./
DATA B1/23., 55., 107., 197., 278./
DATA B2/55., 83., 103., 97., 82./
DATA IROW/1, 2, 3, 3, 4, 4, 5, 5, 5, 5/
DATA JCOL/1, 2, 1, 3, 3, 4, 1, 2, 4, 5/
Select minimum degree ordering
for multifrontal method
IJOB = 3
ITWKSP = 0
MAXSUB = 3*NZ

```
```

    CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER, &
                MAXSUB=MAXSUB, IPER=IPER, ISPACE=ISPACE)
                            Check if NRLNZ is large enough
    IF (NRLNZ .GE. MAXNZ) THEN
        IJOB = 2
    CALL LNFXD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ,&
            IPER, INVPER,ISPACE, DIAGNL, RLNZ, RPARAM, IJOB=IJOB)
                                    Solve A * X1 = B1
    CALL LFSXD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL,&
            IPER, B1, X)
            Print X1
    CALL WRRRN (' x1 ', X, 1, N, 1)
                            Solve A * X2 = B2
    CALL LFSXD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, &
                        DIAGNL, IPER, B2, X)
                            Print X2
    CALL WRRRN (' x2 ' X, 1, N, 1)
    END IF
    !
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{x1} \\
\hline 1 & 2 & 3 & 4 & 5 \\
\hline 1.000 & 2.000 & 3.000 & 4.000 & 5.000 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{\(x 2\)} \\
\hline 1 & 2 & 3 & 4 & 5 \\
\hline 5.000 & 4.000 & 3.000 & 2.000 & 1.000 \\
\hline
\end{tabular}

\section*{LSLZD}

Solves a complex sparse Hermitian positive definite system of linear equations by Gaussian elimination.

\section*{Required Arguments}

A-Complex vector of length NZ containing the nonzero coefficients in the lower triangle of the linear system. (Input)
The sparse matrix has nonzeroes only in entries (IROW \((i)\), \(\operatorname{JCOL}(i)\) ) for \(i=1\) to NZ , and at this location the sparse matrix has value \(A(i)\).

IROW - Vector of length NZ containing the row numbers of the corresponding elements in the lower triangle of \(A\). (Input)
Note \(\operatorname{IROW}(i) \geq \operatorname{JCOL}(i)\), since we are only indexing the lower triangle.
\(\boldsymbol{J C O L}\) - Vector of length NZ containing the column numbers of the corresponding elements in the lower triangle of A. (Input)

B - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution to the linear system. (Output)

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{B}, 1)\).
\(N Z\) - The number of nonzero coefficients in the lower triangle of the linear system. (Input) Default: NZ \(=\operatorname{size}(\mathrm{A}, 1)\).

ITWKSP - The total workspace needed. (Input) If the default is desired, set ITWKSP to zero. Default: \(I T W K S P=0\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLZD (A, IROW, JCOL, B, X [, ...])
Specific: The specific interface names are S_LSLZD and D_LSLZD.

\section*{FORTRAN 77 Interface}

Single: CALL LSLZD (N, NZ, A, IROW, JCOL, B, ITWKSP, X)

Double: The double precision name is DLSLZD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and Hermitian. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be \(n z\). The two integer arrays irow and jcol, each of length \(n z\), contain the row and column indices for these entries in \(A\). That is
\[
\begin{gathered}
A_{\text {irow }(i), \operatorname{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
\operatorname{irow}(i) \geq \operatorname{jcol}(i)
\end{gathered} \quad i=1, \ldots, \mathrm{nz} \quad . ~ \$
\]
with all other entries in the lower triangle of \(A\) zero.
The routine LSLZD solves a system of linear algebraic equations having a complex, sparse, Hermitian and positive definite coefficient matrix. It first uses the routine LSCXD to compute a
symbolic factorization of a permutation of the coefficient matrix. It then calls LNFZD to perform the numerical factorization. The solution of the linear system is then found using LFSZD.

The routine LSCXD computes a minimum degree ordering or uses a user-supplied ordering to set up the sparse data structure for the Cholesky factor, \(L\). Then the routine LNFZD produces the numerical entries in \(L\) so that we have
\[
P A P^{T}=L L^{H}
\]

Here P is the permutation matrix determined by the ordering.
The numerical computations can be carried out in one of two ways. The first method performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

Finally, the solution \(x\) is obtained by the following calculations:
\[
\begin{aligned}
& \text { 1) } L y_{1}=P b \\
& \text { 2) } L^{H} y_{2}=y_{1} \\
& \text { 3) } x=P^{T} y_{2}
\end{aligned}
\]

The routine LFSZD accepts \(b\) and the permutation vector which determines \(P\). It then returns \(x\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LZD/DL2LZD. The reference is:

CALL L2LZD (N, NZ, A, IROW, JCOL, B, X, IPER, IPARAM, RPARAM, WK, LWK, IWK, LIWK)

The additional arguments are as follows:
IPER - Vector of length N containing the ordering.
IPARAM - Integer vector of length 4. See Comment 3.
\(\boldsymbol{R P A R A M}\) - Real vector of length 2. See Comment 3.

WK - Complex work vector of length LWK.
\(\boldsymbol{L W K}\) - The length of wK, LWK should be at least \(2 \mathrm{~N}+6 \mathrm{Nz}\).
\(I W K\) - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least \(15 \mathrm{~N}+15 \mathrm{NZ}+9\).

Note that the parameter ITWKSP is not an argument for this routine.
2. Informational errors

Type Code
41 The coefficient matrix is not positive definite.
42 A column without nonzero elements has been found in the coefficient matrix.
3. If the default parameters are desired for L2LZD, then set IPARAM(1) to zero and call the routine L2LZD. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling L2LZD.

CALL L4LZD (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to L4LZD will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above. The arguments are as follows:
IPARAM - Integer vector of length 4.
```

$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ The numerical factorization method.

```
\begin{tabular}{ll} 
IPARAM(2) & Action \\
0 & Multifrontal \\
1 & Sparse column
\end{tabular}

Default: 0 .
\(\operatorname{IPARAM}(3)=\) The ordering option.

IPARAM(3
0

1

Action
Minimum degree ordering
User's ordering specified in IPER

Default: 0.
\(\operatorname{IPARAM}(4)=\) The total number of nonzeros in the factorization matrix.
RPARAM - Real vector of length 2.
\(\operatorname{RPARAM}(1)=\) The absolute value of the largest diagonal element in the Cholesky factorization.
\(\operatorname{RPARAM}(2)=\) The absolute value of the smallest diagonal element in the Cholesky factorization.

If double precision is required, then DL4LZD is called and RPARAM is declared double precision.

\section*{Example}

As an example, consider the \(3 \times 3\) linear system:
\[
A=\left[\begin{array}{rrr}
2+0 i & -1+i & 0 \\
-1-i & 4+0 i & 1+2 i \\
0 & 1-2 i & 10+0 i
\end{array}\right]
\]

Let \(x^{T}=(1+i, 2+2 i, 3+3 i)\) so that \(A x=(-2+2 i, 5+15 i, 36+28 i)^{T}\). The number of nonzeros in the lower triangle of \(A\) is \(\mathrm{nz}=5\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{crrrrr} 
irow & 1 & 2 & 3 & 2 & 3 \\
jcol & 1 & 2 & 3 & 1 & 2 \\
a & \(2+0 i\) & \(4+0 i\) & \(10+0 i\) & \(-1-i\) & \(1-2 i\)
\end{tabular}
or equivalently by
\begin{tabular}{crrrrr} 
arow & 3 & 2 & 3 & 1 & 2 \\
jcol & 3 & 1 & 2 & 1 & 2 \\
a & \(10+0 i\) & \(-1-i\) & \(1-2 i\) & \(2+0 i\) & \(4+0 i\)
\end{tabular}
```

USE LSLZD INT
USE WRCRN INT
INTEGER - N, NZ
PARAMETER (N=3, NZ=5)
INTEGER IROW(NZ), JCOL(NZ)
COMPLEX A(NZ), B(N), X(N)
DATA A/ (2.0,0.0), (4.0,0.0), (10.0,0.0), (-1.0,-1.0), (1.0,-2.0)/
DATA B/(-2.0,2.0), (5.0,15.0), (36.0,28.0)/
DATA IROW/1, 2, 3, 2, 3/
DATA JCOL/1, 2, 3, 1, 2/
CALL LSLZD (A, IROW, JCOL, B, X)
Print results
CALL WRCRN (' x ', X, 1, N, 1)
END

```

\section*{Output}
\((1.000,1.000)^{1} \quad(2.000,2.000)^{2} \quad(3.000,3.000)^{3}\)

\section*{LNFZD}

Computes the numerical Cholesky factorization of a sparse Hermitian matrix \(A\).

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex vector of length NZ containing the nonzero coefficients of the lower triangle of the linear system. (Input)

IROW - Vector of length NZ containing the row numbers of the corresponding elements in the lower triangle of A. (Input)

JCOL - Vector of length nZ containing the column numbers of the corresponding elements in the lower triangle of A. (Input)

MAXSUB - Number of subscripts contained in array NZSUB as output from subroutine LSCXD/DLSCXD. (Input)

NZSUB - Vector of length MAXSUB containing the row subscripts for the nonzeros in the Cholesky factor in compressed format as output from subroutine LSCXD/DLSCXD. (Input)

INZSUB - Vector of length \(\mathrm{N}+1\) containing pointers for NZSUB as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts for the nonzeros in column \(J\) are stored from location INZSUB( \(J\) ) to \(\operatorname{INZSUB}(J+1)-1\).
\(\boldsymbol{M A X N Z}\) — Length of RLNZ as output from subroutine LSCXD/DLSCXD. (Input)
ILNZ - Vector of length \(\mathrm{N}+1\) containing pointers to the Cholesky factor as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts for the nonzeros in column \(J\) of the factor are stored from location \(\operatorname{ILNZ}(\mathrm{J})\) to \(\operatorname{ILNZ}(\mathrm{J}+1)-1\).
(ILNZ, NZSUB, INZSUB) sets up the compressed data structure in column ordered form for the Cholesky factor.

IPER - Vector of length N containing the permutation as output from subroutine LSCXD/DLSCXD. (Input)

INVPER - Vector of length N containing the inverse permutation as output from subroutine LSCXD/DLSCXD. (Input)

ISPACE - The storage space needed for the stack of frontal matrices as output from subroutine LSCXD/DLSCXD. (Input)

DIAGNL - Complex vector of length N containing the diagonal of the factor. (Output)
\(\boldsymbol{R L N Z}\) - Complex vector of length MAXNZ containing the strictly lower triangle nonzeros of the Cholesky factor. (Output)

RPARAM - Parameter vector containing factorization information. (Output)
RPARAM (1) = smallest diagonal element in absolute value.
RPARAM (2) = largest diagonal element in absolute value.

\section*{Optional Arguments}
\(N\) - Number of equations. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\) IPER,1).
\(N Z\) - The number of nonzero coefficients in the linear system. (Input)
Default: \(N Z=\operatorname{size}(A, 1)\).
IJOB - Integer parameter selecting factorization method. (Input)
IJOB \(=1\) yields factorization in sparse column format.
\(I J O B=2\) yields factorization using multifrontal method.
Default: \(\operatorname{IJOB}=1\).

ITWKSP - The total workspace needed. (Input)
If the default is desired, set ITwKSP to zero. See Comment 1 for the default.
Default: \(I T W K S P=0\).

\section*{FORTRAN 90 Interface}

Generic: CALL LNFZD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ, IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM [, ...])

Specific: The specific interface names are S_LNFZD and D_LNFZD.

\section*{FORTRAN 77 Interface}

Single:
CALL LNFZD (N, NZ, A, IROW, JCOL, IJOB, MAXSUB, NZSUB, INZSUB, MAXNZ, ILNZ, IPER, INVPER, ISPACE, ITWKSP, DIAGNL, RLNZ, RPARAM)

Double: The double precision name is DLNFZD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and Hermitian. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be nz. The two integer arrays irow and jcol, each of length \(n z\), contain the row and column indices for these entries in \(A\). That is
\[
A_{\text {irow }(i), \mathrm{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz}
\]
\[
\operatorname{irow}(i) \geq \operatorname{jcol}(i) \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in the lower triangle of \(A\) zero.
The routine LNFZD produces the Cholesky factorization of \(P A P^{T}\) given the symbolic factorization of \(A\) which is computed by LSCXD. That is, this routine computes \(L\) which satisfies
\[
P A P^{T}=L L^{H}
\]

The diagonal of \(L\) is stored in DIAGNL and the strictly lower triangular part of \(L\) is stored in compressed subscript form in \(R=\) RLNZ as follows. The nonzeros in the \(j\) th column of \(L\) are stored in locations \(R(i), \ldots, R(i+k)\) where \(i=\operatorname{ILNZ}(j)\) and \(k=\operatorname{ILNZ}(j+1)-\operatorname{ILNZ}(j)-1\). The row subscripts are stored in the vector NZSUB from locations INZSUB \((j)\) to \(\operatorname{INZSUB}(j)+k\).

The numerical computations can be carried out in one of two ways. The first method (when \(I J O B=2\) ) performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method (when I JOB \(=1\) ) is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme.

\section*{Comments}
1. Workspace may be explicitly provided by use of L2FZD/DL2FZD. The reference is:
```

CALL L2FZD (N, NZ, A, IROW, JCOL, IJOB, MAXSUB, NZSUB, INZSUB,
MAXNZ, ILNZ, IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM, WK,
LWK, IWK,LIWK)

```

The additional arguments are as follows:
WK - Complex work vector of length LWK.
\(\boldsymbol{L W K}\) - The length of wK, LWK should be at least \(\mathrm{N}+3 \mathrm{Nz}\).
IWK - Integer work vector of length LIWK.
LIWK - The length of IWK, LIWK should be at least 2 N .
Note that the parameter ITWKSP is not an argument to this routine.
2. Informational errors

Type Code
41 The coefficient matrix is not positive definite.
42 A column without nonzero elements has been found in the coefficient matrix.

\section*{Example}

As an example, consider the \(3 \times 3\) linear system:
\[
A=\left[\begin{array}{rrr}
2+0 i & -1+i & 0 \\
-1-i & 4+0 i & 1+2 i \\
0 & 1-2 i & 10+0 i
\end{array}\right]
\]

The number of nonzeros in the lower triangle of \(A\) is \(n z=5\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{crrrrr} 
irow & 1 & 2 & 3 & 2 & 3 \\
jcol & 1 & 2 & 3 & 1 & 2 \\
\(a\) & \(2+0 i\) & \(4+0 i\) & \(10+0 i\) & \(-1-i\) & \(1-2 i\)
\end{tabular}
or equivalently by
\begin{tabular}{crrrrr} 
arow & 3 & 2 & 3 & 1 & 2 \\
jcol & 3 & 1 & 2 & 1 & 2 \\
\(a\) & \(10+0 i\) & \(-1-i\) & \(1-2 i\) & \(2+0 i\) & \(4+0 i\)
\end{tabular}

We first call LSCXD to produce the symbolic information needed to pass on to LNFZD. Then call LNFZD to factor this matrix. The results are displayed below.
```

USE LNFZD INT
USE LSCXD }\mp@subsup{}{}{-}\mathrm{ INT
USE WRCRN_INT
INTEGER N, NZ, NRLNZ
PARAMETER (N=3, NZ=5, NRLNZ=5)
INTEGER IJOB, ILNZ (N+1), INVPER(N), INZSUB(N+1), IPER(N),\&
IROW(NZ), ISPACE, JCOL(NZ), MAXNZ, MAXSUB, \&
NZSUB(3*NZ)
REAL RPARAM(2)
COMPLEX A(NZ), DIAGNL(N), RLNZ(NRLNZ)
DATA A/ (2.0,0.0), (4.0,0.0), (10.0,0.0), (-1.0,-1.0), (1.0,-2.0)/
DATA IROW/1, 2, 3, 2, 3/
DATA JCOL/1, 2, 3, 1, 2/
Select minimum degree ordering
for multifrontal method
IJOB = 3
MAXSUB = 3*NZ
CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER, \&
IJOB=IJOB, MAXSUB=MAXSUB)
Check if NRLNZ is large enough
IF (NRLNZ .GE. MAXNZ) THEN
IJOB = 2
CALL LNFZD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB, MAXNZ, \&
ILNZ, IPER, INVPER, ISPACE, DIAGNL, RLNZ, RPARAM, \&
IJOB=IJOB)
Print results
CALL WRCRN (' diagnl ', DIAGNL, 1, N, 1)
CALL WRCRN (' rlnz ', RLNZ, 1, MAXNZ, 1)
END IF

```
!
END

\section*{Output}

\section*{diagnl}
\((1.414,0.00)^{1} \quad(1.732,0.000)^{2}(2.887,0.000)^{3}\)

> rlnz
\((-0.707,-0.707)^{1} \quad 2\)
\((-0.707,-0.707)(0.577,-1.155)\)

\section*{LFSZD}

Solves a complex sparse Hermitian positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.

\section*{Required Arguments}
\(N-\) Number of equations. (Input)
MAXSUB - Number of subscripts contained in array NZSUB as output from subroutine LSCXD/DLSCXD. (Input)

NZSUB - Vector of length MAXSUB containing the row subscripts for the off-diagonal nonzeros in the factor as output from subroutine LSCXD/DLSCXD. (Input)

INZSUB - Vector of length \(\mathrm{N}+1\) containing pointers for NZSUB as output from subroutine LSCXD/DLSCXD. (Input)
The row subscripts of column \(J\) are stored from location INZSUB(J) to
\(\operatorname{INZSUB}(J+1)-1\).
MAXNZ - Total number of off-diagonal nonzeros in the Cholesky factor as output from subroutine LSCXD/DLSCXD. (Input)
\(\boldsymbol{R L N Z}\) - Complex vector of length MAXNZ containing the off-diagonal nonzeros in the factor in column ordered format as output from subroutine LNFZD/DLNFZD. (Input)
\(\boldsymbol{I L N Z}\) - Vector of length \(\mathrm{N}+1\) containing pointers to RLNZ as output from subroutine
LSCXD/DLSCXD. The nonzeros in column \(J\) of the factor are stored from location \(\operatorname{ILNZ}(\mathrm{J})\) to \(\operatorname{ILNZ}(\mathrm{J}+1)-1\). (Input)
The values (RLNZ, ILNZ, NZSUB, INZSUB) give the off-diagonal nonzeros of the factor in a compressed subscript data format.

DIAGNL - Complex vector of length \(N\) containing the diagonals of the Cholesky factor as output from subroutine LNFZD/DLNFZD. (Input)

IPER - Vector of length \(N\) containing the ordering as output from subroutine LSCXD/DLSCXD. (Input) \(\operatorname{IPER}(I)=K\) indicates that the original row \(K\) is the new row \(I\).

B - Complex vector of length N containing the right-hand side. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL LFSZD (N, MAXZUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL, IPER, B, X)

Specific: The specific interface names are S_LFSZD and D_LFSZD.

\section*{FORTRAN 77 Interface}

Single: CALL LFSZD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL, IPER, B, X)

Double: The double precision name is DLFSZD.

\section*{Description}

Consider the linear equation
\[
A x=b
\]
where \(A\) is sparse, positive definite and Hermitian. The sparse coordinate format for the matrix \(A\) requires one complex and two integer vectors. The complex array a contains all the nonzeros in the lower triangle of \(A\) including the diagonal. Let the number of nonzeros be nz . The two integer arrays irow and jcol, each of length \(n z\), contain the row and column indices for these entries in \(A\). That is
\[
\begin{gathered}
A_{\text {irow }(i), \operatorname{icol}(i)}=\mathrm{a}(i), \quad i=1, \ldots, \mathrm{nz} \\
\operatorname{irow}(i) \geq \operatorname{jcol}(i)
\end{gathered} \quad i=1, \ldots, \mathrm{nz}
\]
with all other entries in the lower triangle of \(A\) zero.
The routine LFSZD computes the solution of the linear system given its Cholesky factorization. The factorization is performed by calling LSCXD followed by LNFZD. The routine LSCXD computes a minimum degree ordering or uses a user-supplied ordering to set up the sparse data structure for the Cholesky factor, \(L\). Then the routine LNFZD produces the numerical entries in L so that we have
\[
P A P^{T}=L L^{H}
\]

Here \(P\) is the permutation matrix determined by the ordering.

The numerical computations can be carried out in one of two ways. The first method performs the factorization using a multifrontal technique. This option requires more storage but in certain cases will be faster. The multifrontal method is based on the routines in Liu (1987). For detailed description of this method, see Liu (1990), also Duff and Reid (1983, 1984), Ashcraft (1987), Ashcraft et al. (1987), and Liu (1986, 1989). The second method is fully described in George and Liu (1981). This is just the standard factorization method based on the sparse compressed storage scheme. Finally, the solution \(x\) is obtained by the following calculations:
\[
\begin{aligned}
& \text { 1) } L y_{1}=P b \\
& \text { 2) } L^{H} y_{2}=y_{1} \\
& \text { 3) } x=P^{T} y_{2}
\end{aligned}
\]

\section*{Comments}

Informational error
\begin{tabular}{cc} 
Type & Code \\
4 & 1
\end{tabular}

\section*{Example}

As an example, consider the \(3 \times 3\) linear system:
\[
A=\left[\begin{array}{rrr}
2+0 i & -1+i & 0 \\
-1-i & 4+0 i & 1+2 i \\
0 & 1-2 i & 10+0 i
\end{array}\right]
\]

Let
\[
x_{1}^{T}=(1+i, 2+2 i, 3+3 i)
\]
so that \(A x_{1}=(-2+2 i, 5+15 i, 36+28 i)^{T}\), and
\[
x_{2}^{T}=(3+3 i, 2+2 i, 1+1 i)
\]
so that \(A x_{2}=(2+6 i, 7-5 i, 16+8 i)^{T}\). The number of nonzeros in the lower triangle of \(A\) is \(\mathrm{nz}=5\). The sparse coordinate form for the lower triangle of \(A\) is given by:
\begin{tabular}{rrrrrr} 
irow & 1 & 2 & 3 & 2 & 3 \\
jcol & 1 & 2 & 3 & 1 & 2 \\
a & \(2+0 i\) & \(4+0 i\) & \(10+0 i\) & \(-1-i\) & \(1-2 i\)
\end{tabular}
or equivalently by
\begin{tabular}{rrrrrr} 
irow & 3 & 2 & 3 & 1 & 2 \\
jcol & 3 & 1 & 2 & 1 & 2 \\
a & \(10+0 i\) & \(-1-i\) & \(1-2 i\) & \(2+0 i\) & \(4+0 i\)
\end{tabular}
```

USE IMSL_LIBRARIES
INTEGER N, NZ, NRLNZ

```
```

PARAMETER (N=3, NZ=5, NRLNZ=5)
INTEGER IJOB, ILNZ(N+1), INVPER(N), INZSUB(N+1), IPER(N),\&
IROW(NZ), ISPACE, JCOL(NZ), MAXNZ, MAXSUB,\&
NZSUB (3*NZ)
COMPLEX A(NZ), B1 (N), B2 (N), DIAGNL (N), RLNZ (NRLNZ), X(N)
REAL RPARAM(2)
!
DATA A/ (2.0,0.0), (4.0,0.0), (10.0,0.0), (-1.0,-1.0), (1.0,-2.0)/
DATA B1/(-2.0,2.0), (5.0,15.0), (36.0,28.0)/
DATA B2/(2.0,6.0), (7.0,5.0), (16.0,8.0)/
DATA IROW/1, 2, 3, 2, 3/
DATA JCOL/1, 2, 3, 1, 2/
Select minimum degree ordering
for multifrontal method
IJOB = 3
MAXSUB = 3*NZ
CALL LSCXD (IROW, JCOL, NZSUB, INZSUB, MAXNZ, ILNZ, INVPER, \&
IJOB=IJOB, MAXSUB=MAXSUB, IPER=IPER, ISPACE=ISPACE)
Check if NRLNZ is large enough
IF (NRLNZ .GE. MAXNZ) THEN
Choose multifrontal method
IJOB = 2
CALL LNFZD (A, IROW, JCOL, MAXSUB, NZSUB, INZSUB,\&
MAXNZ, ILNZ, IPER, INVPER, ISPACE, DIAGNL, \&
RLNZ, RPARAM, IJOB=IJOB)
Solve A * X1 = B1
CALL LFSZD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL,\&
IPER, B1, X)
Print X1
CALL WRCRN (' x1 ', X, 1, N,1)
Solve A * X2 = B2
CALL LFSZD (N, MAXSUB, NZSUB, INZSUB, MAXNZ, RLNZ, ILNZ, DIAGNL,\&
IPER, B2, X)
Print X2
CALL WRCRN (' x2 ', X, 1, N,1)
END IF
!
END

```

\section*{Output}
\begin{tabular}{rrr} 
& \(x^{1}\) \\
\((1.000,1.000)^{2}\) & \((2.000,2.000)^{2}\) & \((3.000,3.000)^{3}\) \\
\((3.000,3.000)^{1}\) & \((2.000,2.000)^{2}\) & \((1.000,1.000)^{3}\)
\end{tabular}

\section*{LSLTO}

Solves a complex sparse Hermitian positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.

\section*{Required Arguments}

A-Real vector of length \(2 \mathrm{~N}-1\) containing the first row of the coefficient matrix followed by its first column beginning with the second element. (Input) See Comment 2.

B - Real vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Real vector of length N containing the solution of the linear system. (Output)
If \(B\) is not needed then \(B\) and \(X\) may share the same storage locations.

\section*{Optional Arguments}
\(N\) - Order of the matrix represented by A. (Input)
Default: \(\mathrm{N}=(\operatorname{size}(\mathrm{A}, 1)+1) / 2\)
IPATH - Integer flag. (Input) IPATH \(=1\) means the system \(A x=B\) is solved. IPATH \(=2\) means the system \(A^{T} x=B\) is solved. Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLTO (A, B, X [, ...])
Specific: The specific interface names are S_LSLTO and D_LSLTO.

\section*{FORTRAN 77 Interface}

Single: CALL LSLTO (N, A, B, IPATH, X)
Double: \(\quad\) The double precision name is DLSLTO.

\section*{Description}

Toeplitz matrices have entries that are constant along each diagonal, for example,
\[
A=\left[\begin{array}{cccc}
p_{0} & p_{1} & p_{2} & p_{4} \\
p_{-1} & p_{0} & p_{1} & p_{2} \\
p_{-2} & p_{-1} & p_{0} & p_{1} \\
p_{-3} & p_{-2} & p_{-1} & p_{0}
\end{array}\right]
\]

The routine LSLTO is based on the routine TSLS in the TOEPLITZ package, see Arushanian et al. (1983). It is based on an algorithm of Trench (1964). This algorithm is also described by Golub and van Loan (1983), pages 125-133.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LTO/DL2LTO. The reference is:

CALL L2LTO (N, A, B, IPATH, X, WK)
The additional argument is:
\(W K-\) Work vector of length \(2 \mathrm{~N}-2\).
2. Because of the special structure of Toeplitz matrices, the first row and the first column of a Toeplitz matrix completely characterize the matrix. Hence, only the elements \(A(1,1), \ldots, A(1, N), A(2,1), \ldots, A(N, 1)\) need to be stored.

\section*{Example}

A system of four linear equations is solved. Note that only the first row and column of the matrix \(A\) are entered.
```

USE LSLTO INT
USE WRRRN INT
INTEGER N
PARAMETER (N=4)
REAL A(2*N-1), B(N), X(N)
Set values for A, and B
A=($$
\begin{array}{llll}{2}&{-3}&{-1}&{6}\end{array}
$$)
([$$
\begin{array}{llll}{1}&{2}&{-3}&{-1}\end{array}
$$)
( }$$
\begin{array}{rrrrr}{4}&{1}&{2}&{-3}\\{(}&{4}&{1}&{2}\end{array}
$$
B = ($$
\begin{array}{llll}{16}&{-29 -7 5 )}\end{array}
$$)
DATA A/2.0, -3.0, -1.0, 6.0, 1.0, 4.0, 3.0/
DATA B/16.0, -29.0, -7.0, 5.0/
Solve AX = B
CALL LSLTO (A, B, X)
CALL WRRRN ('X', X, 1, N, 1)
END

```
!

\section*{Output}
\begin{tabular}{rrrrr} 
& & X & & \\
1 & 2 & 3 & 4 \\
-2.000 & -1.000 & 7.000 & 4.000
\end{tabular}

\section*{LSLTC}

Solves a complex Toeplitz linear system.

\section*{Required Arguments}

A-Complex vector of length \(2 \mathrm{~N}-1\) containing the first row of the coefficient matrix followed by its first column beginning with the second element. (Input) See Comment 2.
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution of the linear system. (Output)

\section*{Optional Arguments}
\(N\) — Order of the matrix represented by A. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{A}, 1)\).
IPATH - Integer flag. (Input)
IPATH \(=1\) means the system \(A x=B\) is solved.
IPATH \(=2\) means the system \(A^{T} x=B\) is solved.
Default: \(\operatorname{IPATH}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLTC (A, B, X [,...])
Specific: The specific interface names are S_LSLTC and D_LSLTC.

\section*{FORTRAN 77 Interface}

Single: CALL LSLTC ( \(\mathrm{N}, \mathrm{A}, \mathrm{B}, \mathrm{IPATH}, \mathrm{X}\) )
Double: The double precision name is DLSLTC.

\section*{Description}

Toeplitz matrices have entries which are constant along each diagonal, for example,
\[
A=\left[\begin{array}{cccc}
p_{0} & p_{1} & p_{2} & p_{4} \\
p_{-1} & p_{0} & p_{1} & p_{2} \\
p_{-2} & p_{-1} & p_{0} & p_{1} \\
p_{-3} & p_{-2} & p_{-1} & p_{0}
\end{array}\right]
\]

The routine LSLTC is based on the routine TSLC in the TOEPLITZ package, see Arushanian et al. (1983). It is based on an algorithm of Trench (1964). This algorithm is also described by Golub and van Loan (1983), pages 125-133.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LTC/DL2LTC. The reference is:

CALL L2LTC ( \(\mathrm{N}, \mathrm{A}, \mathrm{B}\), IPATH, X, WK)
The additional argument is

WK - Complex work vector of length \(2 \mathrm{~N}-2\).
2. Because of the special structure of Toeplitz matrices, the first row and the first column of a Toeplitz matrix completely characterize the matrix. Hence, only the elements \(A(1,1), \ldots, A(1, N), A(2,1), \ldots, A(N, 1)\) need to be stored.

\section*{Example}

A system of four complex linear equations is solved. Note that only the first row and column of the matrix \(A\) are entered.
```

USE LSLTC INT
USE WRCRN INT
PARAMETER (N=4)
COMPLEX A(2*N-1), B(N), X(N)
Set values for A and B
A = ( 2+2i -3 1+4i 6-2i )
( i 2+2i -3 1+4i )
( 4+2i i 2+2i -3 )
( 3-4i 4+2i i 2+2i)
B = ( 6+65i -29-16i 7+i -10+i )
DATA A/ (2.0,2.0), (-3.0,0.0), (1.0,4.0), (6.0,-2.0), (0.0,1.0),\&
(4.0,2.0), (3.0,-4.0)/
DATA B/(6.0,65.0), (-29.0,-16.0), (7.0,1.0), (-10.0,1.0)/
Solve AX = B
CALL LSLTC (A, B, X)
CALL WRCRN ('X', X, 1, N, 1)
END

```
!

\section*{Output}

\section*{x}
\((-2.000,0.000)^{1} \quad(-1.000,-5.000)^{2} \quad(7.000,2.00)^{3} \quad(0.000,4.000)^{4}\)

\section*{LSLCC}

Solves a complex circulant linear system.

\section*{Required Arguments}
\(\boldsymbol{A}\) - Complex vector of length N containing the first row of the coefficient matrix. (Input)
\(\boldsymbol{B}\) - Complex vector of length N containing the right-hand side of the linear system. (Input)
\(\boldsymbol{X}\) - Complex vector of length N containing the solution of the linear system. (Output)

\section*{Optional Arguments}
\(N\) — Order of the matrix represented by A. (Input)
Default: \(N=\operatorname{size}(A, 1)\).
IPATH - Integer flag. (Input)
IPATH \(=1\) means the system \(\mathrm{Ax}=\mathrm{B}\) is solved.
IPATH \(=2\) means the system \(\mathrm{A}^{T} \mathrm{x}=\mathrm{B}\) is solved.
Default: ІРАТН \(=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL LSLCC (A, B, X [, ...])
Specific: The specific interface names are \(S_{-}\)LSLCC and D_LSLCC.

\section*{FORTRAN 77 Interface}

Single: CALL LSLCC (N, A, B, IPATH, X)
Double: The double precision name is DLSLCC.

\section*{Description}

Circulant matrices have the property that each row is obtained by shifting the row above it one place to the right. Entries that are shifted off at the right re-enter at the left. For example,
\[
A=\left[\begin{array}{llll}
p_{1} & p_{2} & p_{3} & p_{4} \\
p_{4} & p_{1} & p_{2} & p_{3} \\
p_{3} & p_{4} & p_{1} & p_{2} \\
p_{2} & p_{3} & p_{4} & p_{1}
\end{array}\right]
\]

If \(q_{k}=p_{-k}\) and the subscripts on \(p\) and \(q\) are interpreted modulo \(N\), then
\[
(A x)_{j}=\sum_{i=1}^{N} p_{i-j+1} x_{i}=\sum_{i=1}^{N} q_{j-i+1} x_{i}=(q * x)_{i}
\]
where \(q * x\) is the convolution of \(q\) and \(x\). By the convolution theorem, if \(q * x=b\), then
\[
\hat{q} \otimes \hat{x}=\hat{b} \text {, where } \hat{q}
\]
is the discrete Fourier transform of \(q\) as computed by the IMSL routine FFTCF and \(\otimes\) denotes elementwise multiplication. By division,
\[
\hat{x}=\hat{b} \varnothing \hat{q}
\]
where \(\varnothing\) denotes elementwise division. The vector \(x\) is recovered from
\[
\hat{x}
\]
through the use of IMSL routine FFTCB.
To solve \(A^{T} x=b\), use the vector \(p\) instead of \(q\) in the above algorithm.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of L2LCC/DL2LCC. The reference is:
```

CALL L2LCC (N,A, B, IPATH, X, ACOPY,WK)

```

The additional arguments are as follows:
ACOPY - Complex work vector of length N. If A is not needed, then A and ACOPY may be the same.

WK - Work vector of length \(6 \mathrm{~N}+15\).
2. Informational error
\begin{tabular}{ccc} 
Type & Code \\
4 & 2 & The input matrix is singular.
\end{tabular}
3. Because of the special structure of circulant matrices, the first row of a circulant matrix completely characterizes the matrix. Hence, only the elements \(A(1,1), \ldots, A(1, N)\) need to be stored.

\section*{Example}

A system of four linear equations is solved. Note that only the first row of the matrix \(A\) is entered.
```

USE LSLCC INT
USE WRCRN_INT
INTEGER N
PARAMETER (N=4)
COMPLEX A(N), B(N), X(N)

```
!
```

    Set values for A, and B
    A = ( 2+2i -3+0i 1+4i 6-2i)
    B = (6+65i -41-10i -8-30i 63-3i)
    DATA A/ (2.0,2.0), (-3.0,0.0), (1.0,4.0), (6.0,-2.0)/
    DATA B/ (6.0,65.0), (-41.0,-10.0), (-8.0,-30.0), (63.0,-3.0)/
        Solve AX = B (IPATH = 1)
    CALL LSLCC (A, B, X)
    CALL WRCRN ('X', X, 1, N, 1)
END

```

\section*{Output}
\((-2.000,0.000)^{1} \quad(-1.000,-5.000)^{2} \quad(7.000,2.00)^{3} \quad(0.000,4.000)^{4}\)

\section*{PCGRC}

Solves a real symmetric definite linear system using a preconditioned conjugate gradient method with reverse communication.

\section*{Required Arguments}

IDO - Flag indicating task to be done. (Input/Output)
On the initial call IDO must be 0 . If the routine returns with IDO \(=1\), then set \(Z=A P\), where \(A\) is the matrix, and call PCGRC again. If the routine returns with \(I D O=2\), then set \(Z\) to the solution of the system \(M Z=R\), where \(M\) is the preconditioning matrix, and call PCGRC again. If the routine returns with IDO \(=3\), then the iteration has converged and X contains the solution.
\(\boldsymbol{X}\) - Array of length N containing the solution. (Input/Output)
On input, x contains the initial guess of the solution. On output, X contains the solution to the system.
\(\boldsymbol{P}\) - Array of length N. (Output) Its use is described under IDO.
\(\boldsymbol{R}\) - Array of length N. (Input/Output)
On initial input, it contains the right-hand side of the linear system. On output, it contains the residual.
\(\mathbf{Z}\) - Array of length N. (Input)
When \(\operatorname{IDO}=1\), it contains AP, where \(A\) is the linear system. When \(I D O=2\), it contains the solution of \(M Z=R\), where \(M\) is the preconditioning matrix. When \(I D O=0\), it is ignored. Its use is described under IDO.

\section*{Optional Arguments}
\(N\) - Order of the linear system. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)\).
RELERR - Relative error desired. (Input)
Default: RELERR \(=1 . \mathrm{e}-5\) for single precision and 1.d-10 for double precision.
ITMAX - Maximum number of iterations allowed. (Input)
Default: \(I T M A X=N\).

\section*{FORTRAN 90 Interface}

Generic: CALL PCGRC (IDO, X, P, R, Z [, ...])
Specific: The specific interface names are S_PCGRC and D_PCGRC.

\section*{FORTRAN 77 Interface}

Single: CALL PCGRC (IDO, N, X, P, R, Z, RELERR, ITMAX)
Double: The double precision name is DPCGRC.

\section*{Description}

Routine PCGRC solves the symmetric definite linear system \(A x=b\) using the preconditioned conjugate gradient method. This method is described in detail by Golub and Van Loan (1983, Chapter 10), and in Hageman and Young (1981, Chapter 7).

The preconditioning matrix, \(M\), is a matrix that approximates \(A\), and for which the linear system \(M z=r\) is easy to solve. These two properties are in conflict; balancing them is a topic of much current research.

The number of iterations needed depends on the matrix and the error tolerance RELERR. As a rough guide, ITMAX \(=N^{1 / 2}\) is often sufficient when \(N \gg 1\). See the references for further information.

Let \(M\) be the preconditioning matrix, let \(b, p, r, x\) and \(z\) be vectors and let \(\tau\) be the desired relative error. Then the algorithm used is as follows.
\(\lambda=-1\)
\(p_{0}=x_{0}\)
\(r_{1}=b-A p\)
For \(k=1, \ldots\), itmax
\(z_{k}=M^{-1} r_{k}\)
If \(k=1\) then
\[
\begin{aligned}
& \beta_{k}=1 \\
& p_{k}=z_{k}
\end{aligned}
\]

Else
\[
\begin{aligned}
& \beta_{k}=z_{k}^{T} r_{k} / z_{k-1}^{T} r_{k-1} \\
& p_{k}=z_{k}+\beta_{k} p_{k}
\end{aligned}
\]

End if
\[
z_{k}=A p
\]
\[
\alpha_{k}=z_{k-1}^{T} r_{k-1} / z_{k}^{T} p_{k}
\]
\[
x_{k}=x_{k}+\alpha_{k} p_{k}
\]
\[
r_{k}=r_{k}-\alpha_{k} z_{k}
\]

If \(\left(\left\|z_{k}\right\|_{2} \leq \tau(1-\lambda)\left\|x_{k}\right\|_{2}\right)\) Then
Recompute \(\lambda\)
\[
\text { If }\left(\left\|z_{k}\right\|_{2} \leq \tau(1-\lambda)\left\|x_{k}\right\|_{2}\right) \text { Exit }
\]

End if end loop
Here \(\lambda\) is an estimate of \(\lambda_{\max }(G)\), the largest eigenvalue of the iteration matrix \(G=I-M^{-1} A\). The stopping criterion is based on the result (Hageman and Young, 1981, pages 148-151)
\[
\frac{\left\|x_{k}-x\right\|_{M}}{\|x\|_{M}} \leq \frac{1}{1-\lambda_{\max }(G)} \frac{\left\|z_{k}\right\|_{M}}{\left\|x_{k}\right\|_{M}}
\]

Where
\[
\|x\|_{M}^{2}=x^{T} M x
\]

It is known that
\[
\lambda_{\max }\left(T_{1}\right) \leq \lambda_{\max }\left(T_{2}\right) \leq \cdots \leq \lambda_{\max }(G)<1
\]
where the \(T_{n}\) are the symmetric, tridiagonal matrices
\[
T_{n}=\left[\begin{array}{ccccc}
\mu_{1} & \omega_{2} & & & \\
\omega_{2} & \mu_{2} & \omega_{3} & & \\
& \omega_{3} & \mu_{3} & \omega_{4} & \\
& & \ddots & \ddots & \ddots
\end{array}\right]
\]
with
\[
\mu_{k}=1-\beta_{k} / \alpha_{k-1}-1 / \alpha_{k}, \mu_{1}=1-1 / \alpha_{1}
\]
and
\[
\omega_{k}=\sqrt{\beta_{k}} / \alpha_{k-1}
\]

The largest eigenvalue of \(T_{k}\) is found using the routine EVASB. Usually this eigenvalue computation is needed for only a few of the iterations.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of P2GRC/DP2GRC. The reference is:
```

CALL P2GRC (IDO, N, X, P, R, Z, RELERR, ITMAX, TRI, WK,IWK)

```

The additional arguments are as follows:
TRI - Workspace of length 2 * ITMAX containing a tridiagonal matrix (in band symmetric form) whose largest eigenvalue is approximately the same as the largest eigenvalue of the iteration matrix. The workspace arrays TRI, WK and IWK should not be changed between the initial call with IDO \(=0\) and \(P C G R C / D P C G R C\) returning with \(I D O=3\).

WK - Workspace of length 5 * ITMAX.
IWK - Workspace of length ITMAX.
2. Informational errors
Type Code
\begin{tabular}{lll}
4 & 1 & The preconditioning matrix is singular. \\
4 & 2 & The preconditioning matrix is not definite. \\
4 & 3 & The linear system is not definite. \\
4 & 4 & The linear system is singular. \\
4 & 5 & No convergence after ITMAX iterations.
\end{tabular}

\section*{Example}

In this example, the solution to a linear system is found. The coefficient matrix \(A\) is stored as a full matrix. The preconditioning matrix is the diagonal of \(A\). This is called the Jacobi preconditioner. It is also used by the IMSL routine JCGRC.
```

USE PCGRC INT
USE MURRV_INT
USE WRRRN_INT
USE SCOPY ${ }^{-}$INT
INTEGER - LDA, N
PARAMETER ( $\mathrm{N}=3$, LDA=N)
INTEGER IDO, ITMAX, J
REAL $\quad A(L D A, N), B(N), P(N), R(N), X(N), Z(N)$
( 1, -3, 2)
$A=\left(\begin{array}{lll}-3, & 10, & -5\end{array}\right)$
( 2, -5, 6 )
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
$B=\quad(\quad 27.0,-78.0,64.0)$
DATA B/27.0, -78.0, 64.0/

```
```

    CALL SCOPY (N, B, 1, R, 1)
    CALL SCOPY (N, B, 1, X, 1)
    !
ITMAX = 100
IDO = 0
1 0 ~ C A L L ~ P C G R C ~ ( I D O , ~ X , ~ P , ~ R , ~ Z , ~ I T M A X = I T M A X ) ,
IF (IDO .EQ. 1) THEN
! Set z = Ap
CALL MURRV (A, P, Z)
GO TO 10
ELSE IF (IDO .EQ. 2) THEN
! Use diagonal of A as the
preconditioning matrix M
and set z = inv(M)*r
DO 20 J=1, N
Z(J) = R(J)/A(J,J)
2 0
CONTINUE
GO TO 10
END IF
! Print the solution
CALL WRRRN ('Solution', X)
END

```

\section*{Output}

Solution
\(1 \quad 1.001\)
\(2-4.000\)
\(3 \quad 7.000\)

\section*{Example 2}

In this example, a more complicated preconditioner is used to find the solution of a linear system which occurs in a finite-difference solution of Laplace's equation on a \(4 \times 4\) grid. The matrix is
\[
A=\left[\begin{array}{ccccccccc}
4 & -1 & 0 & -1 & & & & & \\
-1 & 4 & -1 & 0 & -1 & & & & \\
0 & -1 & 4 & -1 & 0 & -1 & & & \\
-1 & 0 & -1 & 4 & -1 & 0 & -1 & & \\
& -1 & 0 & -1 & 4 & -1 & 0 & -1 & \\
& & -1 & 0 & -1 & 4 & -1 & 0 & -1 \\
& & & -1 & 0 & -1 & 4 & -1 & 0 \\
& & & & -1 & 0 & -1 & 4 & -1 \\
& & & & & -1 & 0 & -1 & 4
\end{array}\right]
\]

The preconditioning matrix \(M\) is the symmetric tridiagonal part of \(A\),
\[
M=\left[\begin{array}{rrrrrrrr}
4 & -1 & & & & & & \\
-1 & 4 & -1 & & & & & \\
& -1 & 4 & -1 & & & & \\
& & -1 & 4 & -1 & & & \\
& & & -1 & 4 & -1 & & \\
& & & & -1 & 4 & -1 & \\
& & & & & -1 & 4 & -1 \\
& & & & & & -1 & 4 \\
& & & & & & -1 \\
& & & & & & & -1
\end{array}\right)
\]

Note that \(M\), called PRECND in the program, is factored once.
```

USE IMSL LIBRARIES
INTEGER LDA, LDPRE, N, NCODA, NCOPRE
PARAMETER }\begin{array}{c}{(\textrm{N}=9, NCODA=3, NCOPRE=1, LDA=2*NCODA+1, \&}<br>{LDPRE=NCOPRE+1)}
INTEGER IDO, ITMAX
REAL A (LDA,N), P(N), PRECND (LDPRE,N), PREFAC (LDPRE,N),\&
R(N), RCOND, RELERR, X(N), Z(N)
Set A in band form
DATA A/3*0.0, 4.0, -1.0, 0.0, -1.0, 2*0.0, -1.0, 4.0, -1.0, 0.0,\&
-1.0, 2*0.0, -1.0, 4.0, -1.0, 0.0, -1.0, -1.0, 0.0, -1.0,\&
4.0, -1.0, 0.0, -1.0, -1.0, 0.0, -1.0, 4.0, -1.0, 0.0,\&
-1.0, -1.0, 0.0, -1.0, 4.0, -1.0, 0.0, -1.0, -1.0, 0.0, \&
-1.0, 4.0, -1.0, 2*0.0, -1.0, 0.0, -1.0, 4.0, -1.0, 2*0.0,\&
-1.0, 0.0, -1.0, 4.0, 3*0.0/
Set PRECND in band symmetric form
DATA PRECND/0.0, 4.0, -1.0, 4.0, -1.0, 4.0, -1.0, 4.0, -1.0, 4.0,\&
-1.0, 4.0, -1.0, 4.0, -1.0, 4.0, -1.0, 4.0/
R = 1.0E0
X = 0.0EO
CALL LFCQS (PRECND, NCOPRE, PREFAC, RCOND)
ITMAX = 100
RELERR = 1.0E-4
IDO = 0
1 0 ~ C A L L ~ P C G R C ~ ( I D O , ~ X , ~ P , ~ R , ~ Z , ~ R E L E R R = R E L E R R , ~ I T M A X = I T M A X ) ~
IF (IDO .EQ. 1) THEN
Set z = Ap
CALL MURBV (A, NCODA, NCODA, P, Z)
GO TO 10
ELSE IF (IDO .EQ. 2) THEN
Solve PRECND*z = r for r
CALL LSLQS (PREFAC, NCOPRE, R, Z)
GO TO 10
END IF
! Print the solution

```
!
END

\section*{Output}

Solution
10.955
21.241
31.349
41.578
\(5 \quad 1.660\)
\(6 \quad 1.578\)
\(7 \quad 1.349\)
\(8 \quad 1.241\)
90.955

\section*{JCGRC}

Solves a real symmetric definite linear system using the Jacobi-preconditioned conjugate gradient method with reverse communication.

\section*{Required Arguments}

IDO - Flag indicating task to be done. (Input/Output)
On the initial call IDO must be 0 . If the routine returns with IDO \(=1\), then set \(Z=A * P\), where \(A\) is the matrix, and call JCGRC again. If the routine returns with \(I D O=2\), then the iteration has converged and \(X\) contains the solution.

DIAGNL - Vector of length \(N\) containing the diagonal of the matrix. (Input) Its elements must be all strictly positive or all strictly negative.
\(\boldsymbol{X}\) - Array of length N containing the solution. (Input/Output)
On input, x contains the initial guess of the solution. On output, x contains the solution to the system.
\(\boldsymbol{P}\) - Array of length N. (Output)
Its use is described under IDO.
\(\boldsymbol{R}\) - Array of length N. (Input/Output)
On initial input, it contains the right-hand side of the linear system. On output, it contains the residual.
\(\mathbf{Z}\) - Array of length N. (Input)
When \(\operatorname{IDO}=1\), it contains AP, where A is the linear system. When IDO \(=0\), it is ignored. Its use is described under IDO.

\section*{Optional Arguments}
\(N\) - Order of the linear system. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{x}, 1)\).
RELERR - Relative error desired. (Input)
Default: ReLerr \(=1 . \mathrm{e}-5\) for single precision and 1.d-10 for double precision.
ITMAX - Maximum number of iterations allowed. (Input)
Default: \(I T\) MAX \(=100\).

\section*{FORTRAN 90 Interface}

Generic: CALL JCGRC (IDO, DIAGNL, \(\mathrm{x}, \mathrm{P}, \mathrm{R}, \mathrm{Z}[, \ldots]\) )
Specific: The specific interface names are S_JCGRC and D_JPCGRC.

\section*{FORTRAN 77 Interface}

Single: CALL JCGRC (IDO, N, DIAGNL, X, P, R, Z, ReLerr, ItMAX)
Double: The double precision name is DJCGRC.

\section*{Description}

Routine JCGRC solves the symmetric definite linear system \(A x=b\) using the Jacobi conjugate gradient method. This method is described in detail by Golub and Van Loan (1983, Chapter 10), and in Hageman and Young (1981, Chapter 7).
This routine is a special case of the routine PCGRC, with the diagonal of the matrix A used as the preconditioning matrix. For details of the algorithm see PCGRC.

The number of iterations needed depends on the matrix and the error tolerance RELERR. As a rough guide, \(\operatorname{ITMAX}=N\) is often sufficient when \(N » 1\). See the references for further information.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of J2GRC/DJ2GRC. The reference is:

CALL J2GRC (IDO, N, DIAGNL, X, P, R, Z, RELERR, ITMAX, TRI, WK, IWK)
The additional arguments are as follows:
TRI - Workspace of length 2 * ITMAX containing a tridiagonal matrix (in band symmetric form) whose largest eigenvalue is approximately the same as the largest eigenvalue of the iteration matrix. The workspace arrays TRI, WK and IWK should not be changed between the initial call with IDO \(=0\) and JCGRC/DJCGRC returning with IDO \(=2\).

WK - Workspace of length 5 * ITMAX.
IWK - Workspace of length ITMAX.
2. Informational errors

Type Code
\(4 \quad 1 \quad\) The diagonal contains a zero.
42 The diagonal elements have different signs.
43 No convergence after ITMAX iterations.
\(4 \quad 4\) The linear system is not definite.
45 The linear system is singular.

\section*{Example}

In this example, the solution to a linear system is found. The coefficient matrix \(A\) is stored as a full matrix.
```

    USE IMSL_LIBRARIES
    INTEGER LDA, N
    PARAMETER (LDA=3, N=3)
    INTEGER IDO, ITMAX
    REAL A(LDA,N), B(N), DIAGNL(N), P(N), R(N), X(N), &
        Z (N)
    ! ! A = (
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
DATA B/27.0, -78.0, 64.0/
CALL SCOPY (N, B, 1, R, 1)
CALL SCOPY (N, B, 1, X, 1)
CALL SCOPY (N, A(:, 1), LDA+1, DIAGNL, 1)
! Set parameters
ITMAX = 100
IDO = 0
10 CALL JCGRC (IDO, DIAGNL, X, P, R, Z, ITMAX=ITMAX)
IF (IDO .EQ. 1) THEN
CALL MURRV (A, P, Z)
GO TO 10
END IF
CALL WRRRN ('Solution', X)
END

```

\section*{Output}
```

Solution
1 1.001
-4.000
7.000

```

\section*{GMRES}

Uses the Generalized Minimal Residual Method with reverse communication to generate an approximate solution of \(A x=b\).

\section*{Required Arguments}

IDO - Flag indicating task to be done. (Input/Output)
On the initial call IDO must be 0 . If the routine returns with IDO \(=1\), then set \(Z=A P\), where \(A\) is the matrix, and call GMRES again. If the routine returns with IDO \(=2\), then set \(Z\) to the solution of the system \(M Z=P\), where \(M\) is the preconditioning matrix, and call GMRES again. If the routine returns with \(I D O=3\), set \(\mathrm{z}=A M^{l} P\), and call GMRES again. If the routine returns with IDO \(=4\), the iteration has converged, and X contains the approximate solution to the linear system.
\(\boldsymbol{X}\) - Array of length N containing an approximate solution. (Input/Output)
On input, x contains an initial guess of the solution. On output, x contains the approximate solution.
\(\boldsymbol{P}\) - Array of length N. (Output)
Its use is described under IDO.
\(\boldsymbol{R}\) - Array of length N. (Input/Output)
On initial input, it contains the right-hand side of the linear system. On output, it contains the residual, \(b-A x\).
\(\mathbf{Z}\) - Array of length N. (Input)
When IDO \(=1\), it contains \(A P\), where \(A\) is the coefficient matrix. When IDO \(=2\), it contains \(M^{l} P\). When IDO \(=3\), it contains \(A M^{l} P\). When IDO \(=0\), it is ignored.

TOL - Stopping tolerance. (Input/Output)
The algorithm attempts to generate a solution \(x\) such that \(|b-A x| \leq T O L \star|b|\). On output, TOL contains the final residual norm.

\section*{Optional Arguments}
\(N\) - Order of the linear system. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL GMRES (IDO, X, P, R, Z, TOL [, ...])
Specific: The specific interface names are S_GMRES and D_GMRES.

\section*{FORTRAN 77 Interface}

Single: CALL GMRES (IDO, N, X, P, R, Z, TOL)
Double: The double precision name is DGMRES.

\section*{Description}

The routine GMRES implements restarted GMRES with reverse communication to generate an approximate solution to \(A x=b\). It is based on GMRESD by Homer Walker.
There are four distinct GMRES implementations, selectable through the parameter vector INFO. The first Gram-Schmidt implementation, \(\operatorname{INFO}(1)=1\), is essentially the original algorithm by Saad and Schultz (1986). The second Gram-Schmidt implementation, developed by Homer Walker and Lou Zhou, is simpler than the first implementation. The least squares problem is constructed in upper-triangular form and the residual vector updating at the end of a GMRES cycle is cheaper. The first Householder implementation is algorithm 2.2 of Walker (1988), but with more efficient correction accumulation at the end of each GMRES cycle. The second Householder implementation is algorithm 3.1 of Walker (1988). The products of Householder transformations are expanded as sums, allowing most work to be formulated as large scale matrix-vector operations. Although BLAS are used wherever possible, extensive use of Level 2 BLAS in the second Householder implementation may yield a performance advantage on certain computing environments.

The Gram-Schmidt implementations are less expensive than the Householder, the latter requiring about twice as much arithmetic beyond the coefficient matrix/vector products. However, the Householder implementations may be more reliable near the limits of residual reduction. See Walker (1988) for details. Issues such as the cost of coefficient matrix/vector products, availability of effective preconditioners, and features of particular computing environments may serve to mitigate the extra expense of the Householder implementations.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of G2RES/DG2RES. The reference is:
```

CALL G2RES (IDO, N, X, P, R, Z, TOL, INFO, USRNPR, USRNRM,
WORK)

```

The additional arguments are as follows:
INFO - Integer vector of length 10 used to change parameters of GMRES.
(Input/Output).

For any components INFO (1) ... INFO (7) with value zero on input, the default value is used.
\(\operatorname{INFO}(1)=\) IMP, the flag indicating the desired implementation.
\begin{tabular}{cl} 
IMP & Action \\
1 & first Gram-Schmidt implementation \\
2 & second Gram-Schmidt implementation \\
3 & first Householder implementation \\
4 & second Householder implementation
\end{tabular}

Default: \(\operatorname{IMP}=1\)
INFO (2) = KDMAX, the maximum Krylor subspace dimension, i.e., the maximum allowable number of GMRES iterations before restarting. It must satisfy \(1 \leq K \operatorname{LMAX} \leq N\).
Default: \(\operatorname{KDMAX}=\min (\mathrm{N}, 20)\)
\(\operatorname{INFO}(3)=\) ITMAX, the maximum number of GMRES iterations allowed.
Default: \(\operatorname{ITMAX}=1000\)
\(\operatorname{INFO}(4)=\operatorname{IRP}\), the flag indicating whether right preconditioning is used. If \(\operatorname{IRP}=0\), no right preconditioning is performed. If \(I R P=1\), right preconditioning is performed. If \(\operatorname{IRP}=0\), then \(\operatorname{IDO}=2\) or 3 will not occur.
Default: \(\operatorname{IRP}=0\)
\(\operatorname{INFO}(5)=\) IRESUP, the flag that indicates the desired residual vector updating prior to restarting or on termination.
\begin{tabular}{ll} 
IRESUP & Action \\
1 & update by linear combination, restarting only \\
2 & update by linear combination, restarting and termination \\
3 & update by direct evaluation, restarting only \\
4 & update by direct evaluation, restarting and termination
\end{tabular}

Updating by direct evaluation requires an otherwise unnecessary matrixvector product. The alternative is to update by forming a linear combination of various available vectors. This may or may not be cheaper and may be less reliable if the residual vector has been greatly reduced. If IRESUP \(=2\) or 4 , then the residual vector is returned in WORK (1) , ... WORK (N). This is useful in some applications but costs another unnecessary residual update. It is recommended that \(\operatorname{IRESUP}=1\) or 2 be used, unless matrix-vector products are inexpensive or great residual reduction is required. In this case use IRESUP \(=3\) or 4 . The meaning of "inexpensive" varies with IMP as follows:
\begin{tabular}{ll} 
IMP & \(\leq\) \\
1 & \((\) KDMAX +1\() \star_{\mathrm{N}}\) flops \\
2 & N flops \\
3 & \((2 \star\) KDMAX +1\() \star_{\mathrm{N}}\) flops \\
4 & \((2 \star\) KDMAX +1\() \star_{\mathrm{N}}\) flops
\end{tabular}
"Great residual reduction" means that TOL is only a few orders of magnitude larger than machine epsilon.
Default: \(\operatorname{IRESUP}=1\)
\(\operatorname{INFO}(6)=\) flag for indicating the inner product and norm used in the GramSchmidt implementations. If \(\operatorname{INFO}(6)=0\), sdot and snrm2, from BLAS, are used. If \(\operatorname{INFO}(6)=1\), the user must provide the routines, as specified under arguments USRNPR and USRNRM.
Default: \(\operatorname{INFO}(6)=0\)
\(\operatorname{INFO}(7)=\operatorname{IPRINT}\), the print flag. If IPRINT \(=0\), no printing is performed. If IPRINT \(=1\), print the iteration numbers and residuals.
Default: \(\operatorname{IPRINT}=0\)
\(\operatorname{INFO}(8)=\) the total number of GMRES iterations on output.
\(\operatorname{INFO}(9)=\) the total number of matrix-vector products in GMRES on output.
\(\operatorname{INFO}(10)=\) the total number of right preconditioner solves in GMRES on output if \(\operatorname{IRP}=1\).

USRNPR - User-supplied FUNCTION to use as the inner product in the Gram-Schmidt implementation, if \(\operatorname{INFO}(6)=1\). If \(\operatorname{INFO}(6)=0\), the dummy function G8RES/DG8RES may be used. The usage is

REAL FUNCTION USRNPR (N, SX, INCX, SY, INCY)
N - Length of vectors X and Y . (Input)
SX - Real vector of length MAX(N*IABS (INCX) , 1). (Input)
INCX - Displacement between elements of SX. (Input)
\(X(I)\) is defined to be \(S X(1+(I-1) * I N C X)\) if INCX is greater than 0 , or \(S X(1+(I-N) * I N C X)\) if INCX is less than 0 .

SY - Real vector of length MAX ( \(N *\) IABS (INXY), 1). (Input)
INCY - Displacement between elements of SY. (Input)
\(Y(I)\) is defined to be \(S Y(1+(I-1) *\) INCY \()\) if INCY is greater than 0 , or \(S Y(1+(I-N) * I N C Y)\) if INCY is less than zero.
USRNPR must be declared EXTERNAL in the calling program.
USRNRM - User-supplied FUNCTION to use as the norm \(\|X\|\) in the Gram-Schmidt implementation, if \(\operatorname{INFO}(6)=1\). If \(\operatorname{INFO}(6)=0\), the dummy function G9RES/DG9RES may be used.The usage is
```

REAL FUNCTION USRNRM (N, SX, INCX)

```

N - Length of vectors X and Y . (Input)
SX - Real vector of length MAX ( \(\mathrm{N} *\) IABS (INCX) , 1). (Input)
INCX - Displacement between elements of SX. (Input)
\(X(I)\) is defined to be \(S X(1+(I-1) * I N C X)\) if INCX is greater than 0 , or \(S X(1+(I-N) * I N C X)\) if INCX is less than 0.
USRNRM must be declared EXTERNAL in the calling program.
WORK - Work array whose length is dependent on the chosen implementation.
```

IMP length of WORK
1 N* (KDMAX + 2) + KDMAX**2 + 3 *KDMAX + 2
2 N* (KDMAX + 2) + KDMAX**2 + 2 *KDMAX + 1
3* (KDMAX + 2) + 3 *KDMAX + 2
N* (KDMAX + 2) + KDMAX** 2 + 2 *KDMAX + 2

```

\section*{Example 1}

This is a simple example of GMRES usage. A solution to a small linear system is found. The coefficient matrix \(A\) is stored as a full matrix, and no preconditioning is used. Typically, preconditioning is required to achieve convergence in a reasonable number of iterations.
```

USE IMSL_LIBRARIES
Declare variables
INTEGER LDA, N
PARAMETER (N=3, LDA=N)
Specifications for local variables
REAL P(N), TOL, X(N), Z(N)
REAL A(LDA,N), R(N)
SAVE A, R
INTRINSIC SQRT
REAL SQRT
( 33.0
A = (-24.0 -10.0 -57.0)
( 18.0 -11.0 7.0)
B = (129.0 -96.0 8.5)
DATA A/33.0, -24.0, 18.0, 16.0, -10.0, -11.0, 72.0, -57.0, 7.0/
DATA R/129.0, -96.0, 8.5/
CALL UMACH (2, NOUT)
Initial guess = (0 ... 0)
X = 0.0EO
TOL = AMACH (4)
Set stopping tolerance to
square root of machine epsilon

```
```

        TOL = SQRT(TOL)
        IDO = 0
    1 0 ~ C O N T I N U E ~
    CALL GMRES (IDO, X, P, R, Z, TOL)
    IF (IDO .EQ. 1) THEN
    ! Set z = A*p
CALL MURRV (A, P, Z)
GO TO 10
END IF
!
CALL WRRRN ('Solution', X, 1, N, 1)
WRITE (NOUT,'(A11, E15.5)') 'Residual = ', TOL
END

```

\section*{Output}
```

        Solution
    1 2 3
    1.000 1.500 1.000
Residual = 0.29746E-05

```

\section*{Additional Examples}

\section*{Example 2}

This example solves a linear system with a coefficient matrix stored in coordinate form, the same problem as in the document example for LSLXG. Jacobi preconditioning is used, i.e. the preconditioning matrix \(M\) is the diagonal matrix with \(M_{i i}=A_{i i}\), for \(i=1, \ldots, n\).
```

USE IMSL_LIBRARIES
INTEGER N, NZ
PARAMETER (N=6, NZ=15)
INTEGER IDO, INFO(10), NOUT
REAL P(N), TOL, WORK(1000), X(N), Z(N)
REAL DIAGIN(N), R(N)
INTRINSIC SQRT
REAL SQRT
EXTERNAL AMULTP
EXTERNAL G8RES, G9RES
DATA DIAGIN/0.1, 0.1, 0.0666667, 0.1, 1.0, 0.16666667/
DATA R/10.0, 7.0, 45.0, 33.0, -34.0, 31.0/
CALL UMACH (2, NOUT)
X = 1.0E0
Set up the options vector INFO
to use preconditioning

```
```

INFO = 0
INFO(4) = 1
TOL = AMACH (4)
TOL = SQRT (TOL)
IDO = 0
10 CONTINUE
CALL G2RES (IDO, N, X, P, R, Z, TOL, INFO, G8RES, G9RES, WORK)
IF (IDO .EQ. 1) THEN
Set z = A*p
CALL AMULTP (P, Z)
GO TO 10
ELSE IF (IDO .EQ. 2) THEN
Set z = inv(M)*p
The diagonal of inv(M) is stored
in DIAGIN
CALL SHPROD (N, DIAGIN, 1, P, 1, Z, 1)
GO TO 10
ELSE IF (IDO .EQ. 3) THEN
Set z = A*inv(M)*p
CALL SHPROD (N, DIAGIN, 1, P, 1, Z, 1)
P = Z
CALL AMULTP (P, Z)
GO TO 10
END IF
CALL WRRRN ('Solution', X)
WRITE (NOUT,'(A11, E15.5)') 'Residual = ', TOL
END
SUBROUTINE AMULTP (P, Z)
USE IMSL_LIBRARIES
INTEGER - NZ
PARAMETER (NZ=15)
REAL P(*), Z(*)
INTEGER N
PARAMETER (N=6)
INTEGER I
INTEGER IROW (NZ), JCOL (NZ)
REAL A(NZ)
SAVE A, IROW, JCOL
SPECIFICATIONS FOR SUBROUTINES
Define the matrix A
DATA A/6.0, 10.0, 15.0, -3.0, 10.0, -1.0, -1.0, -3.0, -5.0, 1.0, \&
10.0, -1.0, -2.0, -1.0, -2.0/
DATA IROW/6, 2, 3, 2, 4, 4, 5, 5, 5, 5, 1, 6, 6, 2, 4/

```
\(!\)
```

    DATA JCOL/6, 2, 3, 3, 4, 5, 1, 6, 4, 5, 1, 1, 2, 4, 1/
    !
CALL SSET(N, 0.0, Z, 1)
! Accumulate the product A*p in z
DO 10 I=1, NZ
Z(IROW(I)) = Z(IROW(I)) + A(I)*P(JCOL(I))
10 CONTINUE
RETURN
END

```

\section*{Output}
```

Solution

```
11.000
\(2 \quad 2.000\)
33.000
    4.000
    5.000
    6.000
Residual \(=0.25882 \mathrm{E}-05\)

\section*{Example 3}

The coefficient matrix in this example corresponds to the five-point discretization of the 2-d Poisson equation with the Dirichlet boundary condition. Assuming the natural ordering of the unknowns, and moving all boundary terms to the right hand side, we obtain the block tridiagonal matrix
\[
A=\left[\begin{array}{cccc}
T & -I & & \\
-I & \ddots & \ddots & \\
& \ddots & \ddots & -I \\
& & -I & T
\end{array}\right]
\]
where
\[
T=\left[\begin{array}{cccc}
4 & -1 & & \\
-1 & \ddots & \ddots & \\
& \ddots & \ddots & -1 \\
& & -1 & 4
\end{array}\right]
\]
and \(I\) is the identity matrix. Discretizing on a \(k \times k\) grid implies that \(T\) and \(I\) are both \(k \times k\), and thus the coefficient matrix \(A\) is \(k^{2} \times k^{2}\).

The problem is solved twice, with discretization on a \(50 \times 50\) grid. During both solutions, use the second Householder implementation to take advantage of the large scale matrix/vector operations done in Level 2 BLAS. Also choose to update the residual vector by direct evaluation since the small tolerance will require large residual reduction.
The first solution uses no preconditioning. For the second solution, we construct a block diagonal preconditioning matrix
\[
M=\left[\begin{array}{lll}
T & & \\
& \ddots & \\
& & T
\end{array}\right]
\]
\(M\) is factored once, and these factors are used in the forward solves and back substitutions necessary when GMRES returns with \(I D O=2\) or 3 .
Timings are obtained for both solutions, and the ratio of the time for the solution with no preconditioning to the time for the solution with preconditioning is printed. Though the exact results are machine dependent, we see that the savings realized by faster convergence from using a preconditioner exceed the cost of factoring M and performing repeated forward and back solves.
```

USE IMSL_LIBRARIES
INTEGER K, N
PARAMETER (K=50, N=K*K)
INTEGER IDO, INFO(10), IR(20), IS(20), NOUT
REAL A (2*N), B (2*N), C (2*N), G8RES, G9RES, P (2*N), R(N), \&
TNOPRE, TOL, TPRE, U(2*N), WORK(100000), X(N), \&
Y(2*N), Z(2*N)
Specifications for subroutines
EXTERNAL AMULTP, G8RES, G9RES
CALL UMACH (2, NOUT)
Specifications for functions
Right hand side and initial guess
to (1 ... 1)
R = 1.0E0
X = 1.0E0
Use the 2nd Householder
implementation and update the
residual by direct evaluation
INFO = 0
INFO(1) = 4
INFO(5) = 3
TOL = AMACH (4)
TOL = 100.0*TOL
IDO = 0
! Time the solution with no
!
TNOPRE = CPSEC()
10 CONTINUE
CALL G2RES (IDO, N, X, P, R, Z, TOL, INFO, G8RES, G9RES, WORK)
IF (IDO .EQ. 1) THEN
Set z = A*p
CALL AMULTP (K, P, Z)
GO TO 10
END IF
TNOPRE = CPSEC() - TNOPRE
!
WRITE (NOUT,'(A32, I4)') 'Iterations, no preconditioner = ', \&
INFO(8)

```
!
                                    Solve again using the diagonal blocks
                                    of A as the preconditioning matrix M
    R = 1.0E0
    X = 1.0E0
!
    CALL SSET (N-1, -1.0, B, 1)
    CALL SSET (N-1, -1.0, C, 1)
        CALL SSET (N, 4.0, A, 1)
        INFO(4) = 1
        TOL = AMACH (4)
        TOL = 100.0*TOL
        IDO = 0
        TPRE = CPSEC()
    CALL LSLCR (C, A, B, Y, U, IR, IS, IJOB=6)
    20 CONTINUE
    CALL G2RES (IDO, N, X, P, R, Z, TOL, INFO, G8RES, G9RES, WORK)
        IF (IDO .EQ. 1) THEN
                                    Set z = A*p
    CALL AMULTP (K, P, Z)
    GO TO 20
ELSE IF (IDO .EQ. 2) THEN
                                    Set z = inv(M)*p
    CALL SCOPY (N, P, 1, Z, 1)
    CALL LSLCR (C, A, B, Z, U, IR, IS, IJOB=5)
    GO TO 20
ELSE IF (IDO .EQ. 3) THEN
                    Set z = A*inv(M)*p
    CALL LSLCR (C, A, B, P, U, IR, IS, IJOB=5)
    CALL AMULTP (K, P, Z)
    GO TO 20
END IF
TPRE = CPSEC() - TPRE
WRITE (NOUT,'(A35, I4)') 'Iterations, with preconditioning = ',&
INFO(8)
WRITE (NOUT,'(A45, F10.5)') '(Precondition time)/(No '// &
                    'precondition time) = ', TPRE/TNOPRE
END
SUBROUTINE AMULTP (K, P, Z)
USE IMSL_LIBRARIES
INTEGER K
REAL P(*), Z(*)
Specifications for local variables
INTEGER I, N
!
```

```
    N = K*K
    CALL SVCAL (N, 4.0, P, 1, Z, 1)
    CALL SAXPY (N-1, -1.0, P(2:(N-1)), 1, Z, 1)
    CALL SAXPY (N-1, -1.0, P, 1, Z(2:(N-1)), 1)
                                    Correct for terms not properly in
                                    block diagonal
    DO 10 I=K, N - K, K
        Z(I) = Z(I) + P(I+1)
        Z(I+1) = Z(I+1) + P(I)
    10 CONTINUE
    Do the super and subdiagonal blocks,
    the -I's
    CALL SAXPY (N-K, -1.0, P(( K+1): (N-K)), 1, Z, 1)
    CALL SAXPY (N-K, -1.0, P, 1, Z((K+1):(N-K)), 1)
    RETURN
    END
```

$!$
$!$

## Output

```
Iterations, no preconditioner = 329
```

Iterations, with preconditioning = 192
(Precondition time) $/($ No precondition time) $=0.66278$

## LSQRR

Solves a linear least-squares problem without iterative refinement.

## Required Arguments

$\boldsymbol{A}$ - NRA by NCA matrix containing the coefficient matrix of the least-squares system to be solved. (Input)
$\boldsymbol{B}$ - Vector of length NRA containing the right-hand side of the least-squares system. (Input)
$\boldsymbol{X}$ - Vector of length NCA containing the solution vector with components corresponding to the columns not used set to zero. (Output)
$\boldsymbol{R E S}$ - Vector of length NRA containing the residual vector $\mathrm{B}-\mathrm{A}$ * X . (Output)
KBASIS - Scalar containing the number of columns used in the solution.

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{size}(A, 1)$.
$N C A$ - Number of columns of A. (Input)
Default: NCA = size (A,2).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
TOL - Scalar containing the nonnegative tolerance used to determine the subset of columns of $A$ to be included in the solution. (Input)
If TOL is zero, a full complement of $\min (N R A, N C A)$ columns is used. See Comments.
Default: TOL $=0.0$

## FORTRAN 90 Interface

Generic: CALL LSQRR (A, B, X, RES, KBASIS [,...])

Specific: The specific interface names are $S_{-} L S Q R R$ and D_LSQRR.

## FORTRAN 77 Interface

Single: CALL LSQRR (NRA, NCA, A, LDA, B, TOL, X, RES, KBASIS)
Double: The double precision name is DLSQRR.

## ScaLAPACK Interface

Generic: CALL LSQRR (A0, B0, X0, RESO, KBASIS [, ...])

Specific: The specific interface names are S_LSQRR and D_LSQRR.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

Routine LSQRR solves the linear least-squares problem. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The routine LQRRR is first used to compute the $Q R$ decomposition of $A$. Pivoting, with all rows free, is used. Column $k$ is in the basis if

$$
\left|R_{k k}\right| \leq \tau\left|R_{11}\right|
$$

with $\tau=$ TOL. The truncated least-squares problem is then solved using IMSL routine LQRSL. Finally, the components in the solution, with the same index as columns that are not in the basis, are set to zero; and then, the permutation determined by the pivoting in IMSL routine LQRRR is applied.

## Comments

1. Workspace may be explicitly provided, if desired, by use of $L 2 Q R R / D L 2 Q R R$. The reference is:

CALL L2QRR (NRA, NCA, A, LDA, B, TOL, X, RES, KBASIS, QR, QRAUX, IPVT, WORK)
The additional arguments are as follows:
$\boldsymbol{Q R}$ - Work vector of length NRA * NCA representing an NRA by NCA matrix that contains information from the $Q R$ factorization of A. The upper trapezoidal part of $Q R$ contains the upper trapezoidal part of $R$ with its diagonal elements ordered in decreasing magnitude. The strict lower trapezoidal part of $Q R$ contains information to recover the orthogonal matrix $Q$ of the factorization. If $A$ is not needed, $Q R$ can share the same storage locations as A.

QRAUX - Work vector of length NCA containing information about the orthogonal factor of the $Q R$ factorization of $A$.

IPVT - Integer work vector of length NCA containing the pivoting information for the $Q R$ factorization of A.

WORK - Work vector of length 2 * NCA -1.
2. Routine LSQRR calculates the $Q R$ decomposition with pivoting of a matrix A and tests the diagonal elements against a user-supplied tolerance TOL. The first integer KBASIS $=k$ is determined for which

$$
\left|r_{k+1, k+1}\right| \leq \text { TOL } *\left|r_{11}\right|
$$

In effect, this condition implies that a set of columns with a condition number approximately bounded by $1.0 / \mathrm{TOL}$ is used. Then, LQRSL performs a truncated fit of the first KBASIS columns of the permuted A to an input vector B. The coefficient of this fit is unscrambled to correspond to the original columns of $A$, and the coefficients corresponding to unused columns are set to zero. It may be helpful to scale the rows and columns of A so that the error estimates in the elements of the scaled matrix are roughly equal to TOL.
3. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine $L 2 Q R R$ the leading dimension of $Q R$ is increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are
temporarily replaced by $\operatorname{IVAL}(1)$ and $\operatorname{IVAL}(2)$, respectively, in LSQRR. Additional memory allocation for $Q R$ and option value restoration are done automatically in LSQRR. Users directly calling L2QRR can allocate additional space for $Q R$ and set $\operatorname{IVAL}(3)$ and $\operatorname{IVAL}(4)$ so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSQRR or L2QRR. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1$.

17 This option has two values that determine if the $L_{1}$ condition number is to be computed. Routine LSQRR temporarily replaces IVAL(2) by IVAL(1). The routine L2CRG computes the condition number if IVAL $(2)=2$. Otherwise L2CRG skips this computation. LSQRR restores the option. Default values for the option are $\operatorname{IVAL}(*)=1,2$.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:
$\boldsymbol{A O}$ - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the coefficient matrix of the least squares system to be solved. (Input)

B0 - Local vector of length MXLDA containing the local portions of the distributed vector B. $B$ contains the right-hand side of the least squares system. (Input)

X0 - Local vector of length MXLDX containing the local portions of the distributed vector X. X contains the solution vector with components corresponding to the columns not used set to zero. (Output)

RESO - Local vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the residual vector $B-A * X$. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXLDX, and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

## Example

Consider the problem of finding the coefficients $c_{i}$ in

$$
f(x)=c_{0}+c_{1} x+c_{2} x^{2}
$$

given data at $x=1,2,3$ and 4 , using the method of least squares. The row of the matrix $A$ contains the value of $1, x$ and $x^{2}$ at the data points. The vector $b$ contains the data, chosen such that $c_{0} \approx 1, c_{1} \approx 2$ and $c_{2} \approx 0$. The routine $L S Q R R$ solves this least-squares problem.

```
USE LSQRR_INT
USE UMACH_INT
USE WRRRN__INT
```

```
! Declare variables
    PARAMETER (NRA=4, NCA=3, LDA=NRA)
REAL A(LDA,NCA), B(NRA), X(NCA), RES (NRA), TOL
    Set values for A
    A=(\begin{array}{llll}{1}&{2}&{4}\end{array})
        (\begin{array}{llll}{1}&{4}&{16}\\{(}&{6}&{36}\end{array})
        ( 1 1 8 % 64 )
    DATA A/4*1.0, 2.0, 4.0, 6.0, 8.0, 4.0, 16.0, 36.0, 64.0/
        Set values for B
    DATA B/ 4.999, 9.001, 12.999, 17.001 /
    Solve the least squares problem
TOL = 1.0E-4
CALL LSQRR (A, B, X, RES, KBASIS, TOL=TOL)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'KBASIS = ', KBASIS
CALL WRRRN ('X', X, 1, NCA, 1)
CALL WRRRN ('RES', RES, 1, NRA, 1)
!
END
```


## Output

```
KBASIS = 3
```

|  | $X$ |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| 0.999 | 2.000 | 0.000 |


| 1 | 2 | 3 | 4 |
| ---: | ---: | ---: | ---: |
| -0.000400 | 0.001200 | $-0.001200^{2}$ | 0.000400 |

## ScaLAPACK Example

The previous example is repeated here as a distributed computing example. Consider the problem of finding the coefficients $c_{i}$ in

$$
f(x)=c_{0}+c_{1} x+c_{2} x^{2}
$$

given data at $x=1,2,3$ and 4 , using the method of least squares. The row of the matrix $A$ contains the value of $1, x$ and $x^{2}$ at the data points. The vector $b$ contains the data, chosen such that $c_{0} \approx 1, c_{1} \approx 2$ and $c_{2} \approx 0$. The routine LSQRR solves this least-squares problem. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.

```
USE MPI_SETUP_INT
USE LSQRR INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
                                    Declare variables
INTEGER LDA, NRA, NCA, DESCA(9), DESCX(9), DESCR(9)
INTEGER INFO, KBASIS, MXCOL, MXLDA, MXCOLX, MXLDX, NOUT
REAL TOI
REAL, ALLOCATABLE :: A(:,:), B(:), X(:), RES(:)
REAL, ALLOCATABLE :: AO(:,:), BO(:), XO(:), RESO(:)
PARAMETER (NRA=4, NCA=3, LDA=NRA)
MP NPROCS = MP SETUP()
IF(MP RANK .EQ. 0) THEN
    ALLOCATE (A(LDA,NCA), B(NRA), X(NCA), RES (NRA))
                            Set values for A and B
    A(1,:) = (/ 1.0, 2.0, 4.0/)
    A(2,:) = (/ 1.0, 4.0, 16.0/)
    A(3,:) = (/ 1.0, 6.0, 36.0/)
    A(4,:) = (/ 1.0, 8.0, 64.0/)
    B = (/4.999, 9.001, 12.999, 17.001/)
ENDIF
    Set up a 1D processor grid and define
    its context ID, MP ICTXT
CALL SCALAPACK_SETUP(NRA, NCA, .TRUE., .FALSE.)
    Get the array descriptor entities MXLDA,
    MXCOL, MXLDX, and MXCOLX
CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
CALL SCALAPACK_GETDIM(NCA, 1, MP_NB, 1, MXZLDX, MXCOLX)
    Set up the array descriptors
CALL DESCINIT(DESCA, NRA, NCA, MP MB, MP NB, 0, 0, MP ICTXT, MXLDA, &
    INFO)
CALL DESCINIT(DESCX, NCA, 1, MP_NB, 1, 0, 0, MP_ICTXT, MXLDX, INFO)
CALL DESCINIT(DESCR, NRA, 1, MP_MB, 1, 0, 0, MP_ICTXT, MXLDA, INFO)
    Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA), X0 (MXLDX), RES0 (MXLDA))
        Map input arrays to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK_MAP(B, DESCR, BO)
        Solve the least squares problem
    TOL = 1.0E-4
CALL LSQRR (A0, B0, X0, RESO, KBASIS, TOL=TOL)
                                    Unmap the results from the distributed
                                    arrays back to a non-distributed array.
                                    After the unmap, only Rank=0 has the full
                                    array.
CALL SCALAPACK UNMAP (X0, DESCX, X)
CALL SCALAPACK UNMAP(RESO, DESCR, RES)
                            Print results.
                                Only Rank=0 has the solution.
IF(MP_RANK .EQ. 0)THEN
```

```
        CALL UMACH (2, NOUT)
        WRITE (NOUT,*) `KBASIS = `, KBASIS
        CALL WRRRN ('X', X, 1, NCA, 1)
        CALL WRRRN ('RES', RES, 1, NRA, 1)
ENDIF
IF (MP RANK .EQ. O) DEALLOCATE (A, B, RES, X)
DEALLOCATE (A0, B0, RES0, X0)
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP(`FINAL')
END
```


## Output

```
KBASIS = 3
```

|  | $X$ |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| 0.999 | 2.000 | 0.000 |

    RES
    | 1 | ${ }^{2}$ | 3 | 4 |
| ---: | ---: | ---: | ---: |
| -0.000400 | 0.001200 | -0.001200 | $0.000400^{4}$ |

## LQRRV

## MPI <br> CAPABLE

Computes the least-squares solution using Householder transformations applied in blocked form.

## Required Arguments

$\boldsymbol{A}$ - Real LDA by (NCA + NUMEXC) array containing the matrix and right-hand sides. (Input) The right-hand sides are input in $A(1:$ NRA, NCA $+j), j=1, \ldots$, NUMEXC. The array A is preserved upon output. The Householder factorization of the matrix is computed and used to solve the systems.
$\boldsymbol{X}$ - Real LDX by NUMEXC array containing the solution. (Output)

## Optional Arguments

NRA - Number of rows in the matrix. (Input)
Default: NRA $=\operatorname{size}(\mathrm{A}, 1)$.
$N C A$ - Number of columns in the matrix. (Input)
Default: NCA $=\operatorname{size}(A, 2)-$ NUMEXC.

NUMEXC - Number of right-hand sides. (Input)
Default: NUMEXC $=\operatorname{size}(\mathrm{x}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
$\boldsymbol{L D} \boldsymbol{X}$ - Leading dimension of the solution array x exactly as specified in the dimension statement of the calling program. (Input)
Default: LDX $=\operatorname{size}(X, 1)$.

## FORTRAN 90 Interface

Generic: CALL LQRRV (A, $X[, \ldots]$ )
Specific: The specific interface names are $S_{-} L Q R R V$ and $D_{-} L Q R R V$.

## FORTRAN 77 Interface

Single: CALL LQRRV (NRA, NCA, NUMEXC, A, LDA, X, LDX)
Double: The double precision name is $D L Q R R V$.

## ScaLAPACK Interface

Generic: CALL LQRRV (A0, XO [,...])
Specific: The specific interface names are $S_{\_} \quad L Q R R V$ and $D \_L Q R R V$.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

The routine $L Q R R V$ computes the $Q R$ decomposition of a matrix $A$ using blocked Householder transformations. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The standard algorithm is based on the storage-efficient $W Y$ representation for products of Householder transformations. See Schreiber and Van Loan (1989).

The routine $L Q R R V$ determines an orthogonal matrix $Q$ and an upper triangular matrix $R$ such that $A=Q R$. The $Q R$ factorization of a matrix $A$ having NRA rows and NCA columns is as follows:

Initialize $A_{1} \leftarrow A$
For $k=1$, min(NRA-1, NCA)
Determine a Householder transformation for column $k$ of $A_{k}$ having the form

$$
H_{k}=I-\tau_{k} \mu_{k} \mu_{k}^{T}
$$

where $u_{k}$ has zeros in the first $k-1$ positions and $\tau_{k}$ is a scalar. Update

$$
A_{k} \leftarrow H_{k} A_{k-1}=A_{k-1}-\tau_{k} \mu_{k}\left(A_{k-1}^{T} \mu_{k}\right)^{T}
$$

End $k$
Thus,

$$
A_{p}=H_{p} H_{p-1} \cdots H_{1} A=Q^{T} A=R
$$

where $p=\min (\mathrm{NRA}-1, \mathrm{NCA})$. The matrix $Q$ is not produced directly by LQRRV. The information needed to construct the Householder transformations is saved instead. If the matrix $Q$ is needed explicitly, $Q^{T}$ can be determined while the matrix is factored. No pivoting among the columns is done. The primary purpose of $L Q R R V$ is to give the user a high-performance $Q R$ least-squares solver. It is intended for least-squares problems that are well-posed. For background, see Golub and Van Loan (1989, page 225). During the $Q R$ factorization, the most time-consuming step is computing the matrix-vector update $A_{k} \leftarrow H_{k} A_{k-1}$. The routine LQRRV constructs "block" of NB Householder transformations in which the update is "rich" in matrix multiplication. The product of nB Householder transformations are written in the form

$$
H_{k} H_{k+1} \cdots H_{k+n b-1}=I+Y T Y^{T}
$$

where $Y_{N R A \times N B}$ is a lower trapezoidal matrix and $T_{N B \times N B}$ is upper triangular. The optimal choice of the block size parameter NB varies among computer systems. Users may want to change it from its default value of 1 .

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2RRV/DL2RRV. The reference is:

CALL L2RRV (NRA, NCA, NUMEXC, A, LDA, X, LDX, FACT, LDFACT, WK)
The additional arguments are as follows:
$\boldsymbol{F A C T}-\operatorname{LDFACT} \times(\mathrm{NCA}+\mathrm{NUMEXC})$ work array containing the Householder factorization of the matrix on output. If the input data is not needed, A and FACT can share the same storage locations.

LDFACT - Leading dimension of the array FACT exactly as specified in the dimension statement of the calling program. (Input) If $A$ and $F A C T$ are sharing the same storage, then $L D A=$ LDFACT is required.
$\boldsymbol{W K}$ - Work vector of length (NCA + NUMEXC +1 ) * $(\mathrm{NB}+1)$. The default value is $N B=1$. This value can be reset. See item 3 below.
2. Informational errors

$$
\begin{array}{cc}
\text { Type } & \text { Code } \\
4 & 1 \text { The input matrix is singular. }
\end{array}
$$

3. Integer Options with Chapter 11 Options Manager

5 This option allows the user to reset the blocking factor used in computing the factorization. On some computers, changing IVAL(*) to a value larger than 1 will result in greater efficiency. The value $\operatorname{IVAL}(*)$ is the maximum value to use. (The software is specialized so that IVAL (*) is reset to an "optimal" used value within routine L2RRV.) The user can control the blocking by resetting IVAL (*) to a smaller value than the default. Default values are $\operatorname{IVAL}(*)=1$, $\operatorname{IMACH}(5)$.

6 This option is the vector dimension where a shift is made from in-line level-2 loops to the use of level-2 BLAS in forming the partial product of Householder transformations. Default value is IVAL $(*)=\operatorname{IMACH}(5)$.

10 This option allows the user to control the factorization step. If the value is 1 the Householder factorization will be computed. If the value is 2 , the factorization will not be computed. In this latter case the decomposition has already been computed. Default value is $\operatorname{IVAL}(*)=1$.

11 This option allows the user to control the solving steps. The rules for IVAL(*) are:

1. Compute $b \leftarrow Q^{T} b$, and $x \leftarrow R^{+} b$.
2. Compute $b \leftarrow Q^{T} b$.
3. Compute $b \leftarrow Q b$.
4. Compute $x \leftarrow R^{+} b$.

Default value is IVAL $(*)=1$. Note that IVAL $(*)=2$ or 3 may only be set when calling L2RRV/DL2RRV.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix
A. A contains the matrix and right-hand sides. (Input)

The right-hand sides are input in $A(1:$ NRA, NCA $+j), j=1, \ldots$, NUMEXC. The array A is preserved upon output. The Householder factorization of the matrix is computed and used to solve the systems.. (Input)

X0 - MXLDX by MXCOLX local matrix containing the local portions of the distributed matrix $\mathrm{X} . \mathrm{X}$ contains the solution. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXLDX, MXCOL, and MXCOLX can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

## Example

Given a real $m \times k$ matrix $B$ it is often necessary to compute the $k$ least-squares solutions of the linear system $A X=B$, where $A$ is an $m \times n$ real matrix. When $m>n$ the system is considered overdetermined. A solution with a zero residual normally does not exist. Instead the minimization problem

$$
\min _{x_{j} \in \mathbf{R}^{\|}}\left\|A x_{j}-b_{j}\right\|_{2}
$$

is solved $k$ times where $x_{j}, b_{j}$ are the $j$-th columns of the matrices $X, B$ respectively. When $A$ is of full column rank there exits a unique solution $X_{L S}$ that solves the above minimization problem. By using the routine $L Q R R V, X_{L S}$ is computed.

```
USE LQRRV INT
USE WRRRN INT
USE SGEMM_INT
Declare variables
INTEGER LDA, LDX, NCA, NRA, NUMEXC
PARAMETER (NCA=3, NRA=5, NUMEXC=2, LDA=NRA, LDX=NCA)
    SPECIFICATIONS FOR LOCAL VARIABLES
REAL X(LDX,NUMEXC)
    SPECIFICATIONS FOR SAVE VARIABLES
REAL A(LDA,NCA+NUMEXC)
SAVE A
    SPECIFICATIONS FOR SUBROUTINES
    Set values for A and the
    righthand sides.
    A = (\begin{array}{llllllll}{1}&{2}&{4}&{|}&{7}&{10}\end{array})
        (\begin{array}{lllllll}{1}&{4}&{16}&{21}&{10}\end{array})
        (\begin{array}{lllllll}{1}&{6}&{36}&{43}&{9}\end{array})
        (\begin{array}{lllllll}{1}&{8}&{64}&{73}&{10}\end{array})
        (\begin{array}{lllllll}{1}&{10}&{100}&{111}&{10)}\end{array})
DATA A/5*1.0, 2.0, 4.0, 6.0, 8.0, 10.0, 4.0, 16.0, 36.0, 64.0, &
    100.0, 7.0, 21.0, 43.0, 73.0, 111.0, 2*10., 9., 2*10./
    QR factorization and solution
    CALL LQRRV (A, X)
    CALL WRRRN ('SOLUTIONS 1-2', X)
                            Compute residuals and print
CALL SGEMM ('N', 'N', NRA, NUMEXC, NCA, 1.E0, A, LDA, X, LDX, &
        -1.E0, A(1:, (NCA+1):),LDA)
CALL WRRRN ('RESIDUALS 1-2', A(1:,(NCA+1):))
END
```


## Output

```
SOLUTIONS 1-2
    1 2
```

```
    1.00 10.80
    1.00 -0.43
    1.00 0.04
RESIDUALS 1-2
    1 2
    0.0000 0.0857
    0.0000 -0.3429
    0.0000 0.5143
    0.0000-0.3429
    0.0000 0.0857
```


## ScaLAPACK Example

The previous example is repeated here as a distributed computing example. Given a real $m \times k$ matrix $B$ it is often necessary to compute the $k$ least-squares solutions of the linear system $A X=B$, where $A$ is an $m \times n$ real matrix. When $m>n$ the system is considered overdetermined. A solution with a zero residual normally does not exist. Instead the minimization problem

$$
\min _{x_{j} \in \mathbf{R}^{n}}\left\|A x_{j}-b_{j}\right\|_{2}
$$

is solved $k$ times where $x_{j}, b_{j}$ are the $j$-th columns of the matrices $X, B$ respectively. When $A$ is of full column rank there exits a unique solution $X_{L S}$ that solves the above minimization problem. By using the routine LQRRV, $X_{L S}$ is computed. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.

```
USE MPI SETUP INT
USE LQR\overline{RV} INT
USE SGEMM_INT
USE WRRRN-INT
USE SCALA\overline{PACK_SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
Declare variables
INTEGER LDA, LDX, NCA, NRA, NUMEXC, DESCA(9), DESCX(9)
INTEGER INFO, MXCOL, MXLDA, MXLDX, MXCOLX
INTEGER K
REAL, ALLOCATABLE :: A(:,:), X(:)
REAL, ALLOCATABLE :: A0(:,:), X0(:)
PARAMETER (NRA=5, NCA=3, NUMEXC=2, LDA=NRA, LDX=NCA)
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
    A\overline{LLOCATE (A (LDA,NCA+NUMEXC), X(LDX, NUMEXC))}
                                    Set values for A and the righthand sides
    A(1,:) = (/ 1.0, 2.0, 4.0, 7.0, 10.0/)
    A(2,:) = (/ 1.0, 4.0, 16.0, 21.0, 10.0/)
    A(3,:) = (/ 1.0, 6.0, 36.0, 43.0, 9.0/)
    A(4,:) = (/ 1.0, 8.0, 64.0, 73.0, 10.0/)
    A(5,:) = (/ 1.0, 10.0, 100.0, 111.0, 10.0/)
ENDIF
```

! Set up a 1D processor grid and define its context ID, MP_ICTXT
CALL SCALAPACK_SETUP (NRA, NCA+NUMEXC, .TRUE., .TRUE.)
Get the array descriptor entities MXLDA, and MXCOL
CALL SCALAPACK_GETDIM(NRA, NCA+NUMEXC, MP_MB, MP_NB, MXLDA, MXCOL)
Set up the array descriptors
CALL DESCINIT (DESCA, NRA, NCA+NUMEXC, MP_MB, MP_NB, 0, 0, MP_ICTXT, \& MXLDA, INFO)
$K=$ MINO (NRA, NCA)
Need to get dimensions of local $x$
separate since x's leading
dimension differs from A's
Get the array descriptor entities
MXLDX, AND MXCOLX
CALL SCALAPACK_GETDIM(K, NUMEXC, MP_MB, MP_NB, MXLDX, MXCOLX)
CALL DESCINIT $\overline{(D E S C X}^{( } \mathrm{K}, \mathrm{NUMEXC}, \mathrm{MP}{ }^{-} \mathrm{NB}, \mathrm{MP}^{-} \mathrm{NB}, 0,0, \mathrm{MP}$ ICTXT, \&
MXLDX, INFO)
ALLOCATE (AO (MXLDA, MXCOL), XO (MXLDX, MXCOLX))
Map input array to the processor grid
CALL SCALAPACK_MAP (A, DESCA, A0)
Solve the least squares problem
CALL LQRRV (A0, X0)
Unmap the results from the distributed
arrays back to a non-distributed array.
After the unmap, only Rank=0 has the full
array.
CALL SCALAPACK_UNMAP (X0, DESCX, X)
Print results.
Only Rank=0 has the solution, X.
IF (MP_RANK .EQ. O) THEN
CALL WRRRN ('SOLUTIONS 1-2', X)
Compute residuals and print
CALL SGEMM ('N', 'N', NRA, NUMEXC, NCA, 1.EO, A, LDA, X, LDX, \& $-1 . E 0, \mathrm{~A}(1:,(N C A+1):), L D A)$
CALL WRRRN ('RESIDUALS $1-2$ ', A(1:, (NCA+1):))
ENDIF
Exit ScaLAPACK usage
CALL SCALAPACK_EXIT (MP_ICTXT)
MP_NPROCS = MP_SETUP('FINAL')
END

## LSBRR

Solves a linear least-squares problem with iterative refinement.

## Required Arguments

A-Real NRA by NCA matrix containing the coefficient matrix of the least-squares system to be solved. (Input)
$\boldsymbol{B}$ - Real vector of length NRA containing the right-hand side of the least-squares system. (Input)
$\boldsymbol{X}$ - Real vector of length NCA containing the solution vector with components corresponding to the columns not used set to zero. (Output)

## Optional Arguments

$N R A$ - Number of rows of A. (Input)
Default: NRA $=\operatorname{size}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{size}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
TOL - Real scalar containing the nonnegative tolerance used to determine the subset of columns of A to be included in the solution. (Input)
If TOL is zero, a full complement of $\min (N R A, N C A)$ columns is used. See Comments. Default: TOL $=0.0$
$\boldsymbol{R E S}$ - Real vector of length NRA containing the residual vector $B-A X$. (Output)
KBASIS - Integer scalar containing the number of columns used in the solution. (Output)

## FORTRAN 90 Interface

Generic: CALL LSBRR (A, B, X [, ...])
Specific: The specific interface names are $S_{\text {_ }}$ LSBRR and D_LSBRR.

## FORTRAN 77 Interface

Single: CALL LSBRR (NRA, NCA, A, LDA, B, TOL, X, RES, KBASIS)
Double: The double precision name is DLSBRR.

## Description

Routine LSBRR solves the linear least-squares problem using iterative refinement. The iterative refinement algorithm is due to $\mathrm{Björck}(1967,1968)$. It is also described by Golub and Van Loan (1983, pages 182-183).

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2BRR/DL2BRR. The reference is:

CALL L2BRR (NRA, NCA, A, LDA, B, TOL, X, RES, KBASIS, QR, BRRUX, IPVT, WK)

The additional arguments are as follows:
$\boldsymbol{Q R}$ - Work vector of length NRA * NCA representing an NRA by NCA matrix that contains information from the $Q R$ factorization of A. See LQRRR for details.
$\boldsymbol{B R R U X}$ - Work vector of length NCA containing information about the orthogonal factor of the $Q R$ factorization of A. See $L Q R R R$ for details.

IPVT - Integer work vector of length NCA containing the pivoting information for the $Q R$ factorization of A. See LQRRR for details.
$W K-$ Work vector of length NRA +2 * NCA -1.
2. Informational error

Type Code
41 The data matrix is too ill-conditioned for iterative refinement to be effective.
3. Routine LSBRR calculates the $Q R$ decomposition with pivoting of a matrix $A$ and tests the diagonal elements against a user-supplied tolerance TOL. The first integer KBASIS $=k$ is determined for which

$$
\left|r_{k+1, k+1}\right| \leq \mathrm{TOL}^{*}\left|r_{11}\right|
$$

In effect, this condition implies that a set of columns with a condition number approximately bounded by $1.0 /$ TOL is used. Then, LQRSL performs a truncated fit of the first KBAS IS columns of the permuted A to an input vector B. The coefficient of this fit is unscrambled to correspond to the original columns of $A$, and the coefficients corresponding to unused columns are set to zero. It may be helpful to scale the rows and columns of A so that the error estimates in the elements of the scaled matrix are roughly equal to TOL. The iterative refinement method of Björck is then applied to this factorization.
4. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2BRR the leading dimension of $Q R$ is increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSBRR. Additional memory allocation for $Q R$ and option value restoration are done
automatically in LSBRR. Users directly calling L2BRR can allocate additional space for $Q R$ and set $\operatorname{IVAL}(3)$ and $\operatorname{IVAL}(4)$ so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSBRR or L2BRR. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1$.

17 This option has two valuess that determine if the $L_{1}$ condition number is to be computed. Routine LSBRR temporarily replaces IVAL(2) by IVAL(1). The routine L2CRG computes the condition number if IVAL $(2)=2$. Otherwise L2CRG skips this computation. LSBRR restores the option. Default values for the option are $\operatorname{IVAL}(*)=1,2$.

## Example

This example solves the linear least-squares problem with $A$, an $8 \times 4$ matrix. Note that the second and fourth columns of $A$ are identical. Routine LSBRR determines that there are three columns in the basis.

```
USE LSBRR_INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
PARAMETER (NRA=8, NCA=4, LDA=NRA)
REAL A(LDA,NCA), B(NRA), X(NCA), RES (NRA), TOL
    Set values for A
    A=(\begin{array}{lllll}{1}&{5}&{15}&{5}\end{array})
        (\begin{array}{llll}{1}&{4}&{17}&{4}\end{array})
        ( 1
        ( 1rlll}
        (\begin{array}{llll}{1}&{1}&{15}&{1}\end{array})
        ( 11 clll
        ( }\begin{array}{lllrl}{1}&{3}&{9}&{3}\end{array}
DATA A/8*1, 5., 4., 7., 3., 1., 8., 3., 4., 15., 17., 14., &
    18., 15., 11., 9., 10., 5., 4., 7., 3., 1., 8., 3., 4. /
    Set values for B
DATA B/ 30., 31., 35., 29., 18., 35., 20., 22./
    Solve the least squares problem
TOL = 1.0E-4
CALL LSBRR (A, B, X, TOL=TOL, RES=RES, KBASIS=KBASIS)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'KBASIS = ', KBASIS
CALL WRRRN ('X', X, 1, NCA, 1)
CALL WRRRN ('RES', RES, 1, NRA, 1)
END
```


## Output

```
KBASIS = 3
\begin{tabular}{rrrr}
1 & \(2^{X}\) & 3 & 4 \\
0.636 & 2.845 & 1.058 & 0.000
\end{tabular}
    RES
1 rrrrrrrrrrer
```


## LCLSQ

Solves a linear least-squares problem with linear constraints.

## Required Arguments

A- Matrix of dimension NRA by NCA containing the coefficients of the NRA least squares equations. (Input)
$\boldsymbol{B}$ - Vector of length NRA containing the right-hand sides of the least squares equations. (Input)
$\boldsymbol{C}$ - Matrix of dimension NCON by NCA containing the coefficients of the NCON constraints. (Input) If $\mathrm{NCON}=0, \mathrm{C}$ is not referenced.
$\boldsymbol{B L}$ - Vector of length NCON containing the lower limit of the general constraints. (Input) If there is no lower limit on the I-th constraint, then $\mathrm{BL}(\mathrm{I})$ will not be referenced.
$\boldsymbol{B U}$ - Vector of length NCON containing the upper limit of the general constraints. (Input) If there is no upper limit on the I-th constraint, then $\mathrm{BU}(\mathrm{I})$ will not be referenced. If there is no range constraint, $B L$ and $B U$ can share the same storage locations.

IRTYPE - Vector of length NCON indicating the type of constraints exclusive of simple bounds, where $\operatorname{IRTYPE}(I)=0,1,2,3$ indicates .EQ., .LE., .GE., and range constraints respectively. (Input)
$\boldsymbol{X L B}$ - Vector of length NCA containing the lower bound on the variables. (Input) If there is no lower bound on the I-th variable, then $\operatorname{XLB}(\mathrm{I})$ should be set to 1.0 E 30 .
$\boldsymbol{X U B}$ - Vector of length NCA containing the upper bound on the variables. (Input)
If there is no upper bound on the I-th variable, then $\mathrm{XUB}(\mathrm{I})$ should be set to -1.0 E 30 .
$\boldsymbol{X}$ - Vector of length NCA containing the approximate solution. (Output)

## Optional Arguments

NRA - Number of least-squares equations. (Input)
Default: NRA $=\operatorname{size}(\mathrm{A}, 1)$.
NCA - Number of variables. (Input)
Default: NCA $=\operatorname{size}(\mathrm{A}, 2)$.
NCON - Number of constraints. (Input)
Default: $\mathrm{NCON}=\operatorname{size}(\mathrm{C}, 1)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
LDA must be at least NRA.
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.

LDC - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
LDC must be at least NCON.
Default: LDC = size ( $\mathrm{C}, 1$ ).
$\boldsymbol{R E S}$ - Vector of length NRA containing the residuals $B-A X$ of the least-squares equations at the approximate solution. (Output)

## FORTRAN 90 Interface

Generic: CALL LCLSQ (A, B, C, BL, BU, IRTYPE, XLB, XUB, X [, ...])
Specific: The specific interface names are S_LCLSQ and D_LCLSQ.

## FORTRAN 77 Interface

Single: CALL LCLSQ (NRA, NCA, NCON, A, LDA, B, C, LDC, BL, BU, IRTYPE, XLB, XUB, X, RES)

Double: $\quad$ The double precision name is DLCLSQ.

## Description

The routine LCLSQ solves linear least-squares problems with linear constraints. These are systems of least-squares equations of the form $A x \cong \mathrm{~b}$
subject to

$$
\begin{gathered}
b_{l} \leq C_{x} \leq b_{u} \\
x_{l} \leq x \leq x_{u}
\end{gathered}
$$

Here, $A$ is the coefficient matrix of the least-squares equations, $b$ is the right-hand side, and $C$ is the coefficient matrix of the constraints. The vectors $b_{l}, b_{u}, x_{l}$ and $x_{u}$ are the lower and upper
bounds on the constraints and the variables, respectively. The system is solved by defining dependent variables $y \equiv C x$ and then solving the least squares system with the lower and upper bounds on $x$ and $y$. The equation $C x-\mathrm{y}=0$ is a set of equality constraints. These constraints are realized by heavy weighting, i.e. a penalty method, Hanson, (1986, pages 826-834).

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2LSQ/DL2LSQ. The reference is:
```
CALL L2LSQ (NRA, NCA, NCON, A, LDA, B, C, LDC, BL, BU, IRTYPE,
XLB, XUB, X, RES, WK, IWK)
```

The additional arguments are as follows:

$$
\begin{aligned}
& W K-\text { Real work vector of length }(\text { NCON }+ \text { MAXDIM }) *(N C A+N C O N+1)+10 * N C A+ \\
& 9 * \text { NCON }+3 \text {. }
\end{aligned}
$$

$\boldsymbol{I W K}$ - Integer work vector of length 3 * (NCON + NCA).
2. Informational errors

Type Code
31 The rank determination tolerance is less than machine precision.
42 The bounds on the variables are inconsistent.
43 The constraint bounds are inconsistent.
44 Maximum number of iterations exceeded.
3. Integer Options with Chapter 11 Options Manager

13 Debug output flag. If more detailed output is desired, set this option to the value 1. Otherwise, set it to 0 . Default value is 0 .

14 Maximum number of add/drop iterations. If the value of this option is zero, up to $5 * \max (\mathrm{nra}, \mathrm{nca})$ iterations will be allowed. Otherwise set this option to the desired iteration limit. Default value is 0 .
4. Floating Point Options with Chapter 11 Options Manager

2 The value of this option is the relative rank determination tolerance to be used. Default value is sqrt(AMACH (4)).

5 The value of this option is the absolute rank determination tolerance to be used. Default value is sqrt(AMACH (4)).

## Example

A linear least-squares problem with linear constraints is solved.

```
USE LCLSQ_INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE SNRM2_INT
```

```
    Solve the following in the least squares sense:
    3x1 + 2x2 + x3 = 3.3
    4x1 + 2x2 + x3 = 2.3
    2x1 + 2x2 + x3 = 1.3
        x1 + x2 + x3 = 1.0
    Subject to: x1 + x2 + x3 <= 1
            0<= x1 <= . 5
            0<= x2 <= . 5
            0<= x3 <= . 5
    D----------------------------------------------------------
    INTEGER NRA, NCA, MCON, LDA, LDC
    PARAMETER (NRA=4, NCA=3, MCON=1, LDC=MCON, LDA=NRA)
!
    INTEGER IRTYPE(MCON), NOUT
    REAL A(LDA,NCA), B(NRA), BC(MCON), C(LDC,NCA), RES (NRA), &
            RESNRM, XSOL (NCA), XLB(NCA), XUB(NCA)
                                    Data initialization!
    DATA A/3.0E0, 4.0E0, 2.0E0, 1.0E0, 2.0E0, &
        2.0EO, 2.0EO, 1.0E0, 1.0E0, 1.0E0, 1.0E0, 1.0E0/, &
        B/3.3E0, 2.3E0, 1.3E0, 1.0E0/, &
        C/3*1.0E0/, &
        BC/1.0E0/, IRTYPE/1/, XLB/3*0.0E0/, XUB/3*.5E0/
                            Solve the bounded, constrained
                    least squares problem.
    CALL LCLSQ (A, B, C, BC, BC, IRTYPE, XLB, XUB, XSOL, RES=RES)
                            Compute the 2-norm of the residuals.
    RESNRM = SNRM2 (NRA, RES, 1)
    CALL UMACH (2, NOUT)
    WRITE (NOUT, 999) XSOL, RES, RESNRM
!
    999 FORMAT (' The solution is ', 3F9.4, //, ' The residuals ', &
            'evaluated at the solution are ', /, 18X, 4F9.4, //, &
            , The norm of the residual vector is ', F8.4)
!
    END
```


## Output

The solution is $0.5000 \quad 0.3000 \quad 0.2000$
The residuals evaluated at the solution are

$$
\begin{array}{llll}
-1.0000 & 0.5000 & 0.5000 & 0.0000
\end{array}
$$

The norm of the residual vector is 1.2247

## LQRRR

## 匀MPI <br> CAPABLE

Computes the $Q R$ decomposition, $A P=Q R$, using Householder transformations.

## Required Arguments

A - Real NRA by NCA matrix containing the matrix whose $Q R$ factorization is to be computed. (Input)
$\boldsymbol{Q R}$ - Real NRA by NCA matrix containing information required for the $Q R$ factorization. (Output)
The upper trapezoidal part of $Q R$ contains the upper trapezoidal part of $R$ with its diagonal elements ordered in decreasing magnitude. The strict lower trapezoidal part of $Q R$ contains information to recover the orthogonal matrix $Q$ of the factorization.
Arguments A and $Q R$ can occupy the same storage locations. In this case, A will not be preserved on output.

QRAUX - Real vector of length NCA containing information about the orthogonal part of the decomposition in the first $\min$ (NRA, NCA) position. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{size}(A, 1)$.
NCA — Number of columns of A. (Input)
Default: NCA $=\operatorname{size}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
PIVOT - Logical variable. (Input)
PIVOT = . TRUE . means column pivoting is enforced.
PIVOT $=$. FALSE. means column pivoting is not done.
Default: PIVOT = .TRUE.
IPVT - Integer vector of length NCA containing information that controls the final order of the columns of the factored matrix A. (Input/Output)
On input, if $\operatorname{IPVT}(\mathrm{K})>0$, then the K -th column of A is an initial column. $\operatorname{If} \operatorname{IPVT}(\mathrm{K})=0$, then the $K$-th column of $A$ is a free column. If $\operatorname{IPVT}(\kappa)<0$, then the $K$-th column of $A$ is a final column. See Comments.
On output, $\operatorname{IPVT(K)~contains~the~index~of~the~column~of~A~that~has~been~interchanged~}$ into the k-th column. This defines the permutation matrix $P$. The array IPVT is
referenced only if PIVOT is equal to .TRUE. Default: $I P V T=0$.

LDQR — Leading dimension of $Q R$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDQR = size $(Q R, 1)$.

CONORM - Real vector of length NCA containing the norms of the columns of the input matrix. (Output)
If this information is not needed, CONORM and QRAUX can share the same storage locations.

## FORTRAN 90 Interface

Generic: CALL LQRRR (A, QR, QRAUX [,...])
Specific: The specific interface names are $S \_L Q R R R$ and $D \_L Q R R R$.

## FORTRAN 77 Interface

Single: CALL LQRRR (NRA, NCA, A, LDA, PIVOT, IPVT, QR, LDQR, QRAUX, CONORM)
Double: The double precision name is $D L Q R R R$.

## ScaLAPACK Interface

Generic: CALL LQRRR (A0, QR0, QRAUX0 [, ...])
Specific: The specific interface names are $S \_L Q R R R$ and $D \_L Q R R R$.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

The routine $L Q R R R$ computes the $Q R$ decomposition of a matrix using Householder transformations. The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

LQRRR determines an orthogonal matrix $Q$, a permutation matrix $P$, and an upper trapezoidal matrix $R$ with diagonal elements of nonincreasing magnitude, such that $A P=Q R$. The Householder transformation for column $k$ is of the form

$$
I-\frac{u_{k} u_{k}^{T}}{p_{k}}
$$

for $k=1,2, \ldots, \min (N R A, N C A)$, where $u$ has zeros in the first $k-1$ positions. The matrix $Q$ is not produced directly by $L Q R R R$. Instead the information needed to reconstruct the Householder transformations is saved. If the matrix $Q$ is needed explicitly, the subroutine LQERR can be called after LQRRR. This routine accumulates $Q$ from its factored form.

Before the decomposition is computed, initial columns are moved to the beginning of the array $A$ and the final columns to the end. Both initial and final columns are frozen in place during the computation. Only free columns are pivoted. Pivoting, when requested, is done on the free columns of largest reduced norm.

## Comments

1. Workspace may be explicitly provided, if desired, by use of $L 2 R R R / D L 2 R R R$. The reference is:

CALL L2RRR (NRA, NCA, A, LDA, PIVOT, IPVT, QR, LDQR, QRAUX, CONORM, WORK)

The additional argument is
WORK - Work vector of length $2 \mathrm{NCA}-1$. Only NCA - 1 locations of WORK are referenced if PIVOT $=$. FALSE.
2. $\quad L Q R R R$ determines an orthogonal matrix $Q$, permutation matrix $P$, and an upper trapezoidal matrix $R$ with diagonal elements of nonincreasing magnitude, such that $A P=Q R$. The Householder transformation for column $k, k=1, \ldots, \min (\mathrm{NRA}, \mathrm{NCA})$ is of the form

$$
I-u_{k}^{-1} u u^{T}
$$

where $u$ has zeros in the first $k-1$ positions. If the explicit matrix $Q$ is needed, the user can call routine LQERR after calling LQRRR. This routine accumulates $Q$ from its factored form.
3. Before the decomposition is computed, initial columns are moved to the beginning and the final columns to the end of the array A. Both initial and final columns are not moved during the computation. Only free columns are moved. Pivoting, if requested, is done on the free columns of largest reduced norm.
4. When pivoting has been selected by having entries of IPVT initialized to zero, an estimate of the condition number of $A$ can be obtained from the output by computing the magnitude of the number $\operatorname{QR}(1,1) / Q R(K, K)$, where $K=M I N(N R A, N C A)$. This estimate can be used to select the number of columns, KBASIS, used in the solution step computed with routine LQRSL.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:

A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix whose $Q R$ factorization is to be computed. (Input)

QRO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix $Q R . Q R$ contains the information required for the $Q R$ factorization. (Output) The upper trapezoidal part of $Q R$ contains the upper trapezoidal part of $R$ with its diagonal elements ordered in decreasing magnitude. The strict lower trapezoidal part of $Q R$ contains information to recover the orthogonal matrix $Q$ of the factorization. Arguments A and QR can occupy the same storage locations. In this case, A will not be preserved on output.

QRAUXO - Real vector of length MXCOL containing the local portions of the distributed matrix QRAUX. QRAUX contains information about the orthogonal part of the decomposition in the first MIN(NRA, NCA) position. (Output)

IPVT0 - Integer vector of length MXLDB containing the local portions of the distributed vector IPVT. IPVT contains the information that controls the final order of the columns of the factored matrix A. (Input/Output)
On input, if $\operatorname{IPVT}(K)>0$, then the $K-t h$ column of $A$ is an initial column. $\operatorname{If} \operatorname{IPVT}(\mathrm{K})=0$, then the $K$-th column of $A$ is a free column. If $\operatorname{IPVT}(\mathrm{K})<0$, then the $K$-th column of $A$ is a final column. See Comments.
On output, $\operatorname{IPVT}(\mathrm{K})$ contains the index of the column of $A$ that has been interchanged into the K -th column. This defines the permutation matrix $P$. The array IPVT is referenced only if PIVOT is equal to .TRUE.
Default: $I P V T=0$.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXLDB, and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

## Example

In various statistical algorithms it is necessary to compute $q=x^{T}\left(A^{T} A\right)^{-1} x$, where $A$ is a rectangular matrix of full column rank. By using the $Q R$ decomposition, $q$ can be computed without forming $A^{T} A$. Note that

$$
A^{T} A=\left(Q R P^{-1}\right)^{T}\left(Q R P^{-1}\right)=P^{-T} R^{T}\left(Q^{T} Q\right) R P^{-1}=P R^{T} R P^{T}
$$

since $Q$ is orthogonal $\left(Q^{T} Q=I\right)$ and $P$ is a permutation matrix. Let

$$
Q^{T} A P=R=\left[\begin{array}{c}
R_{1} \\
0
\end{array}\right]
$$

where $R_{\mathrm{t}}$ is an upper triangular nonsingular matrix. Then

$$
x^{T}\left(A^{T} A\right)^{-1} x=x^{T} P R_{1}^{-1} R_{1}^{-T} P^{-1} x=\left\|R_{1}^{-T} P^{-1} x\right\|_{2}^{2}
$$

In the following program, first the vector $t=P^{-1} x$ is computed. Then

$$
t:=R_{1}^{-T} t
$$

Finally,

$$
q=\|t\|^{2}
$$

USE IMSL_LIBRARIES
INTEGER LDA, LDQR, NCA, NRA
PARAMETER ( $\mathrm{NCA}=3, \mathrm{NRA}=4, ~ L D A=N R A, ~ L D Q R=N R A)$
SPECIFICATIONS FOR PARAMETERS
INTEGER LDQ
PARAMETER (LDQ=NRA)
INTEGER IPVT (NCA), NOUT
REAL CONORM (NCA), Q, QR(LDQR,NCA), QRAUX(NCA), T(NCA)
LOGICAL PIVOT
REAL A (LDA, NCA), X (NCA)
Set values for A
$A=\left(\begin{array}{llll}1 & 2 & 4\end{array}\right)$
$\left(\begin{array}{llll}1 & 6 & 36\end{array}\right)$
DATA A/4*1.0, 2.0, 4.0, 6.0, 8.0, 4.0, 16.0, 36.0, 64.0/
Set values for X
$X=\left(\begin{array}{llll}1 & 2 & 3\end{array}\right)$
DATA X/1.0, 2.0, 3.0/
PIVOT = .TRUE.
IPVT=0
CALL LQRRR (A, QR, QRAUX, PIVOT=PIVOT, IPVT=IPVT)
Set $t=i n v(P) * x$
CALL PERMU (X, IPVT, T, IPATH=1)
Compute $t=i n v(t r a n s(R)) * t$
CALL LSLRT (QR, T, T, IPATH=4)
Compute 2 -norm of $t$, squared.
$\mathrm{Q}=\operatorname{SDOT}(\mathrm{NCA}, \mathrm{T}, 1, \mathrm{~T}, 1)$
Print result
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'Q = ', Q
!
END

## Output

$\mathrm{Q}=0.840624$

## ScaLAPACK Example

The previous example is repeated here as a distributed computing example. In various statistical algorithms it is necessary to compute $q=x^{T}\left(A^{T} A\right)^{-1} x$, where $A$ is a rectangular matrix of full column rank. By using the $Q R$ decomposition, $q$ can be computed without forming $A^{T} A$. Note that

$$
A^{T} A=\left(Q R P^{-1}\right)^{T}\left(Q R P^{-1}\right)=P^{-\mathrm{T}} R^{T}\left(Q^{T} Q\right) R P^{-1}=P R^{T} R P^{T}
$$

since $Q$ is orthogonal $\left(Q^{T} Q=I\right)$ and $P$ is a permutation matrix. Let

$$
Q^{T} A P=R=\left[\begin{array}{c}
R_{1} \\
0
\end{array}\right]
$$

where $R_{\mathrm{t}}$ is an upper triangular nonsingular matrix. Then

$$
x^{T}\left(A^{T} A\right)^{-1} x=x^{T} P R_{1}^{-1} R_{1}^{-T} P^{-1} x=\left\|R_{1}^{-T} P^{-1} x\right\|_{2}^{2}
$$

In the following program, first the vector $t=P^{-1} x$ is computed. Then

$$
t:=R_{1}^{-T} t
$$

Finally,

$$
q=\|t\|^{2}
$$

SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.

```
USE MPI_SETUP_INT
USE LQR\overline{RR_INT}
USE PERMU_INT
USE LSLRT_INT
USE UMACH_INT
USE SCALAPACCK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
    INTEGER LDA, LDQR, NCA, NRA, DESCA(9), DESCB(9), DESCL(9)
    INTEGER INFO, MXCOL, MXLDA, MXLDB, MXCOLB, NOUT
    INTEGER, ALLOCATABLE :: IPVT(:), IPVTO(:)
    LOGICAL PIVOT
    REAL Q
    REAL, ALLOCATABLE :: A(:,:), X(:), T(:)
    REAL, ALLOCATABLE :: AO(:,:), TO(:), QRO(:,:), QRAUXO(:)
    REAL, (KIND(1E0)) SDOT
    PARAMETER (NRA=4, NCA=3, LDA=NRA, LDQR=NRA)
    MP_NPROCS = MP_SETUP()
    IF(MP_RANK .EQ. 0) THEN
        ALLOCATE (A(LDA,NCA), X(NCA), T(NCA), IPVT(NCA))
    Set values for A and the righthand side
```

```
    A(1,:) = (/ 1.0, 2.0, 4.0/)
    A(2,:) = (/ 1.0, 4.0, 16.0/)
    A(3,:) = (/ 1.0, 6.0, 36.0/)
    A(4,:) = (/ 1.0, 8.0, 64.0/)
    X = (/ 1.0, 2.0, 3.0/)
    IPVT = 0
ENDIF
    Set up a 1D processor grid and define
    its context ID, MP ICTXT
CALL SCALAPACK_SETUP(NRA, NCA, .TRUE., .TRUE.)
    Get the array descriptor entities MXLDA,
    MXCOL, MXLDB, MXCOLB
    CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
CALL SCALAPACK_GETDIM(NCA, 1, MP_NB, 1, MXLDB, MXCOLB)
    Set up the array descriptors
    CALL DESCINIT(DESCA, NRA, NCA, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, &
    INFO)
    CALL DESCINIT(DESCL, 1, NCA, 1, MP NB, 0, 0, MP_ICTXT, 1, INFO)
    CALL DESCINIT(DESCB, NCA, 1, MP_NB, 1, 0, 0, MP_ICTXT, MXLDB, &
    INFO)
                            Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), QR0 (MXLDA,MXCOL), QRAUXO (MXCOL), &
        IPVTO (MXCOL), TO (MXLDB))
                            Map input array to the processor grid
    CALL SCALAPACK MAP(A, DESCA, AO)
    PIVOT = .TRUE.
CALL SCALAPACK MAP(IPVT, DESCL, IPVTO)
    QR factorization
    CALL LQRRR (A0, QR0, QRAUX0, PIVOT=PIVOT, IPVT=IPVTO)
        Unmap the results from the distributed
        array back to a non-distributed array.
        After the unmap, only Rank=0 has the full
        array.
    CALL SCALAPACK_UNMAP(IPVTO, DESCL, IPVT, NCA, .FALSE.)
    IF(MP RANK .EQ. O) CALL PERMU (X, IPVT, T, IPATH=1)
    CALL S
    CALL LSLRT (QRO, TO, TO, IPATH=4)
    CALL SCALAPACK_UNMAP(T0, DESCB, T)
        Print results.
        Only Rank=0 has the solution.
    IF (MP RANK .EQ. O)THEN
    Q = SDOT (NCA, T, 1, T, 1)
    CALL UMACH (2, NOUT)
    WRITE (NOUT, *) 'Q = `, Q
ENDIF
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
                    Shut down MPI
MP NPROCS = MP SETUP('FINAL')
END
```


## Output

$Q=0.840624$

## LQERR

Accumulates the orthogonal matrix $Q$ from its factored form given the $Q R$ factorization of a rectangular matrix $A$.

## Required Arguments

$\boldsymbol{Q R}$ - Real NRQR by NCQR matrix containing the factored form of the matrix $Q$ in the first $\min (N R Q R, N C Q R)$ columns of the strict lower trapezoidal part of $Q R$ as output from subroutine LQRRR/DLQRRR. (Input)

QRAUX - Real vector of length NCQR containing information about the orthogonal part of the decomposition in the first $\min (N R Q R, N C Q R)$ position as output from routine LQRRR/DLQRRR. (Input)
$\boldsymbol{Q}$ - Real $N R Q R$ by $N R Q R$ matrix containing the accumulated orthogonal matrix $Q ; Q$ and $Q R$ can share the same storage locations if $Q R$ is not needed. (Output)

## Optional Arguments

$N R Q R$ - Number of rows in $Q R$. (Input) Default: $N R Q R=\operatorname{size}(Q R, 1)$.
$N C Q R$ - Number of columns in QR. (Input) Default: NCQR = size (QR,2).

LDQR - Leading dimension of $Q R$ exactly as specified in the dimension statement of the calling program. (Input)
Default: $L D Q R=\operatorname{size}(Q R, 1)$.

LDQ - Leading dimension of $Q$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDQ = size $(2,1)$.

## FORTRAN 90 Interface

Generic: CALL LQERR (QR, QRAUX, Q [,...])

Specific: The specific interface names are S_LQERR and D_LQERR.

## FORTRAN 77 Interface

Single: CALL LQERR (NRQR, NCQR, QR, LDQR, QRAUX, Q, LDQ)
Double: The double precision name is DLQERR.

## ScaLAPACK Interface

Generic: CALL LQERR (QRO, QRAUXO, Q0 [,...])

Specific: The specific interface names are S_LQERR and D_LQERR.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

The routine LQERR accumulates the Householder transformations computed by IMSL routine $L Q R R R$ to produce the orthogonal matrix $Q$.

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2ERR/DL2ERR. The reference is:

CALL L2ERR (NRQR, NCQR, QR, LDQR, QRAUX, Q, LDQ, WK)
The additional argument is

WK - Work vector of length 2 * NRQR.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:

QR0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix $Q R . Q R$ contains the factored form of the matrix $Q$ in the first min(NRQR, $N C Q R$ ) columns of the strict lower trapezoidal part of $Q R$ as output from subroutine LQRRR/DLQRRR. (Input)

QRAUXO - Real vector of length MXCOL containing the local portions of the distributed matrix QRAUX. QRAUX contains the information about the orthogonal part of the decomposition in the first min(NRA, NCA) positions as output from subroutine LQRRR/DLQRRR. (Input)

Q0 - MXLDA by MXLDA local matrix containing the local portions of the distributed matrix $Q$. $Q$ contains the accumulated orthogonal matrix $; Q$ and $Q R$ can share the same storage locations if $Q R$ is not needed. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

## Example

In this example, the orthogonal matrix $Q$ in the $Q R$ decomposition of a matrix $A$ is computed. The product $X=Q R$ is also computed. Note that $X$ can be obtained from $A$ by reordering the columns of $A$ according to IPVT.

USE IMSL_LIBRARIES

```
! - Declare variables
    INTEGER LDA, LDQ, LDQR, NCA, NRA
    PARAMETER (NCA=3, NRA=4, LDA=NRA, LDQ=NRA, LDQR=NRA)
    INTEGER IPVT(NCA), J
    REAL A(LDA,NCA), CONORM(NCA), Q(LDQ,NRA), QR(LDQR,NCA), &
        QRAUX(NCA), R(NRA,NCA), X(NRA,NCA)
    LOGICAL PIVOT
    DATA A/4*1.0, 2.0, 4.0, 6.0, 8.0, 4.0, 16.0, 36.0, 64.0/
        QR factorization
                            Set IPVT = O (all columns free)
    IPVT = 0
    PIVOT = .TRUE.
    CALL LQRRR (A, QR, QRAUX, IPVT=IPVT, PIVOT=PIVOT)
    CALL LQERR (QR, QRAUX, Q)
    R = 0.0E0
    DO 10 J=1, NCA
        CALL SCOPY (J, QR(:,J), 1, R(:,J), 1)
    1 0 ~ C O N T I N U E ~
    Compute X = Q*R
    CALL MRRRR (Q, R, X)
    Print results
    CALL WRIRN ('IPVT', IPVT, 1, NCA, 1)
    CALL WRRRN ('Q', Q)
    CALL WRRRN ('R', R)
    CALL WRRRN ('X = Q*R', X)
```

!
END

## Output

| IPVT |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| 1 | 23 |  |  |  |
| 3 | 21 |  |  |  |
|  | Q |  |  |  |
|  | 1 | 2 | 3 | 4 |
| 1 | -0.0531 | -0.5422 | 0.8082 | -0.2236 |
| 2 | -0.2126 | -0.6574 | -0.2694 | 0.6708 |
| 3 | -0.4783 | -0.3458 | -0.4490 | -0.6708 |
| 4 | -0.8504 | 0.3928 | 0.2694 | 0.2236 |

R

| 1 | 2 | 3 |
| ---: | ---: | ---: |
| -75.26 | -10.63 | -1.59 |
| 0.00 | -2.65 | -1.15 |
| 0.00 | 0.00 | 0.36 |
| 0.00 | 0.00 | 0.00 |
|  |  |  |
| $X$ | $=2 * R$ |  |
| 4.00 | 2.00 | 1.00 |
| 16.00 | 4.00 | 1.00 |
| 36.00 | 6.00 | 1.00 |
| 64.00 | 8.00 | 1.00 |

## ScaLAPACK Example

In this example, the orthogonal matrix $Q$ in the $Q R$ decomposition of a matrix $A$ is computed. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines (see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.

```
USE MPI_SETUP_INT
USE LQR\overline{RR INT}
USE LQERR_INT
USE WRRRN INT
USE SCALA\overline{PACK SUPPORT}
IMPLICIT NONE
INCLUDE 'mpif.h'
INTEGER LDA, LDQR, NCA, NRA, DESCA(9), DESCL(9), DESCQ(9)
INTEGER INFO, MXCOL, MXLDA, LDQ
INTEGER, ALLOCATABLE :: IPVT(:), IPVTO(:)
LOGICAL PIVOT
REAL, ALLOCATABLE :: A(:,:), QR(:,:), Q(:,:), QRAUX(:)
REAL, ALLOCATABLE :: AO(:,:), QRO(:,:), QO(:,:), QRAUX0(:)
PARAMETER (NRA=4, NCA=3, LDA=NRA, LDQR=NRA, LDQ=NRA)
MP_NPROCS = MP_SETUP()
IF(MP RANK .EQ. 0) THEN
```

```
        ALLOCATE (A (NRA,NCA), Q(NRA,NRA), QR (NRA,NCA), &
        QRAUX(NCA), IPVT (NCA))
        Set values for A and the righthand sides
        A(1,:) = (/ 1.0, 2.0, 4.0/)
        A(2,:) = (/ 1.0, 4.0, 16.0/)
        A(3,:) = (/ 1.0, 6.0, 36.0/)
        A(4,:) = (/ 1.0, 8.0, 64.0/)
        IPVT = 0
    ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
    CALL SCALAPACK_SETUP(NRA, NCA, .FALSE., .TRUE.)
        Get the array descriptor entities MXLDA,
                        and MXCOL
    CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
            Set up the a\overline{rray descriptors}
    CALL DESCINIT(DESCA, NRA, NCA, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, &
    INFO)
    CALL DESCINIT(DESCL, 1, NCA, 1, MP NB, 0, 0, MP ICTXT, 1, INFO)
    CALL DESCINIT(DESCQ, NRA, NRA, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, &
    INFO)
                            Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), QR0 (MXLDA,MXCOL), QRAUX0 (MXCOL), &
        IPVTO (MXCOL), QO (MXLDA,MXLDA))
                                    Map input array to the processor grid
    CALL SCALAPACK MAP(A, DESCA, AO)
    PIVOT = .TRUE.
    CALL SCALAPACK MAP(IPVT, DESCL, IPVT0)
                                    QR factorization
    CALL LQRRR (A0, QRO, QRAUX0, PIVOT=PIVOT, IPVT=IPVTO)
    CALL LQERR (QRO, QRAUXO, QO)
                                    Unmap the results from the distributed
                                    array back to a non-distributed array.
                                    After the unmap, only Rank=0 has the full
                                    array.
    CALL SCALAPACK_UNMAP(Q0, DESCQ, Q)
                                    Print results.
                            Only Rank=0 has the solution, Q.
    IF(MP RANK .EQ. O) CALL WRRRN ('Q', Q)
            Exit Scalapack usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
        Shut down MPI
    MP_NPROCS = MP_SETUP('FINAL')
    END
```


## LQRSL

## (i)MPI <br> CAPABLE

Computes the coordinate transformation, projection, and complete the solution of the least-squares problem $A x=b$.

## Required Arguments

KBASIS - Number of columns of the submatrix $A_{k}$ of $A$. (Input) The value KBASIS must not exceed min(NRA, NCA), where NCA is the number of columns in matrix $A$. The value NCA is an argument to routine LQRRR. The value of KBASIS is normally NCA unless the matrix is rank-deficient. The user must analyze the problem data and determine the value of KBASIS. See Comments.
$\boldsymbol{Q R}$ - NRA by NCA array containing information about the $Q R$ factorization of $A$ as output from routine $L Q R R R / D L Q R R R$. (Input)

QRAUX - Vector of length NCA containing information about the $Q R$ factorization of $A$ as output from routine $L Q R R R / D L Q R R R$. (Input)
$\boldsymbol{B}-$ Vector $b$ of length NRA to be manipulated. (Input)
IPATH - Option parameter specifying what is to be computed. (Input)
The value IPATH has the decimal expansion IJKLM, such that:
I $\neq 0$ means compute $Q b$;
$J \neq 0$ means compute $Q^{T} b$;
$\mathrm{K} \neq 0$ means compute $\mathrm{Q}^{T} \mathrm{~b}$ and $x$;
$\mathrm{L} \neq 0$ means compute $\mathrm{Q}^{T} \mathrm{~b}$ and $b-A x$;
$\mathrm{M} \neq 0$ means compute $Q^{T} b$ and $A x$.
For example, if the decimal number IPATH $=01101$, then $I=0, J=1, K=1$, $\mathrm{L}=0$, and $\mathrm{M}=1$.

## Optional Arguments

$N R A$ - Number of rows of matrix $A$. (Input)
Default: NRA $=\operatorname{size}(Q R, 1)$.
$\boldsymbol{L D Q R}$ - Leading dimension of $Q R$ exactly as specified in the dimension statement of the calling program. (Input)
Default: $L D Q R=\operatorname{size}(Q R, 1)$.
$\boldsymbol{Q B}-$ Vector of length NRA containing $Q b$ if requested in the option IPATH. (Output)
$\boldsymbol{Q T B}$ - Vector of length NRA containing $Q^{T} b$ if requested in the option IPATH. (Output)
$\boldsymbol{X}$ - Vector of length KBASIS containing the solution of the least-squares problem $A_{k} x=b$, if this is requested in the option IPATH. (Output)
If pivoting was requested in routine $L Q R R R / D L Q R R R$, then the $J$-th entry of $X$ will be associated with column IPVT(J) of the original matrix $A$. See Comments.
$\boldsymbol{R E S}$ - Vector of length NRA containing the residuals $(b-A x)$ of the least-squares problem if requested in the option IPATH. (Output)
This vector is the orthogonal projection of $b$ onto the orthogonal complement of the column space of $A$.
$\boldsymbol{A} \boldsymbol{X}$ - Vector of length NRA containing the least-squares approximation $A x$ if requested in the option IPATH. (Output) This vector is the orthogonal projection of $b$ onto the column space of $A$.

## FORTRAN 90 Interface

Generic: CALL LQRSL (KBASIS, QR, QRAUX, B, IPATH [,...])

Specific: The specific interface names are $S_{-} L Q R S L$ and $D_{-} L Q R S L$.

## FORTRAN 77 Interface

Single: CALL LQRSL (NRA, KBASIS, QR, LDQR, QRAUX, B, IPATH, QB, QTB, X, RES, AX)

Double: $\quad$ The double precision name is $D L Q R S L$.

## ScaLAPACK Interface

Generic: CALL LQRSL (KBASIS, QRO, QRAUX0, B0, IPATH [, ...])
Specific: The specific interface names are $S_{-} \quad L Q R S L$ and $D_{-} L Q R S L$.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

The most important use of $L Q R S L$ is for solving the least-squares problem $A x=b$, with coefficient matrix $A$ and data vector $b$. This problem can be formulated, using the normal equations method, as $A^{T} A x=A^{T} b$. Using LQRRR the $Q R$ decomposition of $A, A P=Q R$, is computed. Here $P$ is a
permutation matrix $(P=P), Q$ is an orthogonal matrix $\left(Q=Q^{T}\right)$ and $R$ is an upper trapezoidal matrix. The normal equations can then be written as

$$
\left(P R^{T}\right)\left(Q^{T} Q\right) R\left(P^{T} x\right)=\left(P R^{T}\right) Q^{T} b
$$

If $A^{T} A$ is nonsingular, then $R$ is also nonsingular and the normal equations can be written as $R\left(P^{T} x\right)=Q^{T} b$. LQRSL can be used to compute $Q^{T} b$ and then solve for $P^{T} x$. Note that the permuted solution is returned.

The routine LQRSL can also be used to compute the least-squares residual, $b-A x$. This is the projection of $b$ onto the orthogonal complement of the column space of $A$. It can also compute $Q b$, $Q^{T} b$ and $A x$, the orthogonal projection of $x$ onto the column space of $A$.

## Comments

1. Informational error

Type Code
41 Computation of the least-squares solution of $A K * X=B$ is requested, but the upper triangular matrix R from the $Q R$ factorization is singular.
2. This routine is designed to be used together with $L Q R R R$. It assumes that $L Q R R R / D L Q R R$ has been called to get $Q R$, QRAUX and IPVT. The submatrix $A_{k}$ mentioned above is actually equal to $A_{k}=(A(\operatorname{IPVT}(1)), A(\operatorname{IPVT}(2)), \ldots, A(\operatorname{IPVT}(\mathrm{KBASIS})))$, where $A(\operatorname{IPVT}(\mathrm{I}))$ is the $\operatorname{IPVT}(\mathrm{I})$-th column of the original matrix.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:
QRO - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix $Q R . Q R$ contains the factored form of the matrix $Q$ in the first min(NRQR, $N C Q R$ ) columns of the strict lower trapezoidal part of $Q R$ as output from subroutine LQRRR/DLQRRR. (Input)

QRAUXO - Real vector of length MXCOL containing the local portions of the distributed matrix QRAUX. QRAUX contains the information about the orthogonal part of the decomposition in the first $\min (N R A, N C A)$ positions as output from subroutine $\mathrm{L} Q R R \mathrm{R} / \mathrm{DLQRRR}$. (Input)

BO - Real vector of length MXLDA containing the local portions of the distributed vector B. $B$ contains the vector to be manipulated. (Input)

QBO - Real vector of length MXLDA containing the local portions of the distributed vector $Q b$ if requested in the option IPATH. (Output)

QTBO - Real vector of length MXLDA containing the local portions of the distributed vector $Q^{T} b$ if requested in the option IPATH. (Output)
$\boldsymbol{X O}$ - Real vector of length MXLDX containing the local portions of the distributed vector $\mathrm{X} . \mathrm{X}$ contains the solution of the least-squares problem $A_{k} x=b$, if this is requested in the option IPATH. (Output)
If pivoting was requested in routine $L Q R R R / D L Q R R R$, then the $J$-th entry of $X$ will be associated with column $\operatorname{IPVT}(J)$ of the original matrix $A$. See Comments.

RESO - Real vector of length MXLDA containing the local portions of the distributed vector RES. RES contains the residuals $(b-A x)$ of the least-squares problem if requested in the option IPATH. (Output)
This vector is the orthogonal projection of $b$ onto the orthogonal complement of the column space of $A$.

AXO - Real vector of length MXLDA containing the local portions of the distributed vector $A X$. AX contains the least-squares approximation $A x$ if requested in the option IPATH. (Output)
This vector is the orthogonal projection of $b$ onto the column space of $A$.
All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXLDX and MXCOL can be obtained through a call to SCALAPACK_GETDIM (see Utilities) after a call to SCALAPACK_SETUP (see Utilities) has been made. See the ScaLAPACK Example below.

## Example

Consider the problem of finding the coefficients $\mathrm{c}_{i}$ in

$$
f(x)=c_{0}+c_{1} x+c_{2} x^{2}
$$

given data at $x_{i}=2_{i,}, \quad=1,2,3,4$, using the method of least squares. The row of the matrix $A$ contains the value of $1, x_{i}$ and

$$
x_{i}^{2}
$$

at the data points. The vector $b$ contains the data. The routine $L Q R R R$ is used to compute the $Q R$ decomposition of $A$. Then LQRSL is then used to solve the least-squares problem and compute the residual vector.

USE IMSL_LIBRARIES

```
PARAMETER (NRA=4, NCA=3, KBASIS=3, LDA=NRA, LDQR=NRA)
INTEGER IPVT (NCA)
REAL A(LDA,NCA), QR(LDQR,NCA), QRAUX(NCA), CONORM(NCA), &
    X(KBASIS), QB(1), QTB(NRA), RES(NRA), &
    AX(1), B (NRA)
LOGICAL PIVOT
```

                                    Set values for A
            \(A=\left(\begin{array}{llll}1 & 2 & 4\end{array}\right)\)
    $\left(\begin{array}{lll}1 & 4 & 16\end{array}\right)$

```
! (\begin{array}{cllll}{1}&{8}&{64}\end{array})
DATA A/4*1.0, 2.0, 4.0, 6.0, 8.0, 4.0, 16.0, 36.0, 64.0/
! Set values for B
DATA B/ 16.99, 57.01, 120.99, ( 16.99 50.01 ().01 120.99 209.01 )
QR factorization
PIVOT = .TRUE.
IPVT = 0
CALL LQRRR (A, QR, QRAUX, PIVOT=PIVOT, IPVT=IPVT)
    Solve the least squares problem
IPATH = 00110
CALL LQRSL (KBASIS, QR, QRAUX, B, IPATH, X=X, RES=RES)
    Print results
CALL WRIRN ('IPVT', IPVT, 1, NCA, 1)
CALL WRRRN ('X', X, 1, KBASIS, 1)
CALL WRRRN ('RES', RES, 1, NRA, 1)
!
END
```


## Output

| IPVT |  |  |
| :---: | :---: | :---: |
| 1 | 2 | 3 |
| 3 | 2 | 1 |



Note that since IPVT is $(3,2,1)$ the array X contains the solution coefficients $c_{i}$ in reverse order.

## ScaLAPACK Example

The previous example is repeated here as a distributed example. Consider the problem of finding the coefficients $\mathrm{c}_{i}$ in

$$
f(x)=c_{0}+c_{1} x+c_{2} x^{2}
$$

given data at $x_{i}=2_{i}, t=1,2,3,4$, using the method of least squares. The row of the matrix $A$ contains the value of $1, x_{i}$ and

$$
x_{i}^{2}
$$

at the data points. The vector $b$ contains the data. The routine $L Q R R R$ is used to compute the $Q R$ decomposition of $A$. Then LQRSL is then used to solve the least-squares problem and compute the residual vector. SCALAPACK_MAP and SCALAPACK_UNMAP are IMSL utility routines
(see Chapter 11, "Utilities") used to map and unmap arrays to and from the processor grid. They are used here for brevity. DESCINIT is a ScaLAPACK tools routine which initializes the descriptors for the local arrays.

```
USE MPI SETUP INT
USE LQRRR_INT
USE LQRSL_INT
USE WRIRN_INT
USE WRRRN_INT
USE SCALAPACK SUPPORT
IMPLICIT NONE
INCLUDE 'mpif.h'
    Declare variables
INTEGER KBASIS, LDA, LDQR, NCA, NRA, DESCA(9), DESCL(9), &
    DESCX(9), DESCB(9)
INTEGER INFO, MXCOL, MXCOLX, MXLDA, MXLDX, LDQ, IPATH
INTEGER, ALLOCATABLE :: IPVT(:), IPVTO(:)
REAL, ALLOCATABLE :: A(:,:), B(:), QR(:,:), QRAUX(:), X(:), &
    RES(:)
REAL, ALLOCATABLE :: AO(:,:), QRO(:,:), QRAUXO(:), X0(:), &
    RESO(:), BO(:), QTB0(:)
LOGICAL PIVOT
PARAMETER (NRA=4, NCA=3, LDA=NRA, LDQR=NRA, KBASIS=3)
                            Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP RANK .EQ.- 0) THEN
    A\overline{LLOCATE (A (LDA,NCA), B(NRA), QR(LDQR,NCA), &}
    QRAUX(NCA), IPVT(NCA), X(NCA), RES(NRA))
                    Set values for A and the righthand sides
    A(1,:) = (/ 1.0, 2.0, 4.0/)
    A(2,:) = (/ 1.0, 4.0, 16.0/)
    A(3,:) = (/ 1.0, 6.0, 36.0/)
    A(4,:) = (/ 1.0, 8.0, 64.0/)
    B = (/ 16.99, 57.01, 120.99, 209.01 /)
    IPVT = 0
ENDIF
    Set up a 1D processor grid and define
    its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(NRA, NCA, .TRUE., .TRUE.)
                        Get the array descriptor entities MXLDA,
                    and MXCOL
CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
CALL SCALAPACK_GETDIM(KBASIS, 1, MP_NB, 1, MXLDX, MXCOLX)
    Set up the array descriptors
CALL DESCINIT(DESCA, NRA, NCA, MP_MB, MP_NB, 0, 0, MP_ICTXT, &
            MXLDA, INFO)
CALL DESCINIT(DESCL, 1, NCA, 1, MP_NB, 0, 0, MP_ICTXT, 1, INFO)
CALL DESCINIT(DESCX, KBASIS, 1, MP_NB, 1, 0, 0, MP_ICTXT, MXLDX, INFO)
CALL DESCINIT (DESCB, NRA, 1, MP MB, 1, 0, 0, MP ICTXXT, MXLDA, INFO)
                        Allocate space for the local arrays
ALLOCATE (A0 (MXLDA,MXCOL), QRO (MXLDA,MXCOL), QRAUXO (MXCOL), &
IPVTO (MXCOL), B0 (MXLDA), X0 (MXLDX), RESO (MXLDA), QTBO (MXLDA))
    Map input array to the processor grid
```

```
CALL SCALAPACK_MAP(A, DESCA, A0)
CALL SCALAPACK_MAP(B, DESCB, BO)
PIVOT = .TRUE.
CALL SCALAPACK_MAP(IPVT, DESCL, IPVTO)
    QR factorization
    CALL LQRRR (AO, QR0, QRAUX0, PIVOT=PIVOT, IPVT=IPVTO)
    IPATH = 00110
CALL LQRSL (KBASIS, QR0, QRAUX0, B0, IPATH, QTB=QTB0, X=X0, RES=RES0)
    Unmap the results from the distributed
    array back to a non-distributed array.
    After the unmap, only Rank=0 has the full
    array.
    CALL SCALAPACK UNMAP(IPVTO, DESCL, IPVT, NCA, .FALSE.)
    CALL SCALAPACK UNMAP(X0, DESCX, X)
    CALL SCALAPACK_UNMAP(RESO, DESCB, RES)
                            Print results.
                            Only Rank=0 has the solution, X.
IF(MP_RANK .EQ. O) THEN
    CALLL WRIRN ('IPVT', IPVT, 1, NCA, 1)
    CALL WRRRN ('X', X, 1, KBASIS, 1)
    CALL WRRRN ('RES', RES, 1, NRA, 1)
ENDIF
! Exit ScaLAPACK usage
CALL SCALAPACK_EXIT(MP_ICTXT)
Shut down MPI
MP NPROCS = MP SETUP('FINAL')
END
```


## Output

```
\begin{tabular}{ccc}
\multicolumn{3}{c}{ IPVT } \\
1 & 2 & 3 \\
3 & 2 & 1
\end{tabular}
X
\begin{tabular}{rrr}
1 & 2 & 3 \\
3.000 & 2.002 & 0.990
\end{tabular}
\begin{tabular}{rrrrr} 
& \multicolumn{4}{c}{ RES } \\
-0.00400 & 0.01200 & -0.01200 & 0.00400
\end{tabular}
```

Note that since IPVT is $(3,2,1)$ the array X contains the solution coefficients $c_{i}$ in reverse order.

## LUPQR

Computes an updated $Q R$ factorization after the rank-one matrix $\alpha x y^{T}$ is added.

## Required Arguments

ALPHA - Scalar determining the rank-one update to be added. (Input)
$W$ - Vector of length NROW determining the rank-one matrix to be added. (Input) The updated matrix is $A+\alpha x y^{T}$. If $\mathrm{I}=0$ then $W$ contains the vector $x$. If $\mathrm{I}=1$ then $W$ contains the vector $Q^{T} x$.
$\boldsymbol{Y}$ - Vector of length NCOL determining the rank-one matrix to be added. (Input)
$\boldsymbol{R}$ - Matrix of order NROW by NCOL containing the $R$ matrix from the $Q R$ factorization. (Input)
Only the upper trapezoidal part of $R$ is referenced.
IPATH - Flag used to control the computation of the $Q R$ update. (Input) IPATH has the decimal expansion IJ such that: I $=0$ means $W$ contains the vector $x$. $\mathrm{I}=1$ means w contains the vector $Q^{T} x$.
$\mathrm{J}=0$ means do not update the matrix $Q . \mathrm{J}=1$ means update the matrix $Q$. For example, if $I$ РАТН $=10$ then, $I=1$ and $J=0$.
$\boldsymbol{R N E W}$ - Matrix of order NROW by NCOL containing the updated R matrix in the $Q R$ factorization. (Output)
Only the upper trapezoidal part of RNEW is updated. R and RNEW may be the same.

## Optional Arguments

NROW - Number of rows in the matrix $A=Q * R$. (Input)
Default: NROW $=\operatorname{size}(\mathrm{W}, 1)$.
NCOL - Number of columns in the matrix $A=Q * R$. (Input) Default: $\mathrm{NCOL}=\operatorname{size}(\mathrm{Y}, 1)$.
$\boldsymbol{Q}$ - Matrix of order NRow containing the $Q$ matrix from the $Q R$ factorization. (Input)
Ignored if IPATH $=0$.
Default: Q is 1 x 1 and un-initialized.
LDQ - Leading dimension of $Q$ exactly as specified in the dimension statement of the calling program. (Input)
Ignored if IPATH $=0$.
Default: LDQ $=\operatorname{size}(\Omega, 1)$.
$\boldsymbol{L D R}$ - Leading dimension of $R$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDR $=\operatorname{size}(\mathrm{R}, 1)$.
QNEW - Matrix of order NROW containing the updated Q matrix in the $Q R$ factorization. (Output)
Ignored if $J=0$, see IPATH for definition of $J$.
LDQNEW - Leading dimension of QNEW exactly as specified in the dimension statement of the calling program. (Input)

Ignored if $J=0$; see IPATH for definition of $J$.
Default: LDQNEW = size (QNEW,1).
LDRNEW - Leading dimension of RNEW exactly as specified in the dimension statement of the calling program. (Input)
Default: LDRNEW = size (RNEW,1).

## FORTRAN 90 Interface

Generic: CALL LUPQR (ALPHA, W, Y, R, IPATH, RNEW [,...])

Specific: The specific interface names are $S_{-} \angle U P Q R$ and $D_{-} L U P Q R$.

## FORTRAN 77 Interface

Single: CALL LUPQR (NROW, NCOL, ALPHA, W, Y, Q, LDQ, R, LDR, IPATH, QNEW, LDQNEW, RNEW, LDRNEW)

Double: The double precision name is DLUPQR.

## Description

Let $A$ be an $m \times n$ matrix and let $A=Q R$ be its $Q R$ decomposition. (In the program, $m$ is called NROW and $n$ is called NCOL) Then

$$
A+\alpha x y^{T}=Q R+\alpha x y^{T}=Q\left(R+\alpha Q^{T} x y^{T}\right)=Q\left(R+\alpha w y^{T}\right)
$$

where $w=Q^{T} x$. An orthogonal transformation $J$ can be constructed, using a sequence of $m-1$ Givens rotations, such that $J w=\omega e_{1}$, where $\omega= \pm\|w\|_{2}$ and $e_{1}=(1,0, \ldots, 0)^{T}$. Then

$$
A+\alpha x y^{T}=\left(Q J^{T}\right)\left(J R+\alpha \omega e_{1} y^{T}\right)
$$

Since $J R$ is an upper Hessenberg matrix, $H=J R+\alpha \omega e_{1} y^{T}$ is also an upper Hessenberg matrix. Again using $m-1$ Givens rotations, an orthogonal transformation $G$ can be constructed such that $G H$ is an upper triangular matrix. Then

$$
A+\alpha x y^{T}=\tilde{Q} \tilde{R}, \text { where } \tilde{Q}=Q J^{T} G^{T}
$$

is orthogonal and

$$
\tilde{R}=G H
$$

is upper triangular.
If the last $k$ components of $w$ are zero, then the number of Givens rotations needed to construct $J$ or $G$ is $m-k-1$ instead of $m-1$.

For further information, see Dennis and Schnabel (1983, pages 55-58 and 311-313), or Golub and Van Loan (1983, pages 437-439).

## Comments

1. Workspace may be explicitly provided, if desired, by use of $\angle 2 P Q R / D L 2 P Q R$. The reference is:

CALL L2PQR (NROW, NCOL, ALPHA, W, Y, Q, LDQ, R, LDR, IPATH, QNEW, LDQNEW, RNEW, LDRNEW, Z, WORK)

The additional arguments are as follows:
Z — Work vector of length NROW.

WORK - Work vector of length MIN(NROW - 1, NCOL).

## Example

The $Q R$ factorization of $A$ is found. It is then used to find the $Q R$ factorization of $A+x y^{T}$. Since pivoting is used, the $Q R$ factorization routine finds $A P=Q R$, where $P$ is a permutation matrix determined by IPVT. We compute

$$
A P+\alpha x y^{T}=\left(A+\alpha x(P y)^{T}\right) P=\tilde{Q} \tilde{R}
$$

The IMSL routine Permu (see Utilities) is used to compute $P y$. As a check
$\tilde{Q} \tilde{R}$
is computed and printed. It can also be obtained from $A+x y^{T}$ by permuting its columns using the order given by IPVT.

```
USE IMSL_LIBRARIES
INTEGER LDA, LDAQR, LDQ, LDQNEW, LDQR, LDR, LDRNEW, NCOL, NROW
PARAMETER (NCOL=3, NROW=4, LDA=NROW, LDAQR=NROW, LDQ=NROW, &
        LDQNEW=NROW, LDQR=NROW, LDR=NROW, LDRNEW=NROW)
    INTEGER IPATH, IPVT(NCOL), J, MINO
REAL A(LDA,NCOL), ALPHA, AQR(LDAQR,NCOL), CONORM(NCOL), &
        Q (LDQ,NROW), QNEW(LDQNEW,NROW), QR(LDQR,NCOL), &
        QRAUX(NCOL), R(LDR,NCOL), RNEW(LDRNEW,NCOL), W (NROW) , &
        Y(NCOL)
    LOGICAL 
    LOGICAL 
    Set values for A
            A =(\begin{array}{lllr}{1}&{2}&{4}\end{array})
            ( 1
            (\begin{array}{lll}{1}&{6}&{36}\end{array})
    DATA A/4*1.0, 2.0, 4.0, 6.0, 8.0, 4.0, 16.0, 36.0, 64.0/
    Set values for W and Y
DATA W/1., 2., 3., 4./
```

$!$
$!$
$!$
$!$
$!$
$!$
!

```
    DATA Y/3., 2., 1./
                                    Q factorization
                                    Set IPVT = 0 (all columns free)
    IPVT = 0
    PIVOT = .TRUE.
    CALL LQRRR (A, QR, QRAUX, IPVT=IPVT, PIVOT=PIVOT)
    CALL LQERR (QR, QRAUX, Q)
    CALL PERMU (Y, IPVT, Y)
        R = 0.0E0
        DO 10 J=1, NCOL
        CALL SCOPY (MINO (J,NROW), QR(:,J), 1, R(:,J), 1)
    10 CONTINUE
        ALPHA = 1.0
        IPATH = 01
        CALL LUPQR (ALPHA, W, Y, R, IPATH, RNEW, Q=Q, QNEW=QNEW)
        Compute AQR = Q*R
        CALL MRRRR (QNEW, RNEW, AQR)
        Print results
        CALL WRIRN ('IPVT', IPVT, 1, NCOL,1)
        CALL WRRRN ('QNEW', QNEW)
        CALL WRRRN ('RNEW', RNEW)
        CALL WRRRN ('QNEW*RNEW', AQR)
        END
```


## Output

```
IPVT
    2 3
    2 1
```

        QNEW
    | 1 | 2 | 3 | 4 |
| ---: | ---: | ---: | ---: |
| -0.0620 | -0.5412 | 0.8082 | -0.2236 |
| -0.2234 | -0.6539 | -0.2694 | 0.6708 |
| -0.4840 | -0.3379 | -0.4490 | -0.6708 |
| -0.8438 | 0.4067 | 0.2694 | 0.2236 |

        RNEW
    $-80.59-21.34 \quad-17.62$
$0.00-4.94-4.83$
$0.00 \quad 0.00 \quad 0.36$
$0.00 \quad 0.00 \quad 0.00$
QNEW*RNEW

| 1 | 2 | 3 |
| :--- | :--- | :--- |

    \(8.00-7.00\)
    \(39.00 \quad 12.00 \quad 10.00\)
    \(68.00 \quad 16.00 \quad 13.00\)
    
## LCHRG

Computes the Cholesky decomposition of a symmetric positive definite matrix with optional column pivoting.

## Required Arguments

$\boldsymbol{A}-\mathrm{N}$ by N symmetric positive definite matrix to be decomposed. (Input)
Only the upper triangle of A is referenced.
FACT — N by N matrix containing the Cholesky factor of the permuted matrix in its upper triangle. (Output)
If $A$ is not needed, A and FACT can share the same storage locations.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
PIVOT - Logical variable. (Input)
PIVOT $=$. TRUE. means column pivoting is done. PIVOT $=$. FALSE. means no pivoting is done.
Default: PIVOT = .TRUE.

IPVT - Integer vector of length N containing information that controls the selection of the pivot columns. (Input/Output)
On input, if $\operatorname{IPVT}(K)>0$, then the $K$-th column of A is an initial column; if $\operatorname{IPVT}(K)=0$, then the $K$-th column of A is a free column; if $\operatorname{IPVT}(K)<0$, then the $K$-th column of A is a final column. See Comments. On output, $\operatorname{IPVT}(K)$ contains the index of the diagonal element of $A$ that was moved into the $K$-th position. IPVT is only referenced when PIVOT is equal to .TRUE..

LDFACT - Leading dimension of FACT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFACT $=\operatorname{size}($ FACT, 1$)$.

## FORTRAN 90 Interface

Generic: CALL LCHRG (A, FACT [,...])
Specific: The specific interface names are S_LCHRG and D_LCHRG.

## FORTRAN 77 Interface

Single: CALL LCHRG (N, A, LDA, PIVOT, IPVT, FACT, LDFACT)
Double: The double precision name is DLCHRG.

## Description

Routine LCHRG is based on the LINPACK routine SCHDC; see Dongarra et al. (1979).
Before the decomposition is computed, initial elements are moved to the leading part of $A$ and final elements to the trailing part of $A$. During the decomposition only rows and columns corresponding to the free elements are moved. The result of the decomposition is an upper triangular matrix $R$ and a permutation matrix $P$ that satisfy $P^{T} A P=R^{T} R$, where $P$ is represented by IPVT.

## Comments

1. Informational error Type Code

41 The input matrix is not positive definite.
2. Before the decomposition is computed, initial elements are moved to the leading part of A and final elements to the trailing part of A. During the decomposition only rows and columns corresponding to the free elements are moved. The result of the decomposition is an upper triangular matrix R and a permutation matrix P that satisfy $P^{T} A P=R^{T} R$, where P is represented by IPVT.
3. LCHRG can be used together with subroutines PERMU and LSLDS to solve the positive definite linear system $\mathrm{AX}=\mathrm{B}$ with the solution $X$ overwriting the right-hand side $B$ as follows:

```
CALL ISET (N, 0,IPVT, 1)
CALL LCHRG (A, FACT, N, LDA,.TRUE, IPVT, LDFACT)
CALL PERMU (B, IPVT, B, N, 1)
CALL LSLDS (FACT, B, B,N, LDFACT)
CALL PERMU (B, IPVT, B, N, 2)
```


## Example

Routine LCHRG can be used together with the IMSL routines PERMU (see Chapter 11) and LFSDS to solve a positive definite linear system $A x=b$. Since $A=P R^{T} R P$, the system $A x=b$ is equivalent to $R^{T} R(P x)=P b$. LFSDS is used to solve $R^{T} R y=P b$ for $y$. The routine PERMU is used to compute both $P b$ and $x=P y$.

```
USE IMSL_LIBRARIES
```

PARAMETER ( $\mathrm{N}=3$, LDA=N, LDFACT=N)
INTEGER IPVT (N)

```
REAL A(LDA,N), FACT(LDFACT,N), B(N), X(N)
LOGICAL PIVOT
    Set values for A and B
    A = (\begin{array}{llll}{1}&{-3}&{2}\end{array})
        (}\begin{array}{rrrr}{-3}&{10}&{-5}\end{array}
    B = ( 27 -78 64 )
DATA A/1.,-3.,2.,-3.,10.,-5.,2.,-5.,6./
DATA B/27.,-78.,64./
    Pivot using all columns
    PIVOT = .TRUE.
    IPVT = 0
    Compute Cholesky factorization
    CALL LCHRG (A, FACT, PIVOT=PIVOT, IPVT=IPVT)
    Permute B and store in X
    CALL PERMU (B, IPVT, X, IPATH=1)
    Solve for X
    CALL LFSDS (FACT, X, X)
    CALL PERMU (X, IPVT, X, IPATH=2)
    Print X
    CALL WRRRN ('X', X, 1, N, 1)
    END
```


## Output

|  | $x$ | 3 |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| 1.000 | -4.000 | 7.000 |

## LUPCH

Updates the $R^{T} R$ Cholesky factorization of a real symmetric positive definite matrix after a rankone matrix is added.

## Required Arguments

$\boldsymbol{R}-\mathrm{N}$ by N upper triangular matrix containing the upper triangular factor to be updated. (Input) Only the upper triangle of $R$ is referenced.
$\boldsymbol{X}$ - Vector of length N determining the rank-one matrix to be added to the factorization $R^{T} R$. (Input)

RNE $W-\mathrm{N}$ by N upper triangular matrix containing the updated triangular factor of $R^{T} R+X X^{T}$. (Output)

Only the upper triangle of RNEW is referenced. If $R$ is not needed, $R$ and RNEW can share the same storage locations.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{R}, 2)$.
$\boldsymbol{L D R}$ - Leading dimension of $R$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDR $=\operatorname{size}(R, 1)$.
LDRNEW - Leading dimension of RNEW exactly as specified in the dimension statement of the calling program. (Input)
Default: LDRNEW = size (RNEW,1).
$\boldsymbol{C S}$ - Vector of length N containing the cosines of the rotations. (Output)
$\boldsymbol{S N}$ - Vector of length N containing the sines of the rotations. (Output)

## FORTRAN 90 Interface

Generic: CALL LUPCH (R, X, RNEW [,...])
Specific: The specific interface names are S_LUPCH and D_LUPCH.

## FORTRAN 77 Interface

Single: CALL LUPCH ( $\mathrm{N}, \mathrm{R}, \mathrm{LDR}, \mathrm{X}$, RNEW, LDRNEW, CS, SN)
Double: The double precision name is DLUPCH.

## Description

The routine LUPCH is based on the LINPACK routine SChUD; see Dongarra et al. (1979).
The Cholesky factorization of a matrix is $A=R^{T} R$, where $R$ is an upper triangular matrix. Given this factorization, LUPCH computes the factorization

$$
A+x x^{T}=\tilde{R}^{T} \tilde{R}
$$

In the program
$\tilde{R}$
is called RNEW.
LUPCH determines an orthogonal matrix $U$ as the product $G_{N} \ldots G_{1}$ of Givens rotations, such that

$$
U\left[\begin{array}{l}
R \\
x^{T}
\end{array}\right]=\left[\begin{array}{l}
\tilde{R} \\
0
\end{array}\right]
$$

By multiplying this equation by its transpose, and noting that $U^{T} U=I$, the desired result

$$
R^{T} R+x x^{T}=\tilde{R}^{T} \tilde{R}
$$

is obtained.
Each Givens rotation, $G_{i}$, is chosen to zero out an element in $x^{T}$. The matrix
$G_{i}$ is $(N+1) \times(N+1)$ and has the form

$$
G_{i}=\left[\begin{array}{cccc}
I_{i-1} & 0 & 0 & 0 \\
0 & c_{i} & 0 & s_{i} \\
0 & 0 & I_{N-i} & 0 \\
0 & -s_{i} & 0 & c_{i}
\end{array}\right]
$$

Where $I_{k}$ is the identity matrix of order $k$ and $c_{i}=\cos \theta_{i}=\operatorname{CS}(\mathrm{I}), s_{i}=\sin \theta_{i}=\operatorname{SN}(\mathrm{I})$ for some $\theta_{i}$.

## Example

A linear system $A z=b$ is solved using the Cholesky factorization of $A$. This factorization is then updated and the system $\left(A+x x^{T}\right) z=b$ is solved using this updated factorization.

```
USE IMSL_LIBRARIES
INTEGER LDA, LDFACT, N
PARAMETER (LDA=3, LDFACT=3, N=3)
REAL A(LDA,LDA), FACT(LDFACT,LDFACT), FACNEW(LDFACT,LDFACT), &
    X(N), B(N), CS(N), SN(N), Z(N)
    Set values for A
            A = ( 1.0 -3.0 2.0)
            (-3.0 10.0 - -5.0)
            ( 2.0 -5.0 6.0)
DATA A/1.0, -3.0, 2.0, -3.0, 10.0, -5.0, 2.0, -5.0, 6.0/
                    Set values for X and B
DATA X/3.0, 2.0, 1.0/
DATA B/53.0, 20.0, 31.0/
                        Factor the matrix A
                            Solve the original system
                    Print the results
CALL WRRRN ('FACT', FACT, ITRING=1)
CALL WRRRN ('Z', Z, 1, N, 1)
                                    Update the factorization
CALL LUPCH (FACT, X, FACNEW)
CALL LFSDS (FACNEW, B, Z)
```

```
! Print the results
    CALL WRRRN ('FACNEW', FACNEW, ITRING=1)
    CALL WRRRN ('Z', Z, 1, N, 1)
!
    END
```


## Output


FACNEW
132
$3.162 \quad 0.949 \quad 1.581$
$3.619-1.243$
$-1.719$
Z

| 1 | 2 | 3 |
| :--- | :--- | :--- |

$4.000 \quad 1.000 \quad 2.000$

## LDNCH

Downdates the $R^{T} R$ Cholesky factorization of a real symmetric positive definite matrix after a rank-one matrix is removed.

## Required Arguments

$\boldsymbol{R}-\mathrm{N}$ by N upper triangular matrix containing the upper triangular factor to be downdated. (Input)
Only the upper triangle of $R$ is referenced.
$\boldsymbol{X}$ - Vector of length N determining the rank-one matrix to be subtracted from the factorization $R^{T} R$. (Input)
$\boldsymbol{R N E} W-\mathrm{N}$ by N upper triangular matrix containing the downdated triangular factor of $R^{T} R-X X^{T}$. (Output)
Only the upper triangle of RNEW is referenced. If $R$ is not needed, $R$ and RNEW can share the same storage locations.

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $N=\operatorname{size}(R, 2)$.
$\boldsymbol{L D R}$ - Leading dimension of $R$ exactly as specified in the dimension statement of the calling program. (Input) Default: LDR = size ( $\mathrm{R}, 1$ ).

LDRNEW - Leading dimension of RNEW exactly as specified in the dimension statement of the calling program. (Input)
Default: LDRNEW = size (RNEW,1).
$\boldsymbol{C S}$ - Vector of length N containing the cosines of the rotations. (Output)
$\boldsymbol{S N}$ - Vector of length N containing the sines of the rotations. (Output)

## FORTRAN 90 Interface

Generic: CALL LDNCH (R, X, RNEW [, ...])
Specific: The specific interface names are S_LDNCH and D_LDNCH.

## FORTRAN 77 Interface

Single: CALL LDNCH (N, R, LDR, X, RNEW, LDRNEW, CS, SN)
Double: The double precision name is DLDNCH.

## Description

The routine LDNCH is based on the LINPACK routine SCHDD; see Dongarra et al. (1979).
The Cholesky factorization of a matrix is $A=R^{T} R$, where $R$ is an upper triangular matrix. Given this factorization, $\operatorname{LDNCH}$ computes the factorization

$$
A-x x^{T}=\tilde{R}^{T} \tilde{R}
$$

In the program

$$
\tilde{R}
$$

is called RNEW. This is not always possible, since $A-x x^{T}$ may not be positive definite. LDNCH determines an orthogonal matrix $U$ as the product $G_{N} \ldots G_{1}$ of Givens rotations, such that

$$
U\left[\begin{array}{l}
R \\
0
\end{array}\right]=\left[\begin{array}{c}
\tilde{R}^{T} \\
x^{T}
\end{array}\right]
$$

By multiplying this equation by its transpose and noting that $U^{T} U=I$, the desired result

$$
R^{T} R-x x^{T}=\tilde{R}^{T} \tilde{R}
$$

is obtained.
Let $a$ be the solution of the linear system $R^{T} a=x$ and let

$$
\alpha=\sqrt{1-\|a\|_{2}^{2}}
$$

The Givens rotations, $G_{i}$, are chosen such that

$$
G_{1} \cdots G_{N}\left[\begin{array}{c}
a \\
\alpha
\end{array}\right]=\left[\begin{array}{l}
0 \\
1
\end{array}\right]
$$

The $G_{i}$ are $(N+1) \times(N+1)$ matrices of the form

$$
G_{i}=\left[\begin{array}{cccc}
I_{i-1} & 0 & 0 & 0 \\
0 & c_{i} & 0 & -s_{i} \\
0 & 0 & I_{N-i} & 0 \\
0 & s_{i} & 0 & c_{i}
\end{array}\right]
$$

where $I_{k}$ is the identity matrix of order $k$; and $c_{i}=\cos \theta_{i}=\operatorname{CS}(\mathrm{I}), s_{i}=\sin \theta_{i}=\operatorname{SN}(\mathrm{I})$ for some $\theta_{i}$.
The Givens rotations are then used to form

$$
\tilde{R}, G_{1} \cdots G_{N}\left[\begin{array}{l}
R \\
0
\end{array}\right]=\left[\begin{array}{l}
\tilde{R} \\
\tilde{x}^{T}
\end{array}\right]
$$

The matrix

$$
\tilde{R}
$$

is upper triangular and

$$
\tilde{x}=x
$$

because

$$
x=\left(R^{T} 0\right)\left[\begin{array}{l}
a \\
\alpha
\end{array}\right]=\left(R^{T} 0\right) U^{T} U\left[\begin{array}{l}
a \\
\alpha
\end{array}\right]=\left(\tilde{R}^{T} \tilde{x}\right)\left[\begin{array}{l}
0 \\
1
\end{array}\right]=\tilde{x}
$$

## Comments

Informational error
Type Code
$41 \quad R^{T} R-X X^{T}$ is not positive definite. R cannot be downdated.

## Example

A linear system $A z=b$ is solved using the Cholesky factorization of $A$. This factorization is then downdated, and the system $\left(A-x x^{T}\right) z=b$ is solved using this downdated factorization.

```
USE LDNCH_INT
USE LFTDS ' INT
USE LFSDS_INT
USE WRRRN_INT
```

```
        INTEGER LDA, LDFACT, N
        PARAMETER (LDA=3, LDFACT=3, N=3)
        REAL A(LDA,LDA), FACT(LDFACT,LDFACT), FACNEW(LDFACT,LDFACT), &
        X(N), B(N), CS(N), SN(N), Z(N)
        Set values for A
        A = ( 10.0 3.0 5.0)
            ( 3.0 14.0 -3.0)
            ( 5.0 -3.0 7.0)
    DATA A/10.0, 3.0, 5.0, 3.0, 14.0, -3.0, 5.0, -3.0, 7.0/
    DATA X/3.0, 2.0, 1.0/
    DATA B/53.0, 20.0, 31.0/
        Factor the matrix A
    CALL LFTDS (A, FACT)
    CALL LFSDS (FACT, B, Z)
    CALL WRRRN ('FACT', FACT, ITRING=1)
    CALL WRRRN ('Z', Z, 1, N, 1)
        Downdate the factorization
    CALL LDNCH (FACT, X, FACNEW)
        Solve the updated system
        Print the results
    CALL WRRRN ('FACNEW', FACNEW, ITRING=1)
    CALL WRRRN ('Z', Z, 1, N, 1)
    END
```


## Output

FACT

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 3.162 | 0.949 | 1.581 |
| 2 |  | 3.619 | -1.243 |
| 3 |  |  | 1.719 |


| 1 | 2 | 3 |
| ---: | ---: | ---: |
| 4.000 | 1.000 | 2.000 |

FACNEW

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 1.000 | -3.000 | 2.000 |
| 2 |  | 1.000 | 1.000 |
| 3 |  |  | 1.000 |


| Z |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| 1859.9 | 433.0 | -254.0 |

## LSVRR

## 匀MPI <br> CAPABLE

Computes the singular value decomposition of a real matrix.

## Required Arguments

A - NRA by NCA matrix whose singular value decomposition is to be computed. (Input)
IPATH - Flag used to control the computation of the singular vectors. (Input)
IPATH has the decimal expansion IJ such that:
$I=0$ means do not compute the left singular vectors;
$I=1$ means return the NRA left singular vectors in U;
NOTE: This option is not available for the ScaLAPACK interface. If this option is chosen for ScaLAPACK usage, the min(NRA, NCA) left singular vectors will be returned.
$I=2$ means return only the $\min (\mathrm{NRA}, \mathrm{NCA})$ left singular vectors in U ;
$J=0$ means do not compute the right singular vectors,
$J=1$ means return the right singular vectors in $v$.
For example, IPATH $=20$ means $I=2$ and $J=0$.
$\boldsymbol{S}$ - Vector of length min(NRA $+1, \mathrm{NCA})$ containing the singular values of A in descending order of magnitude in the first $\min (N R A, N C A)$ positions. (Output)

## Optional Arguments

$N R A$ - Number of rows in the matrix A. (Input)
Default: NRA $=$ size (A, 1 ).
NCA - Number of columns in the matrix A. (Input)
Default: NCA $=\operatorname{size}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA = size (A, 1).

TOL - Scalar containing the tolerance used to determine when a singular value is negligible. (Input)
If TOL is positive, then a singular value $\sigma_{i}$ considered negligible if $\sigma_{i} \leq T O L$. If TOL is negative, then a singular value $\sigma_{i}$ considered negligible if $\sigma_{i} \leq|T O L| *\|\AA\|_{\infty}$. In this case, $|T O L|$ generally contains an estimate of the level of the relative error in the data. Default: $\mathrm{TOL}=1.0 \mathrm{e}-5$ for single precision and $1.0 \mathrm{~d}-10$ for double precision.

IRANK - Scalar containing an estimate of the rank of A. (Output)
$\boldsymbol{U}$ - NRA by NCU matrix containing the left singular vectors of A. (Output) NCU must be equal to NRA if $I$ is equal to 1 . NCU must be equal to min(NRA, NCA) if I is equal to 2 . $U$ will not be referenced if $I$ is equal to zero. If NRA is less than or equal to NCU, then U can share the same storage locations as A. See Comments.
$\boldsymbol{L D U}$ - Leading dimension of U exactly as specified in the dimension statement of the calling program. (Input)
Default: LDU = size (U,1).
$\boldsymbol{V}$ — NCA by NCA matrix containing the right singular vectors of A. (Output)
V will not be referenced if J is equal to zero. V can share the same storage location as
A , however, U and V cannot both coincide with A simultaneously.
$\boldsymbol{L D V}$ - Leading dimension of v exactly as specified in the dimension statement of the calling program. (Input)
Default: LDV = size (v,1).

## FORTRAN 90 Interface

Generic: CALL LSVRR (A, IPATH, S [ ,...])
Specific: The specific interface names are S_LSVRR and D_LSVRR.

## FORTRAN 77 Interface

Single: CALL LSVRR (NRA, NCA, A, LDA, IPATH, TOL, IRANK, S, U, LDU, V, LDV)
Double: The double precision name is DLSVRR.

## ScaLAPACK Interface

Generic: CALL LSVRR (A0, IPATH, S [,...])
Specific: The specific interface names are S_LSVRR and D_LSVRR.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

Let $n=$ NRA (the number of rows in $A$ ) and let $p=$ NCA (the number of columns in $A$ ). For any $n \times$ $p$ matrix $A$, there exists an $n \times n$ orthogonal matrix $U$ and a $p \times p$ orthogonal matrix $V$ such that

$$
U^{T} A V= \begin{cases}{\left[\begin{array}{l}
\Sigma \\
0
\end{array}\right]} & \text { if } n \geq p \\
{[\Sigma 0]} & \text { if } n \leq p\end{cases}
$$

where $\Sigma=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{m}\right)$, and $m=\min (n, p)$. The scalars $\sigma_{1} \geq \sigma_{2} \geq \ldots \geq \sigma_{m} \geq 0$ are called the singular values of $A$. The columns of $U$ are called the left singular vectors of $A$. The columns of $V$ are called the right singular vectors of $A$.
The estimated rank of $A$ is the number of $\sigma_{k}$ that is larger than a tolerance $\eta$. If $\tau$ is the parameter TOL in the program, then

$$
\eta= \begin{cases}\tau & \text { if } \tau>0 \\ |\tau|\|A\|_{\infty} & \text { if } \tau<0\end{cases}
$$

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2VRR/DL2VRR. The reference is:
```
CALL L2VRR (NRA, NCA, A, LDA, IPATH, TOL, IRANK, S, U, LDU, V, LDV,
    ACOPY,WK)
```

The additional arguments are as follows:
ACOPY - NRA $\times$ NCA work array for the matrix A. If A is not needed, then A and ACOPY may share the same storage locations.
$W K-$ Work vector of length NRA + NCA $+\max ($ NRA, NCA $)-1$.
2. Informational error

Type Code
41 Convergence cannot be achieved for all the singular values and their corresponding singular vectors.
3. When NRA is much greater than NCA, it might not be reasonable to store the whole matrix $U$. In this case, IPATH with $I=2$ allows a singular value factorization of $A$ to be computed in which only the first NCA columns of $U$ are computed, and in many applications those are all that are needed.
4. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2VRR the leading dimension of ACOPY is increased by $\operatorname{IVAL}(3)$ when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSVRR. Additional memory allocation for ACOPY and option value restoration are done automatically in LSVRR. Users directly calling L2VRR can allocate additional space for ACOPY and set IVAL(3) and IVAL(4) so that memory bank conflicts no
longer cause inefficiencies. There is no requirement that users change existing applications that use LSVRR or L2VRR. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1$.

17 This option has two values that determine if the $L_{1}$ condition number is to be computed. Routine LSVRR temporarily replaces IVAL(2) by IVAL(1). The routine L2CRG computes the condition number if IVAL $(2)=2$. Otherwise L2CRG skips this computation. LSVRR restores the option. Default values for the option are $\operatorname{IVAL}(*)=1,2$.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:
A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix whose singular value decomposition is to be computed. (Input)

UO - MXLDU by MXCOLU local matrix containing the local portions of the left singular vectors of the distributed matrix A. (Output)
UO will not be referenced if $I$ is equal to zero. If NRA is less than or equal to NCU, then U0 can share the same storage locations as A0. See Comments.

V0 - MXLDV by MXCOLV local matrix containing the local portions of the right singular vectors of the distributed matrix A. (Output)
V0 will not be referenced if $J$ is equal to zero. V0 can share the same storage location as $\mathbf{A 0}$, however, $\boldsymbol{U 0}$ and $\mathbf{V O}$ cannot both coincide with $\mathbf{A 0}$ simultaneously.

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXCOL, MXLDU, MXCOLU, MXLDV and MXCOLV can be obtained through a call to ScaLAPACK_GETDIM (see Chapter 11, "Utilities") after a call to ScaLAPACK_SETUP (see Chapter 11, "Utilities") has been made. See the ScaLAPACK Example below.

## Example

This example computes the singular value decomposition of a $6 \times 4$ matrix $A$. The matrices $U$ and $V$ containing the left and right singular vectors, respectively, and the diagonal of $\Sigma$, containing singular values, are printed. On some systems, the signs of some of the columns of $U$ and $V$ may be reversed.

USE IMSL_LIBRARIES


```
! (llllll
    DATA A/1., 3., 4., 2., 1., 1., 2., 2., 3., 1., 5., 2., 3*1., &
    3., 2., 2., 4., 3., 4., 1., 2., 3./
    IPATH = 11
    TOL = AMACH (4)
    TOL = 10.*TOL
    CALL LSVRR(A, IPATH, S, TOL=TOL, IRANK=IRANK, U=U, V=V)
    Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT, *) 'IRANK = ', IRANK
    CALL WRRRN ('U', U, NRA, NCA)
    CALL WRRRN ('S', S, 1, NCA, 1)
    CALL WRRRN ('V', V)
!
END
```


## Output

IRANK $=4$

|  |  | U |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 |
| 1 | -0.3805 | 0.1197 | 0.4391 | -0.5654 |
| 2 | -0.4038 | 0.3451 | -0.0566 | 0.2148 |
| 3 | -0.5451 | 0.4293 | 0.0514 | 0.4321 |
| 4 | -0.2648 | -0.0683 | -0.8839 | -0.2153 |
| 5 | -0.4463 | -0.8168 | 0.1419 | 0.3213 |
| 6 | -0.3546 | -0.1021 | -0.0043 | -0.5458 |

S

| 1 | 2 | 3 | 4 |
| ---: | ---: | ---: | ---: |
| 11.49 | 3.27 | 2.65 | 2.09 |


|  |  | V |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 |
| 1 | -0.4443 | 0.5555 | -0.4354 | 0.5518 |
| 2 | -0.5581 | -0.6543 | 0.2775 | 0.4283 |
| 3 | -0.3244 | -0.3514 | -0.7321 | -0.4851 |
| 4 | -0.6212 | 0.3739 | 0.4444 | -0.5261 |

## ScaLAPACK Example

The previous example is repeated here as a distributed example. This example computes the singular value decomposition of a $6 \times 4$ matrix $A$. The matrices $U$ and $V$ containing the left and right singular vectors, respectively, and the diagonal of $S$, containing singular values, are printed. On some systems, the signs of some of the columns of $U$ and $V$ may be reversed..

```
USE MPI_SETUP_INT
USE IMSL_LIBRARIES
USE SCALAPACK_SUPPORT
```

```
    IMPLICIT NONE
    INCLUDE 'mpif.h'
                Declare variables
    INTEGER KBASIS, LDA, LDQR, NCA, NRA, DESCA(9), DESCU(9), &
    DESCV(9), MXLDV, MXCOLV, NSZ, MXLDU, MXCOLU
    INTEGER INFO, MXCOL, MXLDA, LDU, LDV, IPATH, IRANK
    REAL TOL, AMACH
    REAL, ALLOCATABLE :: A(:,:),U(:,:), V(:,:), S(:)
    REAL, ALLOCATABLE :: AO(:,:), U0(:,:), VO(:,:), SO(:)
    PARAMETER (NRA=6, NCA=4, LDA=NRA, LDU=NRA, LDV=NCA)
    NSZ = MIN(NRA,NCA)
                            Set up for MPI
    MP_NPROCS = MP SETUP()
    IF(MP RANK .EQ- O) THEN
        ALLOCATE (A(LDA,NCA), U(LDU,NCA), V(LDV,NCA), S (NCA))
                            Set values for A
        A(1,:) = (/ 1.0, 2.0, 1.0, 4.0/)
        A(2,:) = (/ 3.0, 2.0, 1.0, 3.0/)
        A(3,:) = (/ 4.0, 3.0, 1.0, 4.0/)
        A(4,:) = (/ 2.0, 1.0, 3.0, 1.0/)
        A(5,:) = (/ 1.0, 5.0, 2.0, 2.0/)
        A(6,:) = (/ 1.0, 2.0, 2.0, 3.0/)
    ENDIF
        Set up a 1D processor grid and define
        its context ID, MP ICTXT
    CALL SCALAPACK_SETUP(NRA, NCA, .TRUE., .TRUE.)
                Get the array descriptor entities MXLDA,
                    MXCOL, MXLDU, MXCOLU, MXLDV, AND MXCOLV
CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
CALL SCALAPACK GETDIM(NRA, NSZ, MP MB, MP NB, MXLDU, MXCOLU)
CALL SCALAPACK_GETDIM(NSZ, NCA, MP_MB, MP_NB, MXLDV, MXCOLV)
                Set up the array descriptors
    CALL DESCINIT(DESCA, NRA, NCA, MP MB, MP NB, 0, 0, MP_ICTXT, &
    MXLDA, INFO)
    CALL DESCINIT(DESCU, NRA, NSZ, MP MB, MP NB, 0, 0, MP ICTXT, &
    MXLDU, INFO)
    CALL DESCINIT(DESCV, NSZ, NCA, MP_MB, MP_NB, 0, 0, MP_ICTXT, &
    MXLDV, INFO)
    Allocate space for the local arrays
    ALLOCATE (A0 (MXLDA,MXCOL), U0 (MXLDU,MXCOLU), VO (MXLDV,MXCOLV), S (NCA))
                            Map input array to the processor grid
    CALL SCALAPACK_MAP(A, DESCA, A0)
        Compute all singular vectors
    IPATH = 11
    TOL = AMACH (4)
    TOL = 10. * TOL
    CALL LSVRR (A0, IPATH, S, TOL=TOL, IRANK=IRANK, U=U0, V=V0)
    Unmap the results from the distributed
    array back to a non-distributed array.
    After the unmap, only Rank=0 has the full
    array.
    CALL SCALAPACK UNMAP(U0, DESCU, U)
    CALL SCALAPACK_UNMAP(VO, DESCV, V)
    Print results.
    Only Rank=0 has the solution.
```

```
    IF(MP_RANK .EQ. O) THEN
        CALL WRRRN ('U', U, NRA, NCA)
        CALL WRRRN ('S', S, 1, NCA, 1)
        CALL WRRRN ('V', V)
        ENDIF
! Exit ScaLAPACK usage
    CALL SCALAPACK_EXIT(MP_ICTXT)
    MP_NPROCS = MP_SETUP('FINAL')
    END
```


## Output

IRANK $=4$
4 U

|  | 1 | 2 | 3 | 4 |
| ---: | ---: | ---: | ---: | ---: |
| 1 | -0.3805 | 0.1197 | 0.4391 | -0.5654 |
| 2 | -0.4038 | 0.3451 | -0.0566 | 0.2148 |
| 3 | -0.5451 | 0.4293 | 0.0514 | 0.4321 |
| 4 | -0.2648 | -0.0683 | -0.8839 | -0.2153 |
| 5 | -0.4463 | -0.8168 | 0.1419 | 0.3213 |
| 6 | -0.3546 | -0.1021 | -0.0043 | -0.5458 |


| S |  |  |  |
| :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 |
| 11.49 | 3.27 | 2.65 | 2.09 |

$\left.\begin{array}{rrrrr} & & \mathrm{V} & & \\ & & 1 & 2 & 3\end{array}\right) 4$

## LSVCR

Computes the singular value decomposition of a complex matrix.

## Required Arguments

$\boldsymbol{A}$ - Complex NRA by NCA matrix whose singular value decomposition is to be computed. (Input)

IPATH - Integer flag used to control the computation of the singular vectors. (Input) IPATH has the decimal expansion IJ such that:

I=0 means do not compute the left singular vectors;
$\mathrm{I}=1$ means return the NCA left singular vectors in U ;
$I=2$ means return only the $\min (N R A, N C A)$ left singular vectors in $U$;
$J=0$ means do not compute the right singular vectors;
$J=1$ means return the right singular vectors in V.

For example, IPATH $=20$ means $I=2$ and $J=0$.
$\boldsymbol{S}$ - Complex vector of length min(NRA $+1, \mathrm{NCA}$ ) containing the singular values of A in descending order of magnitude in the first $\min (N R A, N C A)$ positions. (Output)

## Optional Arguments

$N R A$ - Number of rows in the matrix A. (Input)
Default: NRA $=\operatorname{size}(\mathrm{A}, 1)$.
$N C A$ - Number of columns in the matrix A. (Input)
Default: NCA $=\operatorname{size}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
TOL - Real scalar containing the tolerance used to determine when a singular value is negligible. (Input)
If $T O L$ is positive, then a singular value $S I$ is considered negligible if $S I \leq T O L$. If $T O L$ is negative, then a singular value SI is considered negligible if $S I \leq|T O L| *($ Infinity norm of A). In this case $|T O L|$ should generally contain an estimate of the level of relative error in the data.
Default: TOL $=1.0 \mathrm{e}-5$ for single precision and $1.0 \mathrm{~d}-10$ for double precision.
IRANK - Integer scalar containing an estimate of the rank of A. (Output)
$\boldsymbol{U}$ - Complex NRA by NRA if $I=1$ or NRA by $\min (N R A, N C A)$ if $I=2$ matrix containing the left singular vectors of A. (Output)
$U$ will not be referenced if $I$ is equal to zero. If NRA is less than or equal to NCA or IPATH $=2$, then $U$ can share the same storage locations as A.
$\boldsymbol{L D U}$ - Leading dimension of U exactly as specified in the dimension statement of the calling program. (Input)
Default: LDU = size (U,1).
$\boldsymbol{V}$ - Complex NCA by NCA matrix containing the right singular vectors of A. (Output) $V$ will not be referenced if $J$ is equal to zero. If NCA is less than or equal to NRA, then $V$ can share the same storage locations as $A$; however $U$ and $V$ cannot both coincide with $A$ simultaneously.
$\boldsymbol{L D V}$ - Leading dimension of v exactly as specified in the dimension statement of the calling program. (Input)
Default: LDV $=\operatorname{size}(\mathrm{V}, 1)$.

## FORTRAN 90 Interface

Generic: CALL LSVCR (A, IPATH, S [,...])

Specific: The specific interface names are S_LSVCR and D_LSVCR.

## FORTRAN 77 Interface

Single: CALL LSVCR (NRA, NCA, A, LDA, IPATH, TOL, IRANK, S, U, LDU, V, LDV)
Double: The double precision name is DLSVCR.

## Description

The underlying code is based on either LINPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

Let $n=$ NRA (the number of rows in $A$ ) and let $p=$ NCA (the number of columns in $A$ ). For any $n \times p$ matrix $A$ there exists an $n \times n$ orthogonal matrix $U$ and a $p \times p$ orthogonal matrix $V$ such that

$$
U^{T} A V= \begin{cases}{\left[\begin{array}{l}
\Sigma \\
0
\end{array}\right]} & \text { if } n \geq p \\
{[\Sigma 0]} & \text { if } n \leq p\end{cases}
$$

where $\Sigma=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{m}\right)$, and $m=\min (n, p)$. The scalars $\sigma_{1} \geq \sigma_{2} \geq \ldots \geq 0$ are called the singular values of $A$. The columns of $U$ are called the left singular vectors of $A$. The columns of $V$ are called the right singular vectors of $A$.

The estimated rank of $A$ is the number of $\sigma_{k}$ which are larger than a tolerance $\eta$. If $\tau$ is the parameter TOL in the program, then

$$
\eta= \begin{cases}\tau & \text { if } \tau>0 \\ |\tau|\|A\|_{\infty} & \text { if } \tau<0\end{cases}
$$

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2VCR/DL2VCR. The reference is
```
CALL L2VCR (NRA, NCA, A, LDA, IPATH, TOL, IRANK, S, U, LDU, V, LDV,
    ACOPY,WK)
```

The additional arguments are as follows:
ACOPY - NRA * NCA complex work array of length for the matrix A. If A is not needed, then A and ACOPY can share the same storage locations.
$W K-$ Complex work vector of length NRA + NCA $+\max ($ NRA, NCA $) 1$.
2. Informational error Type Code

41 Convergence cannot be achieved for all the singular values and their corresponding singular vectors.
3. When NRA is much greater than NCA, it might not be reasonable to store the whole matrix $U$. In this case IPATH with $I=2$ allows a singular value factorization of $A$ to be computed in which only the first NCA columns of $U$ are computed, and in many applications those are all that are needed.
4. Integer Options with Chapter 11 Options Manager

16 This option uses four values to solve memory bank conflict (access inefficiency) problems. In routine L2VCR the leading dimension of ACOPY is increased by $\operatorname{IVAL}(3)$ when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in LSVCR. Additional memory allocation for ACOPY and option value restoration are done automatically in LSVCR. Users directly calling L2VCR can allocate additional space for ACOPY and set IVAL(3) and IVAL(4) so that memory bank conflicts no longer cause inefficiencies. There is no requirement that users change existing applications that use LSVCR or L2VCR. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1$.

17 This option has two values that determine if the $L_{1}$ condition number is to be computed. Routine LSVCR temporarily replaces IVAL(2) by IVAL(1). The routine L2CCG computes the condition number if IVAL $(2)=2$. Otherwise L2CCG skips this computation. LSVCR restores the option. Default values for the option are $\operatorname{IVAL}(*)=1,2$.

## Example

This example computes the singular value decomposition of a $6 \times 3$ matrix $A$. The matrices $U$ and $V$ containing the left and right singular vectors, respectively, and the diagonal of $\Sigma$, containing singular values, are printed. On some systems, the signs of some of the columns of $U$ and $V$ may be reversed.

USE IMSL_LIBRARIES
PARAMETER (NRA=6, NCA=3, LDA=NRA, LDU=NRA, LDV=NCA)
COMPLEX A(LDA,NCA), U(LDU,NRA), V(LDV,NCA), S (NCA)
Set values for A
$A=\left(\begin{array}{cc}1+2 i & 3+2 i\end{array} \quad 1-4 i \quad\right)$
$\left(\begin{array}{cc}3-2 i & 2-4 i \\ 1+3 i\end{array}\right)$
$\left(\begin{array}{lll}4+3 i & -2+1 i & 1+4 i\end{array}\right)$
( $2-1 i \quad 3+0 i \quad 3-1 i$ )
$\left(\begin{array}{cc}1-5 i & 2-5 i \\ 2+2 i\end{array}\right)$
$\left(\begin{array}{ccc}1+2 i & 4-2 i & 2-3 i\end{array}\right)$
DATA A/ $(1.0,2.0),(3.0,-2.0),(4.0,3.0),(2.0,-1.0),(1.0,-5.0), \quad \&$

```
(1.0,2.0), (3.0,2.0), (2.0,-4.0), (-2.0,1.0), (3.0,0.0), &
(2.0,-5.0), (4.0,-2.0), (1.0,-4.0), (1.0,3.0), (1.0,4.0), &
(3.0,-1.0), (2.0,2.0), (2.0,-3.0)/
    TOL = AMACH (4)
    TOL = 10. * TOL
    CALL LSVCR(A, IPATH, S, TOL = TOL, IRANK=IRANK, U=U, V=V)
    Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT, *) 'IRANK = ', IRANK
    CALL WRCRN ('U', U, NRA, NCA)
    CALL WRCRN ('S', S, 1, NCA, 1)
    CALL WRCRN ('V', V)
```

!
IPATH $=11$
!
END

## Output

```
IRANK = 3
```

|  |  |  |  |
| ---: | ---: | ---: | ---: | ---: |
| 1 | $(0.1968,0.2186)$ | $(0.5011,0.0217)$ | $(-0.2007,-0.1003)$ |
| 2 | $(0.3443,-0.3542)$ | $(-0.2933,0.0248)$ | $(0.1155,-0.2338)$ |
| 3 | $(0.1457,0.2307)$ | $(-0.5424,0.1381)$ | $(-0.4361,-0.4407)$ |
| 4 | $(0.3016,-0.0844)$ | $(0.2157,0.2659)$ | $(-0.0523,-0.0894)$ |
| 5 | $(0.2283,-0.6008)$ | $(-0.1325,0.1433)$ | $(0.3152,-0.0090)$ |
| 6 | $(0.2876,-0.0350)$ | $(0.4377,-0.0400)$ | $(0.0458,-0.6205)$ |

S
$(11.77,0.0)^{1} \quad(9.30,0.00)^{2}(4.99,0.00)^{3}$

V

| 1 | $(0.6616,0.0000)$ | $(-0.2651,0.0000)$ | $(-0.7014,0.0000)$ |
| ---: | ---: | ---: | ---: | ---: |
| 2 | $(0.7355,0.0379)$ | $(0.3850,-0.0707)$ | $(0.5482,0.0624)$ |
| 3 | $(0.0507,-0.1317)$ | $(0.1724,0.8642)$ | $(-0.0173,-0.4509)$ |

## LSGRR

Computes the generalized inverse of a real matrix.

## Required Arguments

$\boldsymbol{A}$ - NRA by NCA matrix whose generalized inverse is to be computed. (Input)
GINVA - NCA by NRA matrix containing the generalized inverse of A. (Output)

## Optional Arguments

NRA - Number of rows in the matrix A. (Input)
Default: NRA $=\operatorname{size}(A, 1)$.

NCA - Number of columns in the matrix A. (Input)
Default: NCA $=\operatorname{size}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
TOL - Scalar containing the tolerance used to determine when a singular value (from the singular value decomposition of $A$ ) is negligible. (Input)
If TOL is positive, then a singular value $\sigma_{i}$ considered negligible if $\sigma_{i} \leq$ TOL . If TOL is negative, then a singular value $\sigma_{i}$ considered negligible if $\sigma_{i} \leq|T O L| *\|A\|_{\infty}$. In this case, $|T O L|$ generally contains an estimate of the level of the relative error in the data. Default: TOL $=1.0 \mathrm{e}-5$ for single precision and $1.0 \mathrm{~d}-10$ for double precision.

IRANK - Scalar containing an estimate of the rank of A. (Output)
LDGINV — Leading dimension of GINVA exactly as specified in the dimension statement of the calling program. (Input)
Default: LDGINV = size (GINV,1).

## FORTRAN 90 Interface

Generic: CALL LSGRR (A, GINVA [ , ...])
Specific: The specific interface names are S_LSGRR and D_LSGRR.

## FORTRAN 77 Interface

Single: CALL LSGRR (NRA, NCA, A, LDA, TOL, IRANK, GINVA, LDGINV)

Double: The double precision name is DLSGRR.

## ScaLAPACK Interface

Generic: CALL LSGRR (A0, GINVA0 [, ...])

Specific: The specific interface names are S_LSGRR and D_LSGRR.
See the ScaLAPACK Usage Notes below for a description of the arguments for distributed computing.

## Description

Let $k=$ IRANK, the rank of $A$; let $n=$ NRA, the number of rows in $A$; let $p=$ NCA, the number of columns in $A$; and let

$$
A^{\dagger}=\mathrm{GINV}
$$

be the generalized inverse of $A$.
To compute the Moore-Penrose generalized inverse, the routine LSVRR is first used to compute the singular value decomposition of $A$. A singular value decomposition of $A$ consists of an $n \times n$ orthogonal matrix $U$, a $p \times p$ orthogonal matrix $V$ and a diagonal matrix $\Sigma=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{m}\right)$, $m=\min (n, p)$, such that $U^{T} A V=[\Sigma, 0]$ if $n \leq p$ and $U^{T} A V=[\Sigma, 0]^{T}$ if $n \geq p$. Only the first $p$ columns of $U$ are computed. The rank $k$ is estimated by counting the number of nonnegligible $\sigma_{i}$.

The matrices $U$ and $V$ can be partitioned as $U=\left(U_{1}, U_{2}\right)$ and $V=\left(V_{1}, V_{2}\right)$ where both $U_{1}$ and $V_{1}$ are $k \times k$ matrices. Let $\Sigma_{1}=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{k}\right)$. The Moore-Penrose generalized inverse of $A$ is

$$
A^{\dagger}=\mathrm{V}_{1} \sum_{1}^{-1} U_{1}^{T}
$$

The underlying code is based on either LINPACK, LAPACK, or ScaLAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2GRR/DL2GRR. The reference is:

CALL L2GRR (NRA, NCA, A, LDA, TOL, IRANK, GINVA, LDGINV, WKA, WK)

The additional arguments are as follows:
WKA - Work vector of length NRA * NCA used as workspace for the matrix A. If A is not needed, WKA and A can share the same storage locations.
$\boldsymbol{W K}$ - Work vector of length LWK where LWK is equal to $N R A^{2}+\mathrm{NCA}^{2}+\min (\mathrm{NRA}+1, \mathrm{NCA})+\mathrm{NRA}+\mathrm{NCA}+\max (\mathrm{NRA}, \mathrm{NCA})-2$.
2. Informational error

Type Code
41 Convergence cannot be achieved for all the singular values and their corresponding singular vectors.

## ScaLAPACK Usage Notes

The arguments which differ from the standard version of this routine are:

A0 - MXLDA by MXCOL local matrix containing the local portions of the distributed matrix A. A contains the matrix for which the generalized inverse is to be computed. (Input)

GINVA0 - MXLDG by MXCOLG local matrix containing the local portions of the distributed matrix GINVA. GINVA contains the generalized inverse of matrix A. (Output)

All other arguments are global and are the same as described for the standard version of the routine. In the argument descriptions above, MXLDA, MXCOL, MXLDG, and MXCOLG can be obtained through a call to SCALAPACK_GETDIM (see Chapter 11, "Utilities") after a call to SCALAPACK_SETUP (see Chapter 11, "Utilities") has been made. See the ScaLAPACK Example below.

## Example

This example computes the generalized inverse of a $3 \times 2$ matrix $A$. The rank $k=$ IRANK and the inverse

$$
A^{\dagger}=\mathrm{GINV}
$$

are printed.

```
USE IMSL_LIBRARIES
```

! Declare variables
PARAMETER (NRA=3, NCA=2, LDA=NRA, LDGINV=NCA)
REAL A (LDA, NCA), GINV (LDGINV, NRA)
!
!
!
!
$!$
Set values for A
$\left.A=\begin{array}{rrr}( & 1 & 0\end{array}\right)$
DATA A/1., 1., 100., 0., 1., -50./
Compute generalized inverse
TOL $=$ AMACH (4)
TOL = 10.*TOL
CALL LSGRR (A, GINV,TOL=TOL, IRANK=IRANK)
Print results
CALL UMACH (2, NOUT)
WRITE (NOUT, *) 'IRANK = ', IRANK
CALL WRRRN ('GINV', GINV)
!
END

## Output

IRANK $=2$
GINV

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 0.1000 | 0.3000 | 0.0060 |

## ScaLAPACK Example

This example computes the generalized inverse of a $6 \times 4$ matrix $A$ as a distributed example. The rank $k=$ IRANK and the inverse

$$
A^{\dagger}=\mathrm{GINV}
$$

are printed.

```
USE MPI_SETUP INT
USE IMSL_LIBRARIES
USE SCALAPACK_SUPPORT
IMPLICIT NONE
INCLUDE `mpif.h'
                Declare variables
INTEGER IRANK, LDA, NCA, NRA, DESCA(9), DESCG(9), &
                LDGINV, MXLDG, MXCOLG, NOUT
INTEGER INFO, MXCOL, MXLDA
REAL TOL, AMACH
REAL, ALLOCATABLE :: A(:,:),GINVA(:,:)
REAL, ALLOCATABLE :: A0(:,:), GINVAO(:,:)
PARAMETER (NRA=6, NCA=4, LDA=NRA, LDGINV=NCA)
                            Set up for MPI
MP_NPROCS = MP_SETUP()
IF(MP_RANK .EQ. 0) THEN
        ALLLOCATE (A(LDA,NCA), GINVA(NCA,NRA))
                            Set values for A
        A(1,:) = (/ 1.0, 2.0, 1.0, 4.0/)
        A(2,:) = (/ 3.0, 2.0, 1.0, 3.0/)
        A(3,:) = (/ 4.0, 3.0, 1.0, 4.0/)
        A(4,:) = (/ 2.0, 1.0, 3.0, 1.0/)
        A(5,:) = (/ 1.0, 5.0, 2.0, 2.0/)
        A(6,:) = (/ 1.0, 2.0, 2.0, 3.0/)
    ENDIF
                            Set up a 1D processor grid and define
                            its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(NRA, NCA, .TRUE., .TRUE.)
                        Get the array descriptor entities MXLDA,
                        MXCOL, MXLDG, and MXCOLG
    CALL SCALAPACK_GETDIM(NRA, NCA, MP_MB, MP_NB, MXLDA, MXCOL)
    CALL SCALAPACK GETDIM(NCA, NRA, MP NB, MP MB, MXLDG, MXCOLG)
        Set u\overline{p}}\mathrm{ the ārray descriptors
    CALL DESCINIT(DESCA, NRA, NCA, MP_MB, MP_NB, 0, 0, MP_ICTXT, MXLDA, &
    INFO)
    CALL DESCINIT(DESCG, NCA, NRA, MP_NB, MP_MB, 0, 0, MP_ICTXT, MXLDG, &
    INFO)
    ALLOCATE (AO(MXLDA,MXCOL), GINVAO (MXLDG,MXCOLG))
                            Map input array to the processor grid
CALL SCALAPACK_MAP(A, DESCA, A0)
                                    Compute the generalized inverse
TOL = AMACH(4)
```

```
TOL = 10. * TOL
CALL LSGRR (A0, GINVAO, TOL=TOL, IRANK=IRANK)
                                    Unmap the results from the distributed
                                    array back to a non-distributed array.
                                    After the unmap, only Rank=0 has the full
                    array.
    CALL SCALAPACK_UNMAP(GINVAO, DESCG, GINVA)
        Print results.
                Only Rank=0 has the solution, GINVA
IF(MP_RANK .EQ. 0) THEN
    CALLL UMACH (2, NOUT)
    WRITE (NOUT, *) `IRANK = `,IRANK
    CALL WRRRN ('GINVA', GINVA)
ENDIF
CALL SCALAPACK_EXIT(MP_ICTXT)
                                    Shut down MPI
MP_NPROCS = MP_SETUP(`FINAL')
END
```

!

## Chapter 2: Eigensystem Analysis

## Routines

2.1. Eigenvalue Decomposition
2.1.1 Computes theeigenvalues of a self-adjoint matrix, A............................................................... LIN_EIG_SELF ..... 520
2.1.2 Computes the eigenvalues of an $n \times n$ matrix, $A . .$. LIN_EIG_GEN ..... 527
2.1.3 Computes the generalized eigenvalues of an $n \times n$ matrix pencil, $A v=\lambda B v$ LIN_GEIG_GEN ..... 535
2.2. Eigenvalues and (Optionally) Eigenvectors of $A x=\lambda x$
2.2.1 Real General Problem $A x=\lambda x$
All eigenvalues EVLRG ..... 543
All eigenvalues and eigenvectors ....................................... EVCRG ..... 545
Performance index. .EPIRG ..... 548
2.2.2 Complex General Problem $A x=\lambda x$
All eigenvalues EVLCG ..... 550
All eigenvalues and eigenvectors EVCCG ..... 552
Performance index .EPICG ..... 555
2.2.3 Real Symmetric Problem $A x=\lambda x$
All eigenvalues EVLSF ..... 557
All eigenvalues and eigenvectors EVCSF ..... 559
Extreme eigenvalues EVASF ..... 561
Extreme eigenvalues and their eigenvectors EVESF ..... 563
Eigenvalues in an interval. EVBSF ..... 566
Eigenvalues in an interval and their eigenvectors ..... EVFSF ..... 568
Performance index .EPISF ..... 571
2.2.4 Real Band Symmetric Matrices in Band Storage Mode All eigenvalues EVLSB ..... 573
All eigenvalues and eigenvectors EVCSB ..... 575
Extreme eigenvalues ..... EVASB ..... 578
Extreme eigenvalues and their eigenvectors EVESB ..... 581
Eigenvalues in an interval EVBSB ..... 584
Eigenvalues in an interval and their eigenvectors EVFSB ..... 586
Performance index EPISB ..... 589
2.2.5 Complex Hermitian Matrices
All eigenvalues EVLHF ..... 591
All eigenvalues and eigenvectors EVCHF ..... 593
Extreme eigenvalues ..... EVAHF ..... 596
Extreme eigenvalues and their eigenvectors ..... EVEHF ..... 599
Eigenvalues in an interval EVBHF ..... 602
Eigenvalues in an interval and their eigenvectors EVFHF ..... 604
Performance index EPIHF ..... 607
2.2.6 Real Upper Hessenberg Matrices All eigenvalues EVLRH ..... 609
All eigenvalues and eigenvectors. ..... EVCRH ..... 611
2.2.7 Complex Upper Hessenberg Matrices
All eigenvalues ..... EVLCH ..... 614
All eigenvalues and eigenvectors. ..... EVCCH ..... 616
2.3. Eigenvalues and (Optionally) Eigenvectors of $A x=\lambda B x$
2.3.1 Real General Problem $A x=\lambda B x$
All eigenvalues GVLRG ..... 618
All eigenvalues and eigenvectors GVCRG ..... 621
Performance index .GPIRG ..... 625
2.3.2 Complex General Problem $A x=\lambda B x$
All eigenvalues ..... GVLCG ..... 627
All eigenvalues and eigenvectors GVCCG ..... 629
Performance index GPICG ..... 632
2.3.3 Real Symmetric Problem $A x=\lambda B x$
All eigenvalues GVLSP ..... 634
All eigenvalues and eigenvectors GVCSP ..... 636
Performance index GPISP ..... 639

## Usage Notes

This chapter includes routines for linear eigensystem analysis. Many of these are for matrices with special properties. Some routines compute just a portion of the eigensystem. Use of the appropriate routine can substantially reduce computing time and storage requirements compared to computing a full eigensystem for a general complex matrix.

An ordinary linear eigensystem problem is represented by the equation $A x=\lambda x$ where $A$ denotes an
$n \times n$ matrix. The value $\lambda$ is an eigenvalue and $x \neq 0$ is the corresponding eigenvector. The eigenvector is determined up to a scalar factor. In all routines, we have chosen this factor so that $x$ has Euclidean length with value one, and the component of $x$ of smallest index and largest magnitude is positive. In case $x$ is a complex vector, this largest component is real and positive.

Similar comments hold for the use of the remaining Level 1 routines in the following tables in those cases where the second character of the Level 2 routine name is no longer the character " 2 ".

A generalized linear eigensystem problem is represented by $A x=\lambda B x$ where $A$ and $B$ are $n \times n$ matrices. The value $\lambda$ is an eigenvalue, and $x$ is the corresponding eigenvector. The eigenvectors are normalized in the same manner as for the ordinary eigensystem problem. The linear eigensystem routines have names that begin with the letter " $E$ ". The generalized linear eigensystem routines have names that begin with the letter "G". This prefix is followed by a twoletter code for the type of analysis that is performed. That is followed by another two-letter suffix for the form of the coefficient matrix. The following tables summarize the names of the eigensystem routines.

| Symmetric and Hermitian Eigensystems |  |  |  |
| :--- | :--- | :--- | :--- |
|  | Symmetric <br> Full | Symmetric <br> Band | Hermitian <br> Full |
| All eigenvalues | EVLSF | EVLSB | EVLHF |
| All eigenvalues <br> and eigenvectors | EVCSF | EVCSB | EVCHF |
| Extreme eigenvalues | EVASF | EVASB | EVAHF |
| Extreme eigenvalues <br> and eigenvectors | EVESF | EVESB | EVEHF |
| Eigenvalues in <br> an interval | EVBSF | EVBSB | EVBHF |
| Eigenvalues and <br> eigevectors in an interval | EVFSF | EVFSB | EVFHF |
| Performance index | EPISF | EPISB | EPIHF |


| General Eigensystems |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  | Real <br> General | Complex <br> General | Real <br> Hessenberg | Complex <br> Hessenberg |
| All eigenvalues | EVLRG | EVLCG | EVLRH | EVLCH |
| All eigenvalues <br> and eigenvectors | EVCRG | EVCCG | EVCRH | EVCCH |
| Performance <br> index | EPIRG | EPICG | EPIRG | EPICG |


| Generalized Eigensystems Ax= $\mathbf{n} \boldsymbol{B} \boldsymbol{x}$ |  |  |  |
| :--- | :--- | :--- | :--- |
|  | Real <br> General | Complex <br> General | A Symmetric <br> B Positive <br> Definite |
| All eigenvalues | GVLRG | GVLCG | GVLSP |
| All eigenvalues and <br> eigenvectors | GVCRG | GVCCG | GVCSP |
| Performance index | GPIRG | GPICG | GPISP |

## Error Analysis and Accuracy

The remarks in this section are for the ordinary eigenvalue problem. Except in special cases, routines will not return the exact eigenvalue-eigenvector pair for the ordinary eigenvalue problem $A x=\lambda x$. The computed pair

$$
\tilde{x}, \tilde{\lambda}
$$

is an exact eigenvector-eigenvalue pair for a "nearby" matrix $A+E$. Information about $E$ is known only in terms of bounds of the form $\|E\|_{2} \leq f(n)\|A\|_{2} \varepsilon$. The value of $f(n)$ depends on the algorithm but is typically a small fractional power of $n$. The parameter $\varepsilon$ is the machine precision. By a theorem due to Bauer and Fike (see Golub and Van Loan [1989, page 342],

$$
\min |\tilde{\lambda}-\lambda| \leq \kappa(X)\|E\|_{2} \quad \text { for all } \lambda \text { in } \sigma(\mathrm{A})
$$

where $\sigma(A)$ is the set of all eigenvalues of $A$ (called the spectrum of $A$ ), $X$ is the matrix of eigenvectors, $\|\cdot\|_{2}$ is the 2-norm, and $\kappa(X)$ is the condition number of $X$ defined as $\kappa(X)=\|X\|_{2}\left\|X^{-1}\right\|_{2}$. If $A$ is a real symmetric or complex Hermitian matrix, then its eigenvector matrix $X$ is respectively orthogonal or unitary. For these matrices, $\kappa(X)=1$.
The eigenvalues

$$
\tilde{\lambda}_{j}
$$

and eigenvectors

$$
\tilde{x}_{j}
$$

computed by EVC** can be checked by computing their performance index $\tau$ using EPI**. The performance index is defined by Smith et al. (1976, pages 124-126) to be

$$
\tau=\max _{1 \leq j \leq n} \frac{\left\|A \tilde{x}_{j}-\tilde{\lambda}_{j} \tilde{x}_{j}\right\|_{1}}{10 n \varepsilon\|A\|_{1}\left\|\tilde{x}_{j}\right\|_{1}}
$$

No significance should be attached to the factor of 10 used in the denominator. For a real vector $x$, the symbol $\|x\|_{1}$ represents the usual 1-norm of $x$. For a complex vector $x$, the symbol $\|x\|_{1}$ is defined by

$$
\|x\|_{1}=\sum_{k=1}^{N}\left(\left|\mathfrak{\Re} x_{k}\right|+\left|\mathfrak{J} x_{k}\right|\right)
$$

The performance index $\tau$ is related to the error analysis because

$$
\left\|E \tilde{x}_{j}\right\|_{2} \dot{=}\left\|A \tilde{x}_{j}-\tilde{\lambda}_{j} \tilde{x}_{j}\right\|_{2}
$$

where $E$ is the "nearby" matrix discussed above.
While the exact value of $\tau$ is machine and precision dependent, the performance of an eigensystem analysis routine is defined as excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. This is an arbitrary definition, but large values of $\tau$ can serve as a warning that there is a blunder in the
calculation. There are also similar routines GPI** to compute the performance index for generalized eigenvalue problems.

If the condition number $\kappa(X)$ of the eigenvector matrix $X$ is large, there can be large errors in the eigenvalues even if $\tau$ is small. In particular, it is often difficult to recognize near multiple eigenvalues or unstable mathematical problems from numerical results. This facet of the eigenvalue problem is difficult to understand: A user often asks for the accuracy of an individual eigenvalue. This can be answered approximately by computing the condition number of an individual eigenvalue. See Golub and Van Loan (1989, pages 344-345). For matrices $A$ such that the computed array of normalized eigenvectors $X$ is invertible, the condition number of $\lambda_{j}$ is $\kappa_{j} \equiv$ the Euclidean length of row $j$ of the inverse matrix $X^{-1}$. Users can choose to compute this matrix with routine LINCG, see Chapter 1, Linear Systems. An approximate bound for the accuracy of a computed eigenvalue is then given by $\kappa_{j} \varepsilon\|A\|$. To compute an approximate bound for the relative accuracy of an eigenvalue, divide this bound by $\left|\lambda_{j}\right|$.

## Reformulating Generalized Eigenvalue Problems

The generalized eigenvalue problem $A x=\lambda B x$ is often difficult for users to analyze because it is frequently ill-conditioned. There are occasionally changes of variables that can be performed on the given problem to ease this ill-conditioning. Suppose that $B$ is singular but $A$ is nonsingular. Define the reciprocal $\mu=\lambda^{-1}$. Then, the roles of $A$ and $B$ are interchanged so that the reformulated problem
$B x=\mu A x$ is solved. Those generalized eigenvalues $\mu_{j}=0$ correspond to eigenvalues $\lambda_{j}=\infty$. The remaining

$$
\lambda_{j}=\mu_{j}^{-1}
$$

The generalized eigenvectors for $\lambda_{j}$ correspond to those for $\mu_{j}$. Other reformulations can be made: If $B$ is nonsingular, the user can solve the ordinary eigenvalue problem $C x \equiv B^{-1} A x=\lambda x$. This is not recommended as a computational algorithm for two reasons. First, it is generally less efficient than solving the generalized problem directly. Second, the matrix $C$ will be subject to perturbations due to ill-conditioning and rounding errors when computing $B^{-1} A$. Computing the condition numbers of the eigenvalues for $C$ may, however, be helpful for analyzing the accuracy of results for the generalized problem.

There is another method that users can consider to reduce the generalized problem to an alternate ordinary problem. This technique is based on first computing a matrix decomposition $B=P Q$, where both $P$ and $Q$ are matrices that are "simple" to invert. Then, the given generalized problem is equivalent to the ordinary eigenvalue problem $F y=\lambda y$. The matrix $F \equiv P^{-1} A Q^{-1}$. The unnormalized eigenvectors of the generalized problem are given by $x=Q^{-1} y$. An example of this reformulation is used in the case where $A$ and $B$ are real and symmetric with $B$ positive definite. The IMSL routines GVLSP and GVCSP use $P=R^{T}$ and $Q=R$ where $R$ is an upper triangular matrix obtained from a Cholesky decomposition, $B=R^{T} R$. The matrix $F=R^{-\mathrm{T}} A R^{-1}$ is symmetric and real. Computation of the eigenvalue-eigenvector expansion for $F$ is based on routine EVCSF.

## LIN_EIG_SELF

Computes the eigenvalues of a self-adjoint (i.e. real symmetric or complex Hermitian) matrix, $A$. Optionally, the eigenvectors can be computed. This gives the decomposition $A=V D V^{T}$, where $V$ is an $n \times n$ orthogonal matrix and $D$ is a real diagonal matrix.

## Required Arguments

A- Array of size $n \times n$ containing the matrix. (Input [/Output])
$\boldsymbol{D}$ - Array of size $n$ containing the eigenvalues. The values are in order of decreasing absolute value. (Output)

## Optional Arguments

NROWS $=\mathrm{n}$ (Input)
Uses array $A(1: n, 1: n)$ for the input matrix.
Default: $\mathrm{n}=\operatorname{size}(\mathrm{A}, 1)$
$\mathrm{v}=\mathrm{v}(:,:$ ) (Output)
Array of the same type and kind as $\mathrm{A}(1: \mathrm{n}, 1: \mathrm{n})$. It contains the $n \times n$ orthogonal matrix $V$.
iopt = iopt (: ) (Input)
Derived type array with the same precision as the input matrix; used for passing optional data to the routine. The options are as follows:

| Packaged Options for LIN_EIG_SELF |  |  |
| :--- | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| s_, d_, c_, z_ | Lin_eig_self_set_small | 1 |
| s_, d_, c_, z_ | Lin_eig_self_overwrite_input | 2 |
| s_, d_, c_, z_ | Lin_eig_self_scan_for_NaN | 3 |
| s_, d_, c_, z_ | Lin_eig_self_use_QR | 4 |
| s_, d_, c_, z_ | Lin_eig_self_skip_Orth | 5 |
| s_, d_, c_, z_ | Lin_eig_self_use_Gauss_elim | 6 |
| s_, d_, c_, z_ | Lin_eig_self_set_perf_ratio | 7 |

iopt(IO) = ?_options(?_lin_eig_self_set_small, Small)
If a denominator term is smaller in magnitude than the value Small, it is replaced by Small.
Default: the smallest number that can be reciprocated safely

```
iopt(IO) = ?_options(?_lin_eig_self_overwrite_input, ?_dummy)
```

Do not save the input array $A(:,:)$.
iopt(IO) = ?_options(?_lin_eig_self_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(a(i,j)) == .true.

See the isnaN() function, Chapter 10.
Default: The array is not scanned for NaNs.
iopt(IO) $=$ ?_options(?_lin_eig_use_QR, ?_dummy)
Uses a rational $Q R$ algorithm to compute eigenvalues. Accumulate the eigenvectors using this algorithm.
Default: the eigenvectors computed using inverse iteration
iopt(IO) = ?_options(?_lin_eig_skip_Orth, ?_dummy)
If the eigenvalues are computed using inverse iteration, skips the final orthogonalization of the vectors. This will result in a more efficient computation but the eigenvectors, while a complete set, may be far from orthogonal.
Default: the eigenvectors are normally orthogonalized if obtained using inverse iteration.
iopt(IO) = ?_options(?_lin_eig_use_Gauss_elim, ?_dummy)
If the eigenvalues are computed using inverse iteration, uses standard elimination with partial pivoting to solve the inverse iteration problems.
Default: the eigenvectors computed using cyclic reduction
iopt(IO) $=$ ? _options(?_lin_eig_self_set_perf_ratio, perf_ratio)
Uses residuals for approximate normalized eigenvectors if they have a performance index no larger than perf_ratio. Otherwise an alternate approach is taken and the eigenvectors are computed again: Standard elimination is used instead of cyclic reduction, or the standard $Q R$ algorithm is used as a backup procedure to inverse iteration. Larger values of perf_ratio are less likely to cause these exceptions.
Default: perf_ratio $=4$

## FORTRAN 90 Interface

Generic: CALL LIN_EIG_SELF (A, D [, ...])
Specific: The specific interface names are S_LIN_EIG_SELF, D_LIN_EIG_SELF, C_LIN_EIG_SELF, and Z_LIN_EIG_SELF.

## Description

Routine LIN_EIG_SELF is an implementation of the $Q R$ algorithm for self-adjoint matrices. An orthogonal similarity reduction of the input matrix to self-adjoint tridiagonal form is performed. Then, the eigenvalue-eigenvector decomposition of a real tridiagonal matrix is calculated. The expansion of the matrix as $A V=V D$ results from a product of these matrix factors. See Golub and Van Loan (1989, Chapter 8) for details.

## Fatal, Terminal, and Warning Error Messages

See the messages.gls file for error messages for LIN_EIG_SELF. These error messages are numbered 81-90; 101-110; 121-129; 141-149.

## Example 1: Computing Eigenvalues

The eigenvalues of a self-adjoint matrix are computed. The matrix $A=C+C^{T}$ is used, where $C$ is random. The magnitudes of eigenvalues of $A$ agree with the singular values of $A$. Also, see operator_ex 25 , supplied with the product examples.

```
        use lin eig self int
        use lin_sol_svd_int
        use rand_gen_int
    implicit none
! This is Example 1 for LIN_EIG_SELF.
    integer, parameter :: n=64
    real(kind(1e0)), parameter :: one=1e0
    real(kind(le0)) :: A(n,n), b(n,0), D(n), S(n), x(n,0), y(n*n)
! Generate a random matrix and from it
! a self-adjoint matrix.
    call rand_gen(y)
    A = reshape(y,(/n,n/))
    A = A + transpose(A)
! Compute the eigenvalues of the matrix.
    call lin_eig_self(A, D)
! For comparison, compute the singular values.
    call lin_sol_svd(A, b, x, nrhs=0, s=S)
! Check the results: Magnitude of eigenvalues should equal
! the singular values.
    if (sum(abs(abs(D) - S)) <= &
        sqrt(epsilon(one))*S(1)) then
        write (*,*) 'Example 1 for LIN_EIG_SELF is correct.'
    end if
    end
```


## Output

Example 1 for LIN EIG_SELF is correct.

## Additional Examples

## Example 2: Eigenvalue-Eigenvector Expansion of a Square Matrix

A self-adjoint matrix is generated and the eigenvalues and eigenvectors are computed. Thus, $A=V D V^{T}$, where $V$ is orthogonal and $D$ is a real diagonal matrix. The matrix $V$ is obtained using an optional argument. Also, see operator_ex26, Chapter 10.

```
    use lin_eig_self_int
    use rand_gen__int
    implicit none
! This is Example 2 for LIN_EIG_SELF.
    integer, parameter :: n=8
    real(kind(1e0)), parameter : : one=1e0
    real(kind(1e0)) :: a(n,n), d(n), v_s(n,n), y(n*n)
! Generate a random self-adjoint matrix.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
    a = a + transpose(a)
! Compute the eigenvalues and eigenvectors.
    call lin_eig_self(a, d, v=v_s)
! Check the results for small residuals.
    if (sum(abs(matmul(a,v_s)-v_s*spread(d,1,n)))/d(1) <= &
            sqrt(epsilon(one))) then
        write (*,*) 'Example 2 for LIN_EIG_SELF is correct.'
    end if
    end
```


## Output

Example 2 for LIN EIG SELF is correct.

## Example 3: Computing a few Eigenvectors with Inverse Iteration

A self-adjoint $n \times n$ matrix is generated and the eigenvalues, $\left\{d_{i}\right\}$, are computed. The eigenvectors associated with the first $k$ of these are computed using the self-adjoint solver, lin_sol_self, and inverse iteration. With random right-hand sides, these systems are as follows:

$$
\left(A-d_{i} I\right) v_{i}=b_{i}
$$

The solutions are then orthogonalized as in Hanson et al. (1991) to comprise a partial decomposition $A V=V D$ where $V$ is an $n \times k$ matrix resulting from the orthogonalized $\left\{v_{i}\right\}$ and $D$ is the $k \times k$ diagonal matrix of the distinguished eigenvalues. It is necessary to suppress the error message when the matrix is singular. Since these singularities are desirable, it is appropriate to ignore the exceptions and not print the message text. Also, see operator_ex 27 , supplied with the product examples.

```
    use lin_eig_self_int
    use lin_sol_self_int
    use ran\overline{d}ge\overline{n}\mathrm{ int}
    use error_option_packet
    implicit none
! This is Example 3 for LIN_EIG_SELF.
    integer i, j
    integer, parameter :: n=64, k=8
    real(kind(1d0)), parameter :: one=1d0, zero=0d0
    real(kind(1d0)) big, err
    real(kind(1d0)) :: a(n,n), b(n,1), d(n), res(n,k), temp(n,n), &
            v(n,k), y(n*n)
    type(d_options) :: iopti(2)=d_options(0,zero)
! Generate a random self-adjoint matrix.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
    a = a + transpose(a)
! Compute just the eigenvalues.
    call lin_eig_self(a, d)
    do i=1, k
! Define a temporary array to hold the matrices A - eigenvalue*I.
        temp = a
        do j=1, n
            temp(j,j) = temp(j,j) - d(i)
        end do
! Use packaged option to reset the value of a small diagonal.
        iopti(1) = d_options(d_lin_sol_self_set_small,&
                        epsilon(one)*abs(d(\overline{i})))
! Use packaged option to skip singularity messages.
        iopti(2) = d_options(d_lin_sol_self_no_sing_mess,&
                zero)
        call rand gen(b(1:n,1))
        call lin_sol_self(temp, b, v(1:,i:i),&
            iopt=iopti)
    end do
! Orthogonalize the eigenvectors.
    do i=1, k
        big = maxval(abs(v(1:,i)))
        v(1:,i) = v(1:,i)/big
        v(1:,i) = v(1:,i)/sqrt(sum(v(1:,i)**2))
        if (i == k) cycle
        v(1:,i+1:k) = v(1:,i+1:k) + &
            spread(-matmul(v(1:,i),v(1:,i+1:k)),1,n)* &
            spread(v(1:,i),2,k-i)
    end do
```

```
    do i=k-1, 1, -1
    v(1:,i+1:k) = v(1:,i+1:k) + &
            spread(-matmul(v(1:,i),v(1:,i+1:k)),1,n)* &
            spread(v(1:,i),2,k-i)
    end do
! Check the results for both orthogonality of vectors and small
! residuals.
    res(1:k,1:k) = matmul(transpose(v),v)
    do i=1,k
        res(i,i)=res(i,i)-one
    end do
    err = sum(abs(res))/k**2
    res = matmul(a,v) - v*spread (d(1:k),1,n)
    if (err <= sqrt(epsilon(one))) then
        if (sum(abs(res))/abs(d(1)) <= sqrt(epsilon(one))) then
                write (*,*) 'Example 3 for LIN_EIG_SELF is correct.'
        end if
    end if
    end
```


## Output

Example 3 for LIN EIG_SELF is correct.

## Example 4: Analysis and Reduction of a Generalized Eigensystem

A generalized eigenvalue problem is $A x=\lambda B x$, where $A$ and $B$ are $n \times n$ self-adjoint matrices. The matrix $B$ is positive definite. This problem is reduced to an ordinary self-adjoint eigenvalue
problem $C y=\lambda y$ by changing the variables of the generalized problem to an equivalent form. The eigenvalue-eigenvector decomposition $B=V S V^{T}$ is first computed, labeling an eigenvalue too small if it is less than epsilon (1.d0). The ordinary self-adjoint eigenvalue problem is $C y=\lambda y$ provided that the rank of $B$, based on this definition of Small, has the value $n$. In that case,

$$
C=D V^{T} A V D
$$

where

$$
D=S^{-1 / 2}
$$

The relationship between $x$ and $y$ is summarized as $X=V D Y$, computed after the ordinary eigenvalue problem is solved for the eigenvectors $Y$ of $C$. The matrix $X$ is normalized so that each column has Euclidean length of value one. This solution method is nonstandard for any but the most
ill-conditioned matrices $B$. The standard approach is to compute an ordinary self-adjoint problem following computation of the Cholesky decomposition

$$
B=R^{T} R
$$

where $R$ is upper triangular. The computation of $C$ can also be completed efficiently by exploiting its self-adjoint property. See Golub and Van Loan (1989, Chapter 8) for more information. Also, see operator_ex28, Chapter 10.

```
use lin_eig_self_int
```

```
        use rand_gen_int
        implicit none
! This is Example 4 for LIN_EIG_SELF.
    integer i
    integer, parameter :: n=64
    real(kind(1e0)), parameter :: one=1d0
    real(kind(le0)) b sum
    real(kind(1e0)), dimension(n,n) :: A, B, C, D(n), lambda(n), &
        S(n), vb_d, X, ytemp(n*n), res
! Generate random self-adjoint matrices.
    call rand_gen(ytemp)
    A = reshape(ytemp,(/n,n/))
    A = A + transpose(A)
    call rand_gen(ytemp)
    B = reshape(ytemp,(/n,n/))
    B = B + transpose(B)
    b_sum = sqrt(sum(abs(B**2))/n)
! Add a scalar matrix so B is positive definite.
    do i=1, n
        B(i,i) = B(i,i) + b_sum
    end do
! Get the eigenvalues and eigenvectors for B.
    call lin_eig_self(B, S, v=vb_d)
! For full rank problems, convert to an ordinary self-adjoint
! problem. (All of these examples are full rank.)
    if (S(n) > epsilon(one)) then
        D = one/sqrt(S)
        C = spread(D,2,n)*matmul(transpose(vb_d), &
            matmul(A,vb_d))*spread(D,1,n)
! Get the eigenvalues and eigenvectors for C.
            call lin_eig_self(C, lambda, v=X)
! Compute the generalized eigenvectors.
    X = matmul(vb_d,spread (D, 2, n)*X)
! Normalize the eigenvectors for the generalized problem.
    X = X * spread(one/sqrt(sum(X**2,dim=2)),1,n)
    res = matmul(A,X) - &
            matmul (B,X)*spread(lambda,1,n)
! Check the results.
```

```
if (sum(abs(res))/(sum(abs(A)) +sum(abs(B))) <= &
        sqrt(epsilon(one))) then
        write (*,*) 'Example 4 for LIN EIG SELF is correct.'
        end if
```

end if
end

## Output

Example 4 for LIN_EIG_SELF is correct.

## LIN_EIG_GEN

Computes the eigenvalues of an $n \times n$ matrix, $A$. Optionally, the eigenvectors of $A$ or $A^{T}$ are computed. Using the eigenvectors of $A$ gives the decomposition $A V=V E$, where $V$ is an $n \times n$ complex matrix of eigenvectors, and $E$ is the complex diagonal matrix of eigenvalues. Other options include the reduction of $A$ to upper triangular or Schur form, reduction to block upper triangular form with $2 \times 2$ or unit sized diagonal block matrices, and reduction to upper Hessenberg form.

## Required Arguments

A- Array of size $n \times n$ containing the matrix. (Input [/Output])
$\boldsymbol{E}$ - Array of size $n$ containing the eigenvalues. These complex values are in order of decreasing absolute value. The signs of imaginary parts of the eigenvalues are in no predictable order. (Output)

## Optional Arguments

NROWS $=\mathrm{n}$ (Input)
Uses array $A(1: n, 1: n)$ for the input matrix.
Default: $\mathrm{n}=\operatorname{SIZE}(\mathrm{A}, 1)$
$\mathrm{v}=\mathrm{V}(:,:$ ) (Output)
Returns the complex array of eigenvectors for the matrix $A$.
v_adj $=\mathrm{U}(:,:$ ) (Output)
Returns the complex array of eigenvectors for the matrix $A^{T}$. Thus the residuals

$$
S=A^{T} U-U \bar{E}
$$

are small.
tri $=T(:,:)$ (Output)
Returns the complex upper-triangular matrix $T$ associated with the reduction of the matrix $A$ to Schur form. Optionally a unitary matrix $W$ is returned in array $\mathrm{V}(:,:$ ) such that the residuals $Z=A W-W T$ are small.
iopt = iopt(:) (Input)
Derived type array with the same precision as the input matrix. Used for passing optional data to the routine. The options are as follows:

| Packaged Options for LIN_EIG_GEN |  |  |
| :---: | :---: | :---: |
| Option Prefix = ? | Option Name | Option Value |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_set_small | 1 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_overwrite_input | 2 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_scan_for_NaN | 3 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_no_balance | 4 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_set_iterations | 5 |
| $S_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_in_Hess_form | 6 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_out_Hess_form | 7 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_out_block_form | 8 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_out_tri_form | 9 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_continue_with_V | 10 |
| $s_{-}, d_{-}, c_{-}, z_{-}$ | lin_eig_gen_no_sorting | 11 |

iopt (IO) = ?_options(?_lin_eig_gen_set_small, Small)
This is the tolerance used to declare off-diagonal values effectively zero compared with the size of the numbers involved in the computation of a shift.
Default: Small = epsilon(), the relative accuracy of arithmetic
iopt(IO) = ?_options(?_lin_eig_gen_overwrite_input, ?_dummy)
Does not save the input array A(: : ).
Default: The array is saved.
iopt (IO) = ? _options(?_lin_eig_gen_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(a(i,j)) == .true.
See the isNaN () function, Chapter 10.
Default: The array is not scanned for NaNs.
iopt (IO) = ?_options(?_lin_eig_no_balance, ?_dummy)
The input matrix is not preprocessed searching for isolated eigenvalues followed by rescaling. See Golub and Van Loan (1989, Chapter 7) for references. With some optional uses of the routine, this option flag is required.
Default: The matrix is first balanced.
iopt(IO) = ?_options(?_lin_eig_gen_set_iterations, ?_dummy)
Resets the maximum number of iterations permitted to isolate each diagonal block matrix.
Default: The maximum number of iterations is 52 .
iopt(IO) = ?_options(?_lin_eig_gen_in_Hess_form, ?_dummy)
The input matrix is in upper Hessenberg form. This flag is used to avoid the initial reduction phase which may not be needed for some problem classes.
Default: The matrix is first reduced to Hessenberg form.
iopt(IO) = ?_options(?_lin_eig_gen_out_Hess_form, ?_dummy)
The output matrix is transformed to upper Hessenberg form, $H_{1}$. If the optional argument " $\mathrm{v}=\mathrm{V}(:,:)$ " is passed by the calling program unit, then the array $\mathrm{V}(:,:)$ contains an orthogonal matrix $Q_{1}$ such that

$$
A Q_{1}-Q_{1} H_{1} \cong 0
$$

Requires the simultaneous use of option ?_lin_eig_no_balance.
Default: The matrix is reduced to diagonal form.
iopt(IO) = ?_options(?_lin_eig_gen_out_block_form, ?_dummy)
The output matrix is transformed to upper Hessenberg form, $\bar{H}_{2}$, which is block upper triangular. The dimensions of the blocks are either $2 \times 2$ or unit sized. Nonzero subdiagonal values of $\mathrm{H}_{2}$ determine the size of the blocks. If the optional argument " $\mathrm{v}=\mathrm{V}(:,:)$ " is passed by the calling program unit, then the array $\mathrm{V}(:,:)$ contains an orthogonal matrix $Q_{2}$ such that

$$
A Q_{2}-Q_{2} H_{2} \cong 0
$$

Requires the simultaneous use of option ?_lin_eig_no_balance.
Default: The matrix is reduced to diagonal form.

```
iopt(IO) = ?_options(?_lin_eig_gen_out_tri_form, ?_dummy)
```

The output matrix is transformed to upper-triangular form, $T$. If the optional argument $" \mathrm{v}=\mathrm{V}(:,:)$ " is passed by the calling program unit, then the array $\mathrm{V}(:,:)$ contains a unitary matrix $W$ such that
$A W-W T \cong 0$. The upper triangular matrix $T$ is returned in the optional argument "tri=T (:, : )". The eigenvalues of $A$ are the diagonal entries of the matrix $T$. They are in no particular order. The output array $\mathrm{E}(:)$ is blocked with NaNs using this option. This option requires the simultaneous use of option ?_lin_eig_no_balance. Default: The matrix is reduced to diagonal form.

```
iopt(IO) = ?_options(?_lin_eig_gen_continue_with_V, ?_dummy)
```

As a convenience or for maintaining efficiency, the calling program unit sets the optional argument " $\mathrm{v}=\mathrm{V}(:,:)$ " to a matrix that has transformed a problem to the similar matrix, $\dot{A}$. The contents of $\mathrm{v}(:,:)$ are updated by the transformations used in the algorithm. Requires the simultaneous use of option?_lin_eig_no_balance. Default: The array $\mathrm{V}(:,:)$ is initialized to the identity matrix.
iopt(IO) = ?_options(?_lin_eig_gen_no_sorting, ?_dummy)
Does not sort the eigenvalues as they are isolated by solving the $2 \times 2$ or unit sized blocks. This will have the effect of guaranteeing that complex conjugate pairs of
eigenvalues are adjacent in the array $\mathrm{E}(:)$.
Default: The entries of $\mathrm{E}(:)$ are sorted so they are non-increasing in absolute value.

## FORTRAN 90 Interface

Generic: CALL LIN_EIG_GEN (A, E [, ...])
Specific: The specific interface names are S_LIN_EIG_GEN, D_LIN_EIG_GEN, C_LIN_EIG_GEN, and Z_LIN_EIG_GEN.

## Description

The input matrix $A$ is first balanced. The resulting similar matrix is transformed to upper Hessenberg form using orthogonal transformations. The double-shifted $Q R$ algorithm transforms the Hessenberg matrix so that $2 \times 2$ or unit sized blocks remain along the main diagonal. Any off-diagonal that is classified as "small" in order to achieve this block form is set to the value zero. Next the block upper triangular matrix is transformed to upper triangular form with unitary rotations. The eigenvectors of the upper triangular matrix are computed using back substitution. Care is taken to avoid overflows during this process. At the end, eigenvectors are normalized to have Euclidean length one, with the largest component real and positive. This algorithm follows that given in Golub and Van Loan, (1989, Chapter 7), with some novel organizational details for additional options, efficiency and robustness.

## Fatal, Terminal, and Warning Error Messages

See the messages.gls file for error messages for LIN_EIG_GEN. These error messages are numbered 841-858; 861-878; 881-898; 901-918.

## Example 1: Computing Eigenvalues

The eigenvalues of a random real matrix are computed. These values define a complex diagonal matrix $E$. Their correctness is checked by obtaining the eigenvector matrix $V$ and verifying that the residuals $R=A V-V E$ are small. Also, see operator_ex29, supplied with the product examples.

```
    use lin eig_gen int
    use rand_gen_int
    implicit none
! This is Example 1 for LIN_EIG_GEN.
    integer, parameter :: n=32
    real(kind(1d0)), parameter :: one=1d0
    real(kind(ldO)) A(n,n), y(n*n), err
    complex(kind(1d0)) E(n), V(n,n), E_T(n)
    type(d_error) :: d_epack(16) = d_e\overline{rror}(0,0d0)
! Generate a random matrix.
    call rand_gen(y)
    A = reshape(y,(/n,n/))
```

```
! Compute only the eigenvalues.
    call lin_eig_gen(A, E)
! Compute the decomposition, A*V = V*values,
! obtaining eigenvectors.
    call lin_eig_gen(A, E_T, v=V)
! Use values from the first decomposition, vectors from the
! second decomposition, and check for small residuals.
    err = sum(abs(matmul(A,V) - V*spread(E,DIM=1,NCOPIES=n))) &
            / sum(abs(E))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_EIG_GEN is correct.'
    end if
    end
```


## Output

```
Example 1 for LIN_EIG_GEN is correct.
```


## Additional Examples

## Example 2: Complex Polynomial Equation Roots

The roots of a complex polynomial equation,

$$
f(z) \equiv \sum_{k=1}^{n} b_{k} z^{n-k}+z^{n}=0
$$

are required. This algebraic equation is formulated as a matrix eigenvalue problem. The equivalent matrix eigenvalue problem is solved using the upper Hessenberg matrix which has the value zero except in row number 1 and along the first subdiagonal. The entries in the first row are given by $a_{1}, \mathrm{j}=-b_{j}, i=1, \ldots, n$, while those on the first subdiagonal have the value one. This is a companion matrix for the polynomial. The results are checked by testing for small values of $\left|f\left(e_{i}\right)\right|, i=1, \ldots, n$, at the eigenvalues of the matrix, which are the roots of $f(z)$. Also, see operator_ex 30 , supplied with the product examples.

```
    use lin_eig_gen_int
    use rand_gen_int
    implicit none
! This is Example 2 for LIN_EIG_GEN.
integer i
integer, parameter :: n=12
real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
real(kind(1d0)) err, t(2*n)
type(d options) :: iopti(1)=d options(0,zero)
comple\overline{x}(kind(ld0)) a(n,n), b(\overline{n}), e(n), f(n), fg(n)
call rand_gen(t)
b = cmplx(t(1:n),t(n+1:), kind(one))
```

```
! Define the companion matrix with polynomial coefficients
! in the first row.
    a = zero
    do i=2, n
        a(i,i-1) = one
    end do
    a(1,1:n) = -b
! Note that the input companion matrix is upper Hessenberg.
    iopti(1) = d_options(z_lin_eig_gen_in_Hess_form,zero)
! Compute complex eigenvalues of the companion matrix.
    call lin_eig_gen(a, e, iopt=iopti)
    f=one; fg=one
! Use Horner's method for evaluation of the complex polynomial
! and size gauge at all roots.
    do i=1, n
        f = f*e + b(i)
        fg = fg*abs(e) + abs(b(i))
    end do
! Check for small errors at all roots.
    err = sum(abs(f/fg))/n
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 2 for LIN_EIG_GEN is correct.'
    end if
    end
```


## Output

Example 2 for LIN_EIG_GEN is correct.

## Example 3: Solving Parametric Linear Systems with a Scalar Change

The efficient solution of a family of linear algebraic equations is required. These systems are $(A+h I) x=b$. Here $A$ is an $n \times n$ real matrix, $I$ is the identity matrix, and $b$ is the right-hand side matrix. The scalar $h$ is such that the coefficient matrix is nonsingular. The method is based on the Schur form for matrix $A: A W=W T$, where $W$ is unitary and $T$ is upper triangular. This provides an efficient solution method for several values of $h$, once the Schur form is computed. The solution steps solve, for $y$, the upper triangular linear system

$$
(T+h I) y=\bar{W}^{T} b
$$

Then, $x=x(h)=W y$. This is an efficient and accurate method for such parametric systems provided the expense of computing the Schur form has a pay-off in later efficiency. Using the Schur
form in this way, it is not required to compute an $L U$ factorization of $A+h I$ with each new value of $h$. Note that even if the data $A, h$, and $b$ are real, subexpressions for the solution may involve complex intermediate values, with $x(h)$ finally a real quantity. Also, see operator_ex31, supplied with the product examples.

```
        use lin_eig_gen_int
        use lin_sol_gen_int
        use ran\overline{d}_ge\overline{n}_in\overline{t}
    implicit none
! This is Example 3 for LIN_EIG_GEN.
    integer i
    integer, parameter :: n=32, k=2
    real(kind(1e0)), parameter : : one=1.0e0, zero=0.0e0
    real(kind(le0)) a (n,n), b (n,k), x(n,k), temp(n*max(n,k)), h, err
    type(s_options) :: iopti(2)
    comple\overline{x}(kind(le0)) w(n,n), t(n,n), e(n), z(n,k)
    call rand_gen(temp)
    a = reshape(temp,(/n,n/))
    call rand_gen(temp)
    b = reshape(temp,(/n,k/))
    iopti(1) = s_options(s_lin_eig_gen_out_tri_form,zero)
    iopti(2) = s_options(s_lin_eig_gen_no_balañce,zero)
! Compute the Schur decomposition of the matrix.
    call lin_eig_gen(a, e, v=w, tri=t, &
        iopt=iopti)
! Choose a value so that A+h*I is non-singular.
    h = one
! Solve for (A+h*I)x=b using the Schur decomposition.
    z = matmul(conjg(transpose(w)),b)
! Solve intermediate upper-triangular system with implicit
! additive diagonal, h*I. This is the only dependence on
! h in the solution process.
    do i=n,1,-1
        z(i,1:k) = z(i,1:k)/(t(i,i)+h)
        z(1:i-1,1:k) = z(1:i-1,1:k) + &
            spread(-t(1:i-1,i),dim=2,ncopies=k)* &
                spread(z(i,1:k),dim=1,ncopies=i-1)
    end do
! Compute the solution. It should be the same as x, but will not be
! exact due to rounding errors. (The quantity real(z,kind(one)) is
! the real-valued answer when the Schur decomposition method is used.)
```

```
        z = matmul(w,z)
! Compute the solution by solving for x directly.
        do i=1, n
        a(i,i) = a(i,i) + h
        end do
        call lin_sol_gen(a, b, x)
! Check that x and z agree approximately.
    err = sum(abs(x-z))/sum(abs(x))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 3 for LIN_EIG_GEN is correct.'
    end if
    end
```


## Output

Example 3 for LIN_EIG_GEN is correct.

## Example 4: Accuracy Estimates of Eigenvalues Using Adjoint and Ordinary Eigenvectors

A matrix $A$ has entries that are subject to uncertainty. This is expressed as the realization that $A$ can be replaced by the matrix $A+\eta B$, where the value $\eta$ is "small" but still significantly larger than machine precision. The matrix $B$ satisfies $\|B\| \leq\|A\|$. A variation in eigenvalues is estimated using analysis found in Golub and Van Loan, (1989, Chapter 7, p. 344). Each eigenvalue and eigenvector is expanded in a power series in $\eta$. With

$$
e_{i}(\eta) \approx e_{i}+\eta \dot{e}_{i} \eta
$$

and normalized eigenvectors, the bound

$$
\left|\dot{e}_{i}\right| \leq \frac{\|A\|}{\left|u_{i}^{*} v_{i}\right|}
$$

is satisfied. The vectors $u_{i}$ and $v_{i}$ are the ordinary and adjoint eigenvectors associated respectively with $e_{i}$ and its complex conjugate. This gives an upper bound on the size of the change to each $\left|e_{i}\right|$ due to changing the matrix data. The reciprocal

$$
\left|u_{i}^{*} v_{i}\right|^{-1}
$$

is defined as the condition number of $e_{i}$. Also, see operator_ex32, Chapter 10.

```
use lin_eig_gen_int
use rand_gen_int
implicit none
! This is Example 4 for LIN_EIG_GEN.
```

```
    integer i
    integer, parameter :: n=17
    real(kind(1d0)), parameter :: one=1d0
    real(kind(1d0)) a(n,n), c(n,n), variation(n), y(n*n), temp(n), &
        norm of a, eta
    complex(kind(1d\overline{0})), dimension(n,n) :: e(n), d(n), u, v
! Generate a random matrix.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
! Compute the eigenvalues, left- and right- eigenvectors.
    call lin_eig_gen(a, e, v=v, v_adj=u)
! Compute condition numbers and variations of eigenvalues.
    norm_of_a = sqrt(sum(a**2)/n)
    do i=1, n
        variation(i) = norm_of_a/abs(dot_product(u(1:n,i), &
    end do
! Now perturb the data in the matrix by the relative factors
! eta=sqrt(epsilon) and solve for values again. Check the
! differences compared to the estimates. They should not exceed
! the bounds.
    eta = sqrt(epsilon(one))
    do i=1, n
        call rand_gen(temp)
        c(1:n,i) = a(1:n,i) + (2*temp - 1)*eta*a(1:n,i)
    end do
    call lin_eig_gen(c,d)
! Looking at the differences of absolute values accounts for
! switching signs on the imaginary parts.
    if (count(abs(d)-abs(e) > eta*variation) == 0) then
        write (*,*) 'Example 4 for LIN_EIG_GEN is correct.'
    end if
    end
```


## Output

Example 4 for LIN_EIG_GEN is correct.

## LIN_GEIG_GEN

Computes the generalized eigenvalues of an $n \times n$ matrix pencil, $A v=\lambda B v$. Optionally, the generalized eigenvectors are computed. If either of $A$ or $B$ is nonsingular, there are diagonal matrices $\alpha$ and $\beta$, and a complex matrix $V$, all computed such that $A V \beta=B V \alpha$.

## Required Arguments

A- Array of size $n \times n$ containing the matrix $A$. (Input [/Output])
$\boldsymbol{B}$ - Array of size $n \times n$ containing the matrix $B$. (Input [/Output])

ALPHA - Array of size $n$ containing diagonal matrix factors of the generalized eigenvalues. These complex values are in order of decreasing absolute value. (Output)

BETAV - Array of size $n$ containing diagonal matrix factors of the generalized eigenvalues. These real values are in order of decreasing value. (Output)

## Optional Arguments

NROWS $=\mathrm{n}$ (Input)
Uses arrays $A(1: n, 1: n)$ and $B(1: n, 1: n)$ for the input matrix pencil.
Default: $\mathrm{n}=\operatorname{SIZE}(\mathrm{A}, 1)$
$\mathrm{V}=\mathrm{V}(:,:$ ) (Output)
Returns the complex array of generalized eigenvectors for the matrix pencil.
iopt = iopt (: ) (Input)
Derived type array with the same precision as the input matrix. Used for passing optional data to the routine. The options are as follows:

| Packaged Options for LIN_GEIG_GEN |  |  |
| :--- | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| s_, d_, c_, z_ | lin_geig_gen_set_small | 1 |
| s_, d_, c_, z_ | lin_geig_gen_overwrite_input | 2 |
| s_, d_, c_, z_ | lin_geig_gen_scan_for_NaN | 3 |
| s_, d_, c_, z_ | lin_geig_gen_self_adj_pos | 4 |
| s_, d_, c_, z_ | lin_geig_gen_for_lin_sol_self | 5 |
| s_, d_, c_, z_ | lin_geig_gen_for_lin_eig_self | 6 |
| s_, d_, c_, z_ | lin_geig_gen_for_lin_sol_lsq | 7 |
| s_, d_, c_, z_ | lin_geig_gen_for_lin_eig_gen | 8 |

iopt(IO) = ?_options(?_lin_geig_gen_set_small, Small)
This tolerance, multiplied by the sum of absolute value of the matrix $B$, is used to define a small diagonal term in the routines lin_sol_lsq and lin_sol_self. That value can be replaced using the option flags lin_geig_gen_for_lin_sol_lsq, and lin_geig_gen_for_lin_sol_self.
Default: Small = epsilon(.), the relative accuracy of arithmetic

```
iopt(IO) = ?_options(?_lin_geig_gen_overwrite_input, ?_dummy)
```

Does not save the input arrays $\mathrm{A}(:,:)$ and $\mathrm{B}(:,:)$.
Default: The array is saved.
iopt(IO) = ?_options(?_lin_geig_gen_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(a(i,j)) .or. isNaN(b(i,j)) == .true.
See the isnan () function, Chapter 10.
Default: The arrays are not scanned for NaNs.
iopt(IO) = ?_options(?_lin_geig_gen_self_adj_pos, ?_dummy)
If both matrices $A$ and $B$ are self-adjoint and additionally $B$ is positive-definite, then the Cholesky algorithm is used to reduce the matrix pencil to an ordinary self-adjoint eigenvalue problem.

```
iopt(IO) = ?_options(?_lin_geig_gen_for_lin_sol_self, ?_dummy)
iopt(IO+1) = ?_options((k=size of options for lin_sol_self), ?_dummy)
```

    The options for lin_sol_self follow as data in iopt ().
    ```
iopt(IO) = ?_options(?_lin_geig_gen_for_lin_eig_self, ?_dummy)
```

iopt(IO+1) = ?_options((k=size of options for lin_eig_self), ?_dummy)
The options for lin_eig_self follow as data in iopt ().
iopt(IO) = ?_options(?_lin_geig_gen_for_lin_sol_lsq, ?_dummy)
iopt(IO+1) = ?_options((k=size of options for lin_sol_lsq), ?_dummy)
The options for lin_sol_lsq follow as data in iopt ().
iopt(IO) = ?_options(?_lin_geig_gen_for_lin_eig_gen, ?_dummy)
iopt(IO+1) = ?_options((k=size of options for lin_eig_gen), ?_dummy)
The options $\bar{f}$ or lin_eig_gen follow as data in iopt ().

## FORTRAN 90 Interface

Generic: CALL LIN_GEIG_GEN (A, B, ALPHA, BETAV [, ...])
Specific: The specific interface names are S_LIN_GEIG_GEN, D_LIN_GEIG_GEN, C_LIN_GEIG_GEN, and Z_LIN_GEIG_GEN.

## Description

Routine LIN_GEIG_GEN implements a standard algorithm that reduces a generalized eigenvalue or matrix pencil problem to an ordinary eigenvalue problem. An orthogonal decomposition is computed

$$
B P^{T}=H R
$$

The orthogonal matrix $H$ is the product of $n-1$ row permutations, each followed by a Householder transformation. Column permutations, $P$, are chosen at each step to maximize the Euclidian length of the pivot column. The matrix $R$ is upper triangular. Using the default tolerance $\tau=\varepsilon\|B\|$, where $\varepsilon$ is machine relative precision, each diagonal entry of $R$ exceeds $\tau$ in value.
Otherwise, $R$ is singular. In that case $A$ and $B$ are interchanged and the orthogonal decomposition is computed one more time. If both matrices are singular the problem is declared singular and is not solved. The interchange of $A$ and $B$ is accounted for in the output diagonal matrices $\alpha$ and $\beta$. The ordinary eigenvalue problem is $C x=\lambda x$, where

$$
C=H^{T} A P^{T} R^{-1}
$$

and

$$
R P v=x
$$

If the matrices $A$ and $B$ are self-adjoint and if, in addition, $B$ is positive-definite, then a more efficient reduction than the default algorithm can be optionally used to solve the problem: A Cholesky decomposition is obtained, $R^{T} R R=P B P^{T}$. The matrix $R$ is upper triangular and $P$ is a permutation matrix. This is equivalent to the ordinary self-adjoint eigenvalue problem $C x=\lambda x$, where $R P v=x$ and

$$
C=R^{-T} P A P^{T} R^{-1}
$$

The self-adjoint eigenvalue problem is then solved.

## Fatal, Terminal, and Warning Error Messages

See the messages.gls file for error messages for LIN_GEIG_GEN. These error messages are numbered 921-936; 941-956; 961-976; 981-996.

## Example 1: Computing Generalized Eigenvalues

The generalized eigenvalues of a random real matrix pencil are computed. These values are checked by obtaining the generalized eigenvectors and then showing that the residuals

$$
A V-B V \alpha \beta^{-1}
$$

are small. Note that when the matrix $B$ is nonsingular $\beta=I$, the identity matrix. When $B$ is singular and $A$ is nonsingular, some diagonal entries of $\beta$ are essentially zero. This corresponds to "infinite eigenvalues" of the matrix pencil. This random matrix pencil example has all finite eigenvalues. Also, see operator_ex33, Chapter 10.

```
    use lin_geig_gen_int
    use ran\overline{d_gen_int}
    implicit none
! This is Example 1 for LIN_GEIG_GEN.
    integer, parameter :: n=32
    real(kind(1d0)), parameter :: one=1d0
    real(kind(1d0)) A(n,n), B(n,n), betav(n), beta t(n), err, y(n*n)
    complex(kind(1d0)) alpha(n), alpha_t(n), V(n,n)
```

```
! Generate random matrices for both A and B.
    call rand_gen(y)
    A = reshape(y,(/n,n/))
    call rand_gen(y)
    B = reshape(y,(/n,n/))
! Compute the generalized eigenvalues.
    call lin_geig_gen(A, B, alpha, betav)
! Compute the full decomposition once again, A*V = B*V*values.
    call lin_geig_gen(A, B, alpha_t, beta_t, &
        v=V)
! Use values from the first decomposition, vectors from the
! second decomposition, and check for small residuals.
    err = sum(abs(matmul(A,V) - &
                matmul(B,V)*spread(alpha/betav,DIM=1,NCOPIES=n))) / &
            sum(abs(a)+abs (b))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_GEIG_GEN is correct.'
    end if
    end
```


## Output

Example 1 for LIN_GEIG_GEN is correct.

## Additional Examples

## Example 2: Self-Adjoint, Positive-Definite Generalized Eigenvalue Problem

This example illustrates the use of optional flags for the special case where $A$ and B are complex self-adjoint matrices, and $B$ is positive-definite. For purposes of maximum efficiency an option is passed to routine LIN_SOL_SELF so that pivoting is not used in the computation of the Cholesky decomposition of matrix $B$. This example does not require that secondary option. Also, see operator_ex34, supplied with the product examples.

```
    use lin_geig_gen_int
    use lin_sol_self_int
    use rand_gen_int
    implicit none
! This is Example 2 for LIN_GEIG_GEN.
    integer i
    integer, parameter :: n=32
    real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
    real(kind(1d0)) betav(n), temp_c(n,n), temp_d(n,n), err
    type(d_options) :: iopti(4)=d_options(0, zero)
    comple\overline{x}(kind(ld0)), dimension(n,n) :: A, B, C, D, V, alpha(n)
```

```
! Generate random matrices for both A and B.
    do i=1, n
        call rand_gen(temp_c(1:n,i))
        call rand_gen(temp_d(1:n,i))
    end do
    c = temp_c; d = temp_c
    do i=1, \overline{n}
            call rand_gen(temp_c(1:n,i))
            call rand_gen(temp_d(1:n,i))
    end do
    c = cmplx(real(c),temp_c,kind(one))
    d = cmplx(real(d),temp_d,kind(one))
    a = conjg(transpose(c)) + c
    b = matmul(conjg(transpose(d)),d)
! Set option so that the generalized eigenvalue solver uses an
! efficient method for well-posed, self-adjoint problems.
    iopti(1) = d_options(z_lin_geig_gen_self_adj_pos,zero)
    iopti(2) = d_options(z_lin_geig_gen_for_lin_sol_self,zero)
! Number of secondary optional data items and the options:
    iopti(3) = d_options(1,zero)
    iopti(4) = d_options(z_lin_sol_self_no_pivoting,zero)
    call lin_geig_gen(a, b, alpha, betav, v=v, &
        iopt=iopti)
! Check that residuals are small. Use the real part of alpha
! since the values are known to be real.
    err = sum(abs(matmul(a,v) - matmul(b,v)* &
                spread(real(alpha,kind(one))/betav,dim=1,ncopies=n))) / &
                sum(abs(a)+abs(b))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 2 for LIN_GEIG_GEN is correct.'
    end if
    end
```


## Output

Example 2 for LIN_GEIG_GEN is correct.

## Example 3: A Test for a Regular Matrix Pencil

In the classification of Differential Algebraic Equations (DAE), a system with linear constant coefficients is given by $A \dot{x}+B x=f$. Here $A$ and $B$ are $n \times n$ matrices, and $f$ is an $n$-vector that is not part of this example. The DAE system is defined as solvable if and only if the quantity det $(\mu A+B)$ does not vanish identically as a function of the dummy parameter $\mu$. A sufficient condition for solvability is that the generalized eigenvalue problem $A v=\lambda B v$ is nonsingular. By constructing $A$ and $B$ so that both are singular, the routine flags nonsolvability in the DAE by returning NaN for the generalized eigenvalues. Also, see operator_ex 35 , supplied with the product examples.

```
    use lin_geig_gen_int
    use rand_gen_int
    use error
    use isnan_int
    implicit none
! This is Example 3 for LIN_GEIG_GEN.
    integer, parameter :: n=6
    real(kind(1d0)), parameter :: one=1.0d0, zero=0.0d0
    real(kind(1d0)) a(n,n), b (n,n), betav(n), y(n*n)
    type(d_options) iopti(1)
    type(d_error) epack(1)
    comple\overline{x}(kind(1d0)) alpha(n)
! Generate random matrices for both A and B.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
    call rand_gen(y)
    b = reshape(y, (/n,n/))
! Make columns of A and B zero, so both are singular.
    a(1:n,n)=0; b (1:n,n)=0
! Set internal tolerance for a small diagonal term.
    iopti(1) = d_options(d_lin_geig_gen_set_small,sqrt(epsilon(one)))
! Compute the generalized eigenvalues.
    call lin_geig_gen(a, b, alpha, betav, &
        iopt=iopti,epack=epack)
! See if singular DAE system is detected.
! (The size of epack() is too small for the message, so
! output is blocked with NaNs.)
    if (isnan(alpha)) then
        write (*,*) 'Example 3 for LIN_GEIG_GEN is correct.'
    end if
    end
```


## Output

Example 3 for LIN_GEIG_GEN is correct.

## Example 4: Larger Data Uncertainty than Working Precision

Data values in both matrices $A$ and $B$ are assumed to have relative errors that can be as large as $\varepsilon^{1 / 2}$ where $\varepsilon$ is the relative machine precision. This example illustrates the use of an optional flag that resets the tolerance used in routine $1 i n \_s o l \_l$ sq for determining a singularity of either matrix. The tolerance is reset to the new value $\varepsilon^{1 / 2}\|B\|$ and the generalized eigenvalue problem is
solved. We anticipate that $B$ might be singular and detect this fact. Also, see operator_ex 36 , Chapter 10.

```
    use lin geig gen int
    use lin_sol_lsq_int
    use rand_gen_int
    use isNa\overline{N}_in\overline{t}
    implicit none
! This is Example 4 for LIN_GEIG_GEN.
    integer, parameter :: n=32
    real(kind(1d0)), parameter :: one=1d0, zero=0d0
    real(kind(1d0)) a(n,n), b(n,n), betav(n), y(n*n), err
    type(d_options) iopti(4)
    type(d_error) epack(1)
    complex(kind(1d0)) alpha(n), v(n,n)
! Generate random matrices for both A and B.
    call rand_gen(y)
    a = reshape(y,(/n,n/))
    call rand_gen(y)
    b = reshape(y, (/n,n/))
! Set the option, a larger tolerance than default for lin_sol_lsq.
    iopti(1) = d_options(d_lin_geig_gen_for_lin_sol_lsq,zerō)
! Number of secondary optional data items
    iopti(2) = d_options(2,zero)
    iopti(3) = d_options(d_lin_sol_lsq_set_small,sqrt(epsilon(one))*&
                    sqrt(sum(b\overline{**2)/n))}
    iopti(4) = d_options(d_lin_sol_lsq_no_sing_mess,zero)
! Compute the generalized eigenvalues.
    call lin_geig_gen(A, B, alpha, betav, v=v, &
            iopt=iopti, epack=epack)
    if(.not. isNaN(alpha)) then
! Check the residuals.
        err = sum(abs (matmul (A,V)*spread(betav,dim=1,ncopies=n) - &
                        matmul(B,V)*spread(alpha,dim=1,ncopies=n))) / &
            sum(abs(a)+abs(b))
        if (err <= sqrt(epsilon(one))) then
            write (*,*) 'Example 4 for LIN_GEIG_GEN is correct.'
        end if
    end if
    end
```


## Output

Example 4 for LIN_GEIG_GEN is correct.

## EVLRG

Computes all of the eigenvalues of a real matrix.

## Required Arguments

$\boldsymbol{A}-$ Real full matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Complex vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLRG (A, EVAL [,...])
Specific: The specific interface names are S_EVLRG and D_EVLRG.

## FORTRAN 77 Interface

Single: CALL EVLRG (N, A, LDA, EVAL)
Double: The double precision name is DEVLRG.

## Description

Routine EVLRG computes the eigenvalues of a real matrix. The matrix is first balanced. Elementary or Gauss similarity transformations with partial pivoting are used to reduce this balanced matrix to a real upper Hessenberg matrix. A hybrid double-shifted LR-QR algorithm is used to compute the eigenvalues of the Hessenberg matrix, Watkins and Elsner (1990).

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. The LR-QR algorithm is based on software work of Watkins and Haag. Further details, some timing data, and credits are given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LRG/DE3LRG. The reference is:

CALL E3LRG (N, A, LDA, EVAL, ACOPY, WK, IWK)
The additional arguments are as follows:
ACOPY - Real work array of length $N^{2}$. A and ACOPY may be the same, in which case the first $N^{2}$ elements of A will be destroyed.
$W K-$ Floating-point work array of size 4 N .
$I W K$ - Integer work array of size 2 N .
2. Informational error

Type Code
41 The iteration for an eigenvalue failed to converge.
3. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E3LRG, the internal or working leading dimension of ACOPY is increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine EVLRG. Additional memory allocation and option value restoration are automatically done in EVLRG. There is no requirement that users change existing applications that use EVLRG or E3LRG. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$ Items 5-8 in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVLRG.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 85 ). The eigenvalues of this real matrix are computed and printed. The exact eigenvalues are known to be $\{4,3,2,1\}$.

```
USE EVLRG_INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (N=4, LDA=N)
REAL A(LDA,N)
COMPLEX EVAL (N)
```

$!$
!
$!$

```
!
!
    A=(\begin{array}{llll}{-2.0 2.0 2.0 2.0}\end{array})
    (\begin{array}{llll}{-3.0}&{3.0}&{2.0}&{2.0}\end{array})
    (\begin{array}{llll}{-2.0}&{0.0}&{4.0}&{2.0}\end{array})
    (\begin{array}{llll}{-1.0 0.0 0.0 5.0 )}\end{array})
    DATA A/-2.0, -3.0, -2.0, -1.0, 2.0, 3.0, 0.0, 0.0, 2.0, 2.0, &
    4.0, 0.0, 2.0, 2.0, 2.0, 5.0/
! Find eigenvalues of A
    CALL EVLRG (A, EVAL)
END
```


## Output

$(4.000,0.000)^{1} \quad(3.000,0.000)^{2} \quad(2.000,0.000)^{3} \quad(1.000,0.000)^{4}$

## EVCRG

Computes all of the eigenvalues and eigenvectors of a real matrix.

## Required Arguments

$\boldsymbol{A}$ - Floating-point array containing the matrix. (Input)
$\boldsymbol{E V A L}$ - Complex array of size N containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Complex array containing the matrix of eigenvectors. (Output) The J-th eigenvector, corresponding to $\operatorname{EVAL}(J)$, is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVCRG (A, EVAL, EVEC [,...])

Specific: The specific interface names are S_EVCRG and D_EVCRG.

## FORTRAN 77 Interface

Single: CALL EVCRG (N, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCRG.

## Description

Routine EVCRG computes the eigenvalues and eigenvectors of a real matrix. The matrix is first balanced. Orthogonal similarity transformations are used to reduce the balanced matrix to a real upper Hessenberg matrix. The implicit double-shifted QR algorithm is used to compute the eigenvalues and eigenvectors of this Hessenberg matrix. The eigenvectors are normalized such that each has Euclidean length of value one. The largest component is real and positive.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. Further details, some timing data, and credits are given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E8CRG/DE8CRG. The reference is:

CALL E8CRG (N, A, LDA, EVAL, EVEC, LDEVEC, ACOPY, ECOPY, WK, IWK)
The additional arguments are as follows:

ACOPY - Floating-point work array of size $n$ by n. The arrays A and ACOPY may be the same, in which case the first $N^{2}$ elements of A will be destroyed. The array ACOPY can have its working row dimension increased from N to a larger value. An optional usage is required. See Item 3 below for further details.
$\boldsymbol{E C O P Y}$ - Floating-point work array of default size N by $\mathrm{N}+1$. The working, leading dimension of ECOPY is the same as that for ACOPY. To increase this value, an optional usage is required. See Item 3 below for further details.

WK — Floating-point work array of size 6 N .

IWK - Integer work array of size N.
2. Informational error

Type Code

1 The iteration for the eigenvalues failed to converge. No eigenvalues or eigenvectors are computed.
3. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E8CRG, the internal or working leading dimensions of ACOPY and ECOPY are both increased by IVAL(3) when $N$ is a multiple of $\operatorname{IVAL}(4)$. The values $\operatorname{IVAL}(3)$ and $\operatorname{IVAL}(4)$ are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine EVCRG. Additional memory allocation and option value restoration are automatically done in EVCRG. There is no requirement that users change existing applications that use EVCRG or E8CRG. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$. Items $5-8$ in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVCRG.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 82). The eigenvalues and eigenvectors of this real matrix are computed and printed. The performance index is also computed and printed. This serves as a check on the computations. For more details, see IMSL routine EPIRG.

```
USE EVCRG_INT
USE EPIRG }\mp@subsup{}{}{-}\mathrm{ INT
USE UMACH -}\mp@subsup{}{}{-
USE WRCRN_INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, N
PARAMETER (N=3, LDA=N, LDEVEC=N)
INTEGER NOUT
REAL PI
COMPLEX EVAL(N), EVEC(LDEVEC,N)
REAL A(LDA,N)
    Define values of A:
        A=( 8.0 -1.0 -5.0 )
            ( -4.0 4.0 -2.0 )
            (18.0 -5.0 -7.0)
DATA A/8.0, -4.0, 18.0, -1.0, 4.0, -5.0, -5.0, -2.0, -7.0/
CALL EVCRG (A, EVAL, EVEC)
PI = EPIRG(N,A,EVAL,EVEC)
CALL UMACH (2, NOUT)
CALL WRCRN ('EVAL', EVAL, 1, N, 1)
CALL WRCRN ('EVEC', EVEC)
```

        WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
        END
    
## Output

EVAL
$(2.000,4.000)^{1} \quad(2.000,-4.000)^{2} \quad(1.000,0.000)^{3}$

EVEC

|  |  | ${ }^{1}$ | ${ }^{2}$ |
| :--- | ---: | ---: | ---: | ---: | ---: |
| 1 | $(0.3162,0.3162)$ | $(0.3162,-0.3162)$ | $(0.4082,0.0000)$ |
| 2 | $(-0.0000,0.6325)$ | $(-0.0000,-0.6325)$ | $(0.8165,0.0000)$ |
| 3 | $(0.6325,0.0000)$ | $(0.6325,0.0000)$ | $(0.4082,0.0000)$ |

Performance index $=0.026$

## EPIRG

This function computes the performance index for a real eigensystem.

## Function Return Value

EPIRG - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs on which the performance index computation is based. (Input)
$\boldsymbol{A}$ - Matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Complex vector of length NEVAL containing eigenvalues of A. (Input)
$\boldsymbol{E V E C}$ - Complex n by NEVAL array containing eigenvectors of A. (Input) The eigenvector corresponding to the eigenvalue $\operatorname{EVAL}(J)$ must be in the $J$-th column of EVEC.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: EPIRG (NEVAL, A, EVAL, EVEC [, ...])
Specific: The specific interface names are S_EPIRG and D_EPIRG.

## FORTRAN 77 Interface

Single: EPIRG (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision function name is DEPIRG.

## Description

Let $M=$ NEVAL, $\lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, \mathrm{~J})$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision given by АМАСн(4). The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} x_{j}\right\|_{1}}{10 N \varepsilon\|A\|_{1}\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$.
The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Smith et al. (1976, pages 124-125).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E2IRG/DE2IRG. The reference is:

E2IRG (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC, CWK)
The additional argument is:
$\boldsymbol{C W K}$ - Complex work array of length N .
2. Informational errors
Type Code

| 3 | 1 | The performance index is greater than 100. |
| :--- | :--- | :--- |
| 3 | 2 | An eigenvector is zero. |
| 3 | 3 | The matrix is zero. |

## Example

For an example of EPIRG, see IMSL routine EVCRG.

## EVLCG

Computes all of the eigenvalues of a complex matrix.

## Required Arguments

$\boldsymbol{A}-$ Complex matrix of order N. (Input)
EVAL - Complex vector of length $N$ containing the eigenvalues of $A$ in decreasing order of magnitude. (Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLCG (A, EVAL [,...])
Specific: The specific interface names are S_EVLCG and D_EVLCG.

## FORTRAN 77 Interface

Single: CALL EVLCG (N, A, LDA, EVAL)

Double: The double precision name is EVLCG.

## Description

Routine EVLCG computes the eigenvalues of a complex matrix. The matrix is first balanced.
Unitary similarity transformations are used to reduce this balanced matrix to a complex upper Hessenberg matrix. The shifted QR algorithm is used to compute the eigenvalues of this Hessenberg matrix.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LCG/DE3LCG. The reference is:

CALL E3LCG (N, A, LDA, EVAL, ACOPY, RWK, CWK, IWK)
The additional arguments are as follows:

ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of A will be destroyed.
$\boldsymbol{R W K}$ - Work array of length N .

CWK - Complex work array of length 2 N .
IWK - Integer work array of length N.
2. Informational error

Type Code
41 The iteration for an eigenvalue failed to converge.
3. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E3LCG, the internal or working, leading dimension of ACOPY is increased by IVAL( 3 ) when $N$ is a multiple of IVAL(4). The values $\operatorname{IVAL}(3)$ and IVAL (4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine EVLCG . Additional memory allocation and option value restoration are automatically done in EVLCG. There is no requirement that users change existing applications that use EVLCG or E3LCG. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$. Items $5-8$ in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVLCG.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 115). The program computes the eigenvalues of this matrix.

```
USE EVLCG_INT
USE WRCRN_INT
! Declare variables
INTEGER LDA, N
PARAMETER (N=3, LDA=N)
```

!

```
    COMPLEX A(LDA,N), EVAL (N)
        Set values of A
        A = ( 1+2i 3+4i 21+22i)
    (43+44i 13+14i 15+16i)
    ( 5+6i 7+8i 25+26i)
    DATA A/ (1.0,2.0), (43.0,44.0), (5.0,6.0), (3.0,4.0), &
        (13.0,14.0), (7.0,8.0), (21.0,22.0), (15.0,16.0), &
        (25.0,26.0)/
    CALL EVLCG (A, EVAL)
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
    END
```


## Output



## EVCCG

Computes all of the eigenvalues and eigenvectors of a complex matrix.

## Required Arguments

$\boldsymbol{A}$ - Complex matrix of order N. (Input)

EVAL - Complex vector of length $N$ containing the eigenvalues of $A$ in decreasing order of magnitude. (Output)

EVEC - Complex matrix of order N. (Output)
The J-th eigenvector, corresponding to EVAL(J), is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL EVCCG (A, EVAL, EVEC [,...])

Specific: The specific interface names are S_EVCCG and D_EVCCG.

## FORTRAN 77 Interface

Single: CALL EVCCG (N, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCCG.

## Description

Routine EVCCG computes the eigenvalues and eigenvectors of a complex matrix. The matrix is first balanced. Unitary similarity transformations are used to reduce this balanced matrix to a complex upper Hessenberg matrix. The QR algorithm is used to compute the eigenvalues and eigenvectors of this Hessenberg matrix. The eigenvectors of the original matrix are computed by transforming the eigenvectors of the complex upper Hessenberg matrix.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E6CCG/DE6CCG. The reference is:

CALL E6CCG (N, A, LDA, EVAL, EVEC, LDEVEC, ACOPY, RWK, CWK, IWK)
The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. The arrays A and ACOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of A will be destroyed.
$\boldsymbol{R W K}$ - Work array of length N .
CWK — Complex work array of length 2 N .
$I W K$ - Integer work array of length N.
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge. No eigenvalues or eigenvectors are computed.
3. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E 6 CCG , the internal or working leading dimensions of $A C O P Y$ and ECOPY are both increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values $\operatorname{IVAL}(3)$ and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine EVCCG. Additional memory allocation and option value restoration are automatically done in EVCCG. There is no requirement that users change existing applications that use EVCCG or E6CCG. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$. Items 5-8 in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVCCG.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 116). Its eigenvalues are known to be $\{1+5 i, 2+6 i, 3+7 i, 4+8 i\}$. The program computes the eigenvalues and eigenvectors of this matrix. The performance index is also computed and printed. This serves as a check on the computations, for more details, see IMSL routine EPICG.

```
USE EVCCG INT
USE EPICG - INT
USE WRCRN_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, N
PARAMETER (N=4, LDA=N, LDEVEC=N)
INTEGER NOUT
REAL PI
COMPLEX A(LDA,N), EVAL(N), EVEC(LDEVEC,N)
                                    Set values of A
                                    A = (5+9i 5+5i -6-6i -7-7i)
                                    (3+3i 6+10i -5-5i -6-6i)
                                    (2+2i 3+3i -1+3i -5-5i)
                                    (1+i 2+2i -3-3i 4i)
DATA A/ (5.0,9.0), (3.0,3.0), (2.0,2.0), (1.0,1.0), (5.0,5.0), &
        (6.0,10.0), (3.0,3.0), (2.0,2.0), (-6.0,-6.0), (-5.0,-5.0), &
        (-1.0,3.0), (-3.0,-3.0), (-7.0,-7.0), (-6.0,-6.0), &
        (-5.0,-5.0), (0.0,4.0)/
            Find eigenvalues and vectors of A
    CALL EVCCG (A, EVAL, EVEC)
        Compute performance index
PI = EPICG(N,A,EVAL,EVEC)
CALL UMACH (2, NOUT)
CALL WRCRN ('EVAL', EVAL, 1, N, 1)
CALL WRCRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```

$!$

## Output

```
    (4.000, 8.000) (3.000, 7.000) ( %VAL (2.000, 6.000) 3)
                                    EVEC
                                1 2 3
4
    1 ( 0.5774, 0.0000) ( 0.5774, 0.0000) ( 0.3780, 0.0000) ( 0.7559,
0.0000)
    2 (0.5774,-0.0000) ( 0.5773,-0.0000) ( 0.7559, 0.0000) ( 0.3780,
0.0000)
    3 ( 0.5774,-0.0000) (-0.0000,-0.0000) ( 0.3780, 0.0000) ( 0.3780,
0.0000)
    4 ( 0.0000, 0.0000) ( 0.5774, 0.0000) ( 0.3780, 0.0000) ( 0.3780,
0.0000)
    Performance index = 0.016
```


## EPICG

This function computes the performance index for a complex eigensystem.

## Function Return Value

EPICG - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs on which the performance index computation is based. (Input)

A-Complex matrix of order N . (Input)
EVAL - Complex vector of length N containing the eigenvalues of A . (Input)
EVEC - Complex matrix of order N containing the eigenvectors of A. (Input) The J -th eigenvalue/eigenvector pair should be in $\operatorname{EVAL}(\mathrm{J})$ and in the J -th column of evec.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{n}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: EPICG (NEVAL, A, EVAL, EVEC [, ...])
Specific: The specific interface names are S_EPICG and D_EPICG.

## FORTRAN 77 Interface

Single: EPICG (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision function name is DEPICG.

## Description

Let $M=$ NEVAL, $\lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, J)$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision given by АМАСН(4). The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} x_{j}\right\|_{1}}{10 N \varepsilon\|A\|_{1}\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Smith et al. (1976, pages 124-125).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E2ICG/DE2ICG. The reference is:

E2ICG (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC, WK)
The additional argument is:
$W K-$ Complex work array of length $N$.
2. Informational errors

Type Code
$31 \quad$ Performance index is greater than 100.
32 An eigenvector is zero.

## Example

For an example of EPICG, see IMSL routine EVCCG.

## EVLSF

Computes all of the eigenvalues of a real symmetric matrix.

## Required Arguments

A-Real symmetric matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Real vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLSF (A, EVAL [,...])
Specific: The specific interface names are S_EVLSF and D_EVLSF.

## FORTRAN 77 Interface

Single: CALL EVLSF (N, A, LDA, EVAL)
Double: The double precision name is DEVLSF.

## Description

Routine EVLSF computes the eigenvalues of a real symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. Then, an implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E4LSF/DE4LSF. The reference is:

CALL E4LSF (N, A, LDA, EVAL, WORK, IWORK)
The additional arguments are as follows:
WORK - Work array of length 2 N .
IWORK - Integer array of length N.
2. Informational error

Type Code
31 The iteration for the eigenvalue failed to converge in 100 iterations before deflating.

## Example

In this example, the eigenvalues of a real symmetric matrix are computed and printed. This matrix is given by Gregory and Karney (1969, page 56).

```
USE EVLSF INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (N=4, LDA=N)
REAL A(LDA,N), EVAL(N)
    Set values of A
        A = ( 6.0 4.0 4.0 1.0)
                            ( 4.0 6.0 1.0 4.0)
                            ( 4.0 1.0 6.0 4.0)
            (\begin{array}{llll}{1.0}&{4.0}&{4.0}&{6.0}\end{array})
DATA A /6.0, 4.0, 4.0, 1.0, 4.0, 6.0, 1.0, 4.0, 4.0, 1.0, 6.0, &
        4.0, 1.0, 4.0, 4.0, 6.0 /
CALL EVLSF (A, EVAL)
    Find eigenvalues of A
CALL WRRRN ('EVAL', EVAL, 1, N, 1)
END
```


## Output

|  | EVAL |  |  |  |
| ---: | ---: | ---: | ---: | :---: |
| 1 | 2 | 3 | 4 |  |
| 15.00 | 5.00 | 5.00 | -1.00 |  |

## EVCSF

Computes all of the eigenvalues and eigenvectors of a real symmetric matrix.

## Required Arguments

A-Real symmetric matrix of order N. (Input)
EVAL - Real vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)
$\boldsymbol{E V E C}$ - Real matrix of order N. (Output)
The J-th eigenvector, corresponding to EVAL(J), is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}($ EVEC, 1$)$.

## FORTRAN 90 Interface

Generic: CALL EVCSF (A, EVAL, EVEC [, ...])
Specific: The specific interface names are S_EVCSF and D_EVCSF.

## FORTRAN 77 Interface

Single: CALL EVCSF (N, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCSF.

## Description

Routine EVCSF computes the eigenvalues and eigenvectors of a real symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. These transformations are accumulated. An implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix. The eigenvectors are computed using the eigenvalues as perfect shifts, Parlett (1980, pages 169, 172). The underlying code is based on
either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. Further details, some timing data, and credits are given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E5CSF/DE5CSF. The reference is:

CALL E5CSF (N, A, LDA, EVAL, EVEC, LDEVEC, WORK, IWK)
The additional argument is:
WORK - Work array of length 3 N .
IWK - Integer array of length N .
2. Informational error

Type Code
31 The iteration for the eigenvalue failed to converge in 100 iterations before deflating.

## Example

The eigenvalues and eigenvectors of this real symmetric matrix are computed and printed. The performance index is also computed and printed. This serves as a check on the computations. For more details, see EPISF.

```
USE EVCSF INT
USE EPISF INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDEVEC, N
PARAMETER (N=3, LDA=N, LDEVEC=N)
INTEGER NOUT
REAL A(LDA,N), EVAL (N), EVEC(LDEVEC,N), PI
    Set values of A
    A=(\begin{array}{lll}{7.0}&{-8.0}&{-8.0)}\end{array})
        (-8.0 -16.0 -18.0)
        ( -8.0 -18.0 13.0)
DATA A/7.0, -8.0, -8.0, -8.0, -16.0, -18.0, -8.0, -18.0, 13.0/
Find eigenvalues and vectors of A
CALL EVCSF (A, EVAL, EVEC)
```

```
! Compute performance index
    PI = EPISF (N, A, EVAL, EVEC)
    CALL UMACH (2, NOUT)
    CALL WRRRN ('EVAL', EVAL, 1, N, 1)
    CALL WRRRN ('EVEC', EVEC)
    WRITE (NOUT, '(/,A,F6.3)') ' Performance index = ', PI
    END
```

Output

| EVAL |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -27.90 | 22.68 | 9.22 |

            EVEC
    |  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 0.2945 | -0.2722 | 0.9161 |
| 2 | 0.8521 | -0.3591 | -0.3806 |
| 3 | 0.4326 | 0.8927 | 0.1262 |

Performance index $=0.019$

## EVASF

Computes the largest or smallest eigenvalues of a real symmetric matrix.

## Required Arguments

NEVAL - Number of eigenvalues to be computed. (Input)
A-Real symmetric matrix of order N. (Input)

SMALL - Logical variable. (Input)
If .TRUE., the smallest NEVAL eigenvalues are computed. If .FALSE., the largest NEVAL eigenvalues are computed.

EVAL - Real vector of length NEVAL containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVASF (NEVAL, A, SMALL, EVAL [, ...])
Specific: The specific interface names are S_EVASF and D_EVASF.

## FORTRAN 77 Interface

Single: CALL EVASF (N, NEVAL, A, LDA, SMALL, EVAL)
Double: The double precision name is DEVASF.

## Description

Routine EVASF computes the largest or smallest eigenvalues of a real symmetric matrix.
Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. Then, an implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix.

The reduction routine is based on the EISPACK routine TRED2. See Smith et al. (1976). The rational QR algorithm is called the PWK algorithm. It is given in Parlett (1980, page 169).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E4ASF/DE4ASF. The reference is:

CALL E4ASF (N, NEVAL, A, LDA, SMALL, EVAL, WORK, IWK)
WORK - Work array of length 4 N .
IWK - Integer work array of length N .
2. Informational error

Type Code
31 The iteration for an eigenvalue failed to converge. The best estimate will be returned.

## Example

In this example, the three largest eigenvalues of the computed Hilbert matrix $a_{i j}=1 /(i+j-1)$ of order $N=10$ are computed and printed.

```
USE EVASF INT
USE WRRRN__INT
IMPLICIT NONE
\begin{tabular}{lll} 
& & Declare variables \\
INTEGER & LDA, \(N, \quad\) NEVAL & \\
PARAMETER & \((N=10, \quad\) NEVAL \(=3, \quad L D A=N)\)
\end{tabular}
```

!
$!$

```
INTEGER I, J
    REAL A(LDA,N), EVAL (NEVAL), REAL
    LOGICAL SMALL
    INTRINSIC REAL
    DO 20 J=1, N
        DO 10 I=1, N
                A(I,J) = 1.0/REAL(I+J-1)
    CONTINUE
    CONTINUE
    SMALL = .FALSE.
    CALL EVASF (NEVAL, A, SMALL, EVAL)
                                    Print results
    CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
    END
```


## Output

EVAL

| 1 | 2 | 3 |
| ---: | ---: | ---: |
| 1.752 | 0.343 | 0.036 |

## EVESF

Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix.

## Required Arguments

NEVEC - Number of eigenvalues to be computed. (Input)
A-Real symmetric matrix of order N. (Input)
SMALL - Logical variable. (Input)
If .TRUE., the smallest NEVEC eigenvalues are computed. If .FALSE., the largest NEVEC eigenvalues are computed.
$\boldsymbol{E V A L}$ - Real vector of length NEVEC containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Real matrix of dimension $n$ by NEVEC. (Output)
The J-th eigenvector, corresponding to EVAL(J), is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDEVEC - Leading dimension of evEC exactly as specified in the dimension statement in the calling program. (Input)
Default: $\operatorname{LDEVEC}=\operatorname{SIZE}(\operatorname{EVEC}, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVESF (NEVEC, A, SMALL, EVAL, EVEC [,$\ldots]$ )
Specific: The specific interface names are S_EVESF and D_EVESF.

## FORTRAN 77 Interface

Single: CALL EVESF (N, NEVEC, A, LDA, SMALL, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVESF.

## Description

Routine EVESF computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. Then, an implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix. Inverse iteration is used to compute the eigenvectors of the tridiagonal matrix. This is followed by orthogonalization of these vectors. The eigenvectors of the original matrix are computed by back transforming those of the tridiagonal matrix.

The reduction routine is based on the EISPACK routine TRED2. See Smith et al. (1976). The rational QR algorithm is called the PWK algorithm. It is given in Parlett (1980, page 169). The inverse iteration and orthogonalization computation is discussed in Hanson et al. (1990). The back transformation routine is based on the EISPACK routine TRBAK1.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E5ESF/DE5ESF. The reference is:

CALL E5ESF (N, NEVEC, A, LDA, SMALL, EVAL, EVEC, LDEVEC, WK, IWK)
The additional arguments are as follows:

WK-Work array of length 9 N .
IWK - Integer array of length N.
2. Informational errors

Type Code
31 The iteration for an eigenvalue failed to converge. The best estimate will be returned.
32 Inverse iteration did not converge. Eigenvector is not correct for the specified eigenvalue.
33 The eigenvectors have lost orthogonality.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 55). The largest two eigenvalues and their eigenvectors are computed and printed. The performance index is also computed and printed. This serves as a check on the computations. For more details, see IMSL routine EPISF.

```
USE EVESF INT
USE EPISF INT
USE UMACH_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, N
PARAMETER ( }\textrm{N}=4,LDA=N, LDEVEC=N
INTEGER NEVEC, NOUT
REAL A(LDA,N), EVAL(N), EVEC(LDEVEC,N), PI
LOGICAL SMALL
                                    Set values of A
                                    A = ( 5.0 4.0 1.0 1.0)
            ( 4.0 5.0 1.0 1.0)
            (\begin{array}{llll}{1.0}&{1.0}&{4.0}&{2.0}\end{array})
            (1.0 1.0 2.0 4.0)
DATA A/5.0, 4.0, 1.0, 1.0, 4.0, 5.0, 1.0, 1.0, 1.0, 1.0, 4.0, &
            2.0, 1.0, 1.0, 2.0, 4.0/
                                    Find eigenvalues and vectors of A
NEVEC = 2
SMALL = . FALSE.
CALL EVESF (NEVEC, A, SMALL, EVAL, EVEC)
                            Compute performance index
PI = EPISF(NEVEC,A,EVAL,EVEC)
                                    Print results
CALL UMACH (2, NOUT)
CALL WRRRN ('EVAL', EVAL, 1, NEVEC, 1)
CALL WRRRN ('EVEC', EVEC, N, NEVEC, LDEVEC)
```

```
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output



## EVBSF

Computes selected eigenvalues of a real symmetric matrix.

## Required Arguments

MXEVAL - Maximum number of eigenvalues to be computed. (Input)
$\boldsymbol{A}$ - Real symmetric matrix of order N. (Input)
$\boldsymbol{E L O W}$ - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)
NEVAL - Number of eigenvalues found. (Output)
$\boldsymbol{E V A L}$ - Real vector of length MXEVAL containing the eigenvalues of $A$ in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are significant.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVBSF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL [, ...])

Specific: The specific interface names are S_EVBSF and D_EVBSF.

## FORTRAN 77 Interface

Single: CALL EVBSF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL)
Double: The double precision name is DEVBSF.

## Description

Routine EVBSF computes the eigenvalues in a given interval for a real symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. Then, an implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix. The reduction step is based on the EISPACK routine TRED1. See Smith et al. (1976). The rational QR algorithm is called the PWK algorithm. It is given in Parlett (1980, page 169).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E5BSF/DE5BSF. The reference is

CALL E5BSF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL, WK, IWK)
The additional arguments are as follows:
$\boldsymbol{W} \boldsymbol{K}$ - Work array of length 5 N .
IWK - Integer work array of length 1 N .
2. Informational error

Type Code
31 The number of eigenvalues in the specified interval exceeds MXEVAL. NEVAL contains the number of eigenvalues in the interval. No eigenvalues will be returned.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 56). The eigenvalues of $A$ are known to be $-1,5,5$ and 15 . The eigenvalues in the interval $[1.5,5.5]$ are computed and printed. As a test, this example uses MXEVAL $=4$. The routine EVBSF computes NEVAL, the number of eigenvalues in the given interval. The value of NEVAL is 2 .

```
USE EVBSF_INT
USE UMACH_INT
```

```
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, MXEVAL, N
PARAMETER (MXEVAL=4, N=4, LDA=N)
INTEGER NEVAL, NOUT
REAL A(LDA,N), EHIGH, ELOW, EVAL (MXEVAL)
    Set values of A
    A =( 6.0 4.0 4.0 1.0)
            ( 4.0 6.0 1.0 4.0)
            ( 4.0 1.0 6.0 4.0)
            (1.0 4.0 4.0 6.0)
DATA A/6.0, 4.0, 4.0, 1.0, 4.0, 6.0, 1.0, 4.0, 4.0, 1.0, 6.0, &
    4.0, 1.0, 4.0, 4.0, 6.0/
    Find eigenvalues of A
ELOW = 1.5
EHIGH = 5.5
CALL EVBSF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL)
    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,'(/,A,I2)') ' NEVAL = ', NEVAL
CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
END
```


## Output

```
NEVAL = 2
    EVAL
    1 2
    5.000 5.000
```


## EVFSF

Computes selected eigenvalues and eigenvectors of a real symmetric matrix.

## Required Arguments

MXEVAL - Maximum number of eigenvalues to be computed. (Input)
$\boldsymbol{A}$ - Real symmetric matrix of order N. (Input)
$\boldsymbol{E L O W}$ - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)

NEVAL - Number of eigenvalues found. (Output)
EVAL - Real vector of length MXEVAL containing the eigenvalues of $A$ in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are significant.
$\boldsymbol{E V E C}$ - Real matrix of dimension n by MXEVAL. (Output)
The $J$-th eigenvector corresponding to EVAL(J), is stored in the $J$-th column. Only the first NEVAL columns of EVEC are significant. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVFSF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL, EVEC [,...])

Specific: The specific interface names are S_EVFSF and D_EVFSF.

## FORTRAN 77 Interface

Single: CALL EVFSF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL, EVEC, LDEVEC)

Double: The double precision name is DEVFSF.

## Description

Routine EVFSF computes the eigenvalues in a given interval and the corresponding eigenvectors of a real symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. Then, an implicit rational QR algorithm is used to compute the eigenvalues of this tridiagonal matrix. Inverse iteration is used to compute the eigenvectors of the tridiagonal matrix. This is followed by orthogonalization of these vectors. The eigenvectors of the original matrix are computed by back transforming those of the tridiagonal matrix.

The reduction step is based on the EISPACK routine TRED1. The rational QR algorithm is called the PWK algorithm. It is given in Parlett (1980, page 169). The inverse iteration and orthogonalization processes are discussed in Hanson et al. (1990). The transformation back to the users's input matrix is based on the EISPACK routine TRBAK1. See Smith et al. (1976) for the EISPACK routines.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3FSF/DE3FSF. The reference is:
```
CALL E3FSF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, VAL, EVEC, LDEVEC, WK,
IWK)
```

The additional arguments are as follows:
WK - Work array of length 9 N .
IWK - Integer work array of length N .
2. Informational errors

Type Code
31 The number of eigenvalues in the specified range exceeds MXEVAL. NEVAL contains the number of eigenvalues in the range. No eigenvalues will be computed.
32 Inverse iteration did not converge. Eigenvector is not correct for the specified eigenvalue.
33 The eigenvectors have lost orthogonality.

## Example

In this example, $A$ is set to the computed Hilbert matrix. The eigenvalues in the interval $[0.001,1]$ and their corresponding eigenvectors are computed and printed. This example uses MXEVAL $=3$. The routine EVFSF computes the number of eigenvalues NEVAL in the given interval. The value of NEVAL is 2 . The performance index is also computed and printed. For more details, see IMSL routine EPISF.

```
USE EVFSF INT
USE EPISF}\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, MXEVAL, N, J, I
PARAMETER (MXEVAL=3, N=3, LDA=N, LDEVEC=N)
INTEGER NEVAL, NOUT
REAL A(LDA,N), EHIGH, ELOW, EVAL (MXEVAL), &
    EVEC (LDEVEC,MXEVAL), PI
```

```
!
                                    Compute Hilbert matrix
DO 20 J=1,N
        DO 10 I=1,N
                A(I,J) = 1.0/FLOAT(I+J-1)
    1 0
        CONTINUE
    CONTINUE
! Find eigenvalues and vectors
    ELOW = 0.001
    EHIGH = 1.0
    CALL EVFSF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL, EVEC, LDEVEC)
                            Compute performance index
    PI = EPISF(NEVAL,A,EVAL,EVEC)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,'(/,A,I2)') ' NEVAL = ', NEVAL
    CALL WRRRN ('EVAL', EVAL, 1, NEVAl, 1)
    CALL WRRRN ('EVEC', EVEC, N, NEVAL, LDEVEC)
    WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
    END
```


## Output

```
NEVAL = 2
        EVAL
        1 2
    0.1223 0.0027
            EVEC
1 -0.5474 -0.1277
0.5283 0.7137
3 0.6490 -0.6887
Performance index = 0.008
```


## EPISF

This function computes the performance index for a real symmetric eigensystem.

## Function Return Value

EPISF - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs on which the performance index computation is based on. (Input)
$\boldsymbol{A}-$ Symmetric matrix of order N . (Input)
$\boldsymbol{E V A L}$ - Vector of length NEVAL containing eigenvalues of A. (Input)
$\boldsymbol{E V E C}-\mathrm{N}$ by NEVAL array containing eigenvectors of A. (Input)
The eigenvector corresponding to the eigenvalue EVAL(J) must be in the $J$-th column of EVEC.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: EPISF (NEVAL, A, EVAL, EVEC [, ...])

Specific: The specific interface names are S_EPISF and D_EPISF.

## FORTRAN 77 Interface

Single: EPISF (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision function name is DEPISF.

## Description

Let $M=$ NEVAL, $\lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, \mathrm{~J})$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision, given by $\operatorname{AMACH}(4)$, see the Reference chapter of this manual. The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} x_{j}\right\|_{1}}{10 N \varepsilon\|A\|_{1}\left\|x_{j}\right\|_{1}}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Smith et al. (1976, pages 124-125).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E2ISF/DE2ISF. The reference is:
```
E2ISF (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC,WORK)
```

The additional argument is:
WORK - Work array of length N .

E2ISF - Performance Index.
2. Informational errors

Type Code
$31 \quad$ Performance index is greater than 100.
32 An eigenvector is zero.
33 The matrix is zero.

## Example

For an example of EPISF, see routine EVCSF.

## EVLSB

Computes all of the eigenvalues of a real symmetric matrix in band symmetric storage mode.

## Required Arguments

$\boldsymbol{A}$ - Band symmetric matrix of order N. (Input)

NCODA - Number of codiagonals in A. (Input)
$\boldsymbol{E V A L}$ - Vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: N = SIZE (A,2).
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLSB (A, NCODA, EVAL [,...])
Specific: The specific interface names are S_EVLSB and D_EVLSB.

## FORTRAN 77 Interface

Single: CALL EVLSB (N, A, LDA, NCODA, EVAL)
Double: The double precision name is DEVLSB.

## Description

Routine EVLSB computes the eigenvalues of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. The implicit QL algorithm is used to compute the eigenvalues of the resulting tridiagonal matrix.

The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1977). The QL routine is based on the EISPACK routine IMTQL1; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LSB/DE3LSB. The reference is:

CALL E3LSB (N, A, LDA, NCODA, EVAL, ACOPY, WK)
The additional arguments are as follows:

ACOPY - Work array of length $\mathrm{N}(\mathrm{NCODA}+1)$. The arrays A and ACOPY may be the same, in which case the first $\mathrm{N}(\mathrm{NCODA}+1)$ elements of A will be destroyed.

WK - Work array of length N .
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 77). The eigenvalues of this matrix are given by

$$
\lambda_{k}=\left(1-2 \cos \frac{k \pi}{N+1}\right)^{2}-3
$$

Since the eigenvalues returned by EVLSB are in decreasing magnitude, the above formula for $k=1, \ldots, N$ gives the values in a different order. The eigenvalues of this real band symmetric matrix are computed and printed.

```
USE EVLSB_INT
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDEVEC, N, NCODA
```

```
PARAMETER (N=5, NCODA=2, LDA=NCODA+1, LDEVEC=N)
REAL A(LDA,N), EVAL (N)
Define values of A:
        A = (\begin{array}{llll}{-1}&{2}&{1}&{)}\end{array})=[\begin{array}{l}{2}\end{array})
            (\begin{array}{llll}{2}&{0}&{2}&{1}\end{array})
            (\begin{array}{lllll}{1}&{2}&{0}&{2}&{1}\end{array})
            (}\begin{array}{llllll}{1}&{2}&{0}&{2}\end{array}
            (
                Represented in band symmetric
                form this is:
                A =(\begin{array}{llllll}{0}&{0}&{1}&{1}&{1}\end{array})
                    ([\begin{array}{lllll}{0}&{2}&{2}&{2}&{2}\end{array})
                    (-1
    DATA A/0.0, 0.0, -1.0, 0.0, 2.0, 0.0, 1.0, 2.0, 0.0, 1.0, 2.0, &
        0.0, 1.0, 2.0, -1.0/
    CALL EVLSB (A, NCODA, EVAL)
        Print results
    CALL WRRRN ('EVAL', EVAL, 1, N, 1)
    END
```


## Output

| 1 | 2 | EVAL |  |  |
| ---: | ---: | ---: | ---: | ---: |
| 4.464 | -3.000 | -2.464 | -2.000 | 1.000 |

## EVCSB

Computes all of the eigenvalues and eigenvectors of a real symmetric matrix in band symmetric storage mode.

## Required Arguments

A - Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
$\boldsymbol{E V A L}$ - Vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Matrix of order N containing the eigenvectors. (Output)
The $J$-th eigenvector, corresponding to EVAL(J), is stored in the $J$-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input) Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=$ SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL EVCSB (A, NCODA, EVAL, EVEC [,...])
Specific: The specific interface names are S_EVCSB and D_EVCSB.

## FORTRAN 77 Interface

Single: CALL EVCSB (N, A, LDA, NCODA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCSB.

## Description

Routine EVCSB computes the eigenvalues and eigenvectors of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. These transformations are accumulated. The implicit QL algorithm is used to compute the eigenvalues and eigenvectors of the resulting tridiagonal matrix.

The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1977). The QL routine is based on the EISPACK routine IMTQL2; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E4CSB/DE4CSB. The reference is:
```
CALL E4CSB (N, A, LDA, NCODA, EVAL, EVEC, LDEVEC, COPY,WK, IWK)
```

The additional arguments are as follows:
ACOPY - Work array of length $\mathrm{N}(\mathrm{NCODA}+1)$. A and ACOPY may be the same, in which case the first N * NCODA elements of A will be destroyed.
$\boldsymbol{W} \boldsymbol{K}$ - Work array of length N.

IWK - Integer work array of length N.
2. Informational error

Type Code

1 The iteration for the eigenvalues failed to converge.
3. The success of this routine can be checked using EPISB.

## Example

In this example, a DATA statement is used to set $A$ to a band matrix given by Gregory and Karney (1969, page 75 ). The eigenvalues, $\lambda_{k}$, of this matrix are given by

$$
\lambda_{k}=16 \sin ^{4}\left(\frac{k \pi}{2 N+2}\right)
$$

The eigenvalues and eigenvectors of this real band symmetric matrix are computed and printed. The performance index is also computed and printed. This serves as a check on the computations, for more details, see IMSL routine EPISB.

```
USE EVCSB INT
USE EPISB_-INT
USE UMACH_INT
USE WRRRN__INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, N, NCODA
PARAMETER (N=6, NCODA=2, LDA=NCODA+1, LDEVEC=N)
INTEGER NOUT
REAL A(LDA,N), EVAL (N), EVEC(LDEVEC,N), PI
                                    Define values of A:
                                    A =(\begin{array}{llll}{5}&{-4}&{1}\end{array})
                                    ( -4
                                    (
                                    Represented in band symmetric
                                    form this is:
                                    A =(\begin{array}{llllllll}{0}&{0}&{1}&{1}&{1}&{1}\end{array})
                                    ( ( 0 %rrrrre
DATA A/0.0, 0.0, 5.0, 0.0, -4.0, 6.0, 1.0, -4.0, 6.0, 1.0, -4.0, &
        6.0, 1.0, -4.0, 6.0, 1.0, -4.0, 5.0/
            Find eigenvalues and vectors
CALL EVCSB (A, NCODA, EVAL, EVEC)
                            Compute performance index
PI = EPISB(N,A,NCODA,EVAL,EVEC)
    Print results
CALL UMACH (2, NOUT)
CALL WRRRN ('EVAL', EVAL, 1, N, 1)
CALL WRRRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output



## EVASB

Computes the largest or smallest eigenvalues of a real symmetric matrix in band symmetric storage mode.

## Required Arguments

NEVAL - Number of eigenvalues to be computed. (Input)
A-Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
SMALL - Logical variable. (Input)
If .TRUE., the smallest NEVAL eigenvalues are computed. If .FALSE., the largest NEVAL eigenvalues are computed.

EVAL - Vector of length NEVAL containing the computed eigenvalues in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVASB (NEVAL, A, NCODA, SMALL, EVAL [, ...])

Specific: The specific interface names are S_EVASB and D_EVASB.

## FORTRAN 77 Interface

Single: CALL EVASB (N, NEVAL, A, LDA, NCODA, SMALL, EVAL)
Double: The double precision name is DEVASB.

## Description

Routine EVASB computes the largest or smallest eigenvalues of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. The rational QR algorithm with Newton corrections is used to compute the extreme eigenvalues of this tridiagonal matrix.

The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1978). The QR routine is based on the EISPACK routine RATQR; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3ASB/DE3ASB. The reference is:

CALL E3ASB (N, NEVAL, A, LDA, NCODA, SMALL, EVAL, ACOPY, WK)
The additional arguments are as follows:
ACOPY - Work array of length $\mathrm{N}(\mathrm{NCODA}+1)$. A and ACOPY may be the same, in which case the first $\mathrm{N}(\mathrm{NCODA}+1)$ elements of A will be destroyed.
$W K-$ Work array of length 3 N .
2. Informational error

Type Code
31 The iteration for an eigenvalue failed to converge. The best estimate will be returned.

## Example

The following example is given in Gregory and Karney (1969, page 63). The smallest four eigenvalues of the matrix

$$
A=\left[\begin{array}{lllllllllll}
5 & 2 & 1 & 1 & & & & & & & \\
2 & 6 & 3 & 1 & 1 & & & & & & \\
1 & 3 & 6 & 3 & 1 & 1 & & & & & \\
1 & 1 & 3 & 6 & 3 & 1 & 1 & & & & \\
& 1 & 1 & 3 & 6 & 3 & 1 & 1 & & & \\
& & 1 & 1 & 3 & 6 & 3 & 1 & 1 & & \\
& & & 1 & 1 & 3 & 6 & 3 & 1 & 1 & \\
& & & & 1 & 1 & 3 & 6 & 3 & 1 & 1 \\
& & & & & 1 & 1 & 3 & 6 & 3 & 1 \\
& & & & & & 1 & 1 & 3 & 6 & 2 \\
& & & & & & & 1 & 1 & 2 & 5
\end{array}\right]
$$

are computed and printed.

```
USE EVASB_INT
USE WRRRN_INT
USE SSET_\overline{INT}
IMPLICIT NONE
INTEGER LDA, N, NCODA, NEVAL
PARAMETER (N=11, NCODA=3, NEVAL=4, LDA=NCODA+1)
REAL A (LDA,N), EVAL (NEVAL)
LOGICAL SMALL
                                    Set up matrix in band symmetric
                                    storage mode
CALL SSET (N, 6.0, A(4:,1), LDA)
CALL SSET (N-1, 3.0, A(3:,2), LDA)
CALL SSET (N-2, 1.0, A(2:,3), LDA)
CALL SSET (N-3, 1.0, A(1:,4), LDA)
CALL SSET (NCODA, 0.0, A(1:,1), 1)
CALL SSET (NCODA-1, 0.0, A (1:,2), 1)
CALL SSET (NCODA-2, 0.0, A (1:, 3), 1)
A (4,1) = 5.0
A(4,N)}=5.
A (3,2) = 2.0
A(3,N)=2.0
SMALL = .TRUE.
CALL EVASB (NEVAL, A, NCODA, SMALL, EVAL)
                                    Print results
CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
END
```


## Output

|  | EVAL |  |  |
| ---: | ---: | ---: | ---: |
| 1 | 2 | 3 | 4 |
| 4.000 | 3.172 | 1.804 | 0.522 |

## EVESB

Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix in band symmetric storage mode.

## Required Arguments

NEVEC - Number of eigenvectors to be calculated. (Input)
A-Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
SMALL - Logical variable. (Input)
If .TRUE., the smallest NEVEC eigenvectors are computed. If .FALSE., the largest NEVEC eigenvectors are computed.
$\boldsymbol{E V A L}$ - Vector of length NEVEC containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Real matrix of dimension $n$ by NEVEC. (Output)
The J-th eigenvector, corresponding to $\operatorname{EVAL}(J)$, is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL EVESB (NEVEC, A, NCODA, SMALL, EVAL, EVEC [,...])
Specific: The specific interface names are S_EVESB and D_EVESB.

## FORTRAN 77 Interface

Single: CALL EVESB (N, NEVEC, A, LDA, NCODA, SMALL, EVAL, EVEC, LDEVEC)

Double: The double precision name is DEVESB.

## Description

Routine EVESB computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. The rational QR algorithm with Newton corrections is used to compute the extreme eigenvalues of this tridiagonal matrix. Inverse iteration and orthogonalization are used to compute the eigenvectors of the given band matrix. The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1977). The QR routine is based on the EISPACK routine RATQR; see Smith et al. (1976). The inverse iteration and orthogonalization steps are based on EISPACK routine BANDV using the additional steps given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E4ESB/DE4ESB. The reference is:
```
CALL E4ESB (N,NEVEC, A, LDA, NCODA, SMALL, EVAL, EVEC, LDEVEC, ACOPY,WK,
IWK)
```

The additional argument is:
ACOPY - Work array of length N (NCODA +1 ).
$W K-$ Work array of length $N(2 N C O D A+5)$.
IWK - Integer work array of length N .
2. Informational errors

Type Code
31 Inverse iteration did not converge. Eigenvector is not correct for the specified eigenvalue.
32 The eigenvectors have lost orthogonality.
3. The success of this routine can be checked using EPISB.

## Example

The following example is given in Gregory and Karney (1969, page 75). The largest three eigenvalues and the corresponding eigenvectors of the matrix are computed and printed.

```
USE EVESB INT
USE EPISB INT
USE UMACH_INT
USE WRRRN_`_INT
IMPLICIT NONE
```

```
!
    INTEGER LDA, LDEVEC, N, NCODA, NEVEC
    PARAMETER (N=6, NCODA=2, NEVEC=3, LDA=NCODA+1, LDEVEC=N)
    INTEGER NOUT
    REAL A (LDA,N), EVAL (NEVEC), EVEC(LDEVEC,NEVEC), PI
    LOGICAL SMALL
        Define values of A:
        A = (\begin{array}{llll}{5}&{-4}&{1}\end{array})
            ( }\begin{array}{r}{-4}\\{}\\{(1}
            (
    Represented in band symmetric
    form this is:
    A = (\begin{array}{llllllll}{0}&{0}&{1}&{1}&{1}&{1}\end{array})
    ( (rrrrrra
    DATA A/0.0, 0.0, 5.0, 0.0, -4.0, 6.0, 1.0, -4.0, 6.0, 1.0, -4.0, &
    6.0, 1.0, -4.0, 6.0, 1.0, -4.0, 5.0/
    Find the 3 largest eigenvalues
    and their eigenvectors.
    SMALL = .FALSE
    CALL EVESB (NEVEC, A, NCODA, SMALL, EVAL, EVEC)
        Compute performance index
    PI = EPISB(NEVEC,A,NCODA,EVAL,EVEC)
        Print results
    CALL UMACH (2, NOUT)
    CALL WRRRN ('EVAL', EVAL, 1, NEVEC, 1)
    CALL WRRRN ('EVEC', EVEC)
    WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
    END
```


## Output

| EVAL |  |  |  |
| :---: | :---: | :---: | :---: |
| 14.45 |  | 2 | 3 |
|  |  | 45.98 |  |
| EVEC |  |  |  |
|  | 1 | 2 | 3 |
| 1 | 0.2319 | -0.4179 | 0.5211 |
| 2 | -0.4179 | 0.5211 | -0.2319 |
| 3 | 0.5211 | -0.2319 | -0.4179 |
| 4 | -0.5211 | -0.2319 | 0.4179 |
| 5 | 0.4179 | 0.5211 | 0.2319 |
| 6 | -0.2319 | -0.4179 | -0.5211 |
| Performance index $=0.175$ |  |  |  |

## EVBSB

Computes the eigenvalues in a given interval of a real symmetric matrix stored in band symmetric storage mode.

## Required Arguments

MXEVAL - Maximum number of eigenvalues to be computed. (Input)
A-Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
$\boldsymbol{E L O W}$ - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)
NEVAL - Number of eigenvalues found. (Output)
$\boldsymbol{E V A L}$ - Real vector of length MXEVAL containing the eigenvalues of A in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are set.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVBSB (MXEVAL, A, NCODA, ELOW, EHIGH, NEVAL, EVAL [, ...])
Specific: The specific interface names are S_EVBSB and D_EVBSB.

## FORTRAN 77 Interface

Single: CALL EVBSB (N, MXEVAL, A, LDA, NCODA, ELOW, EHIGH, NEVAL, EVAL)
Double: The double precision name is DEVBSB.

## Description

Routine EVBSB computes the eigenvalues in a given range of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. A bisection algorithm is used to compute the eigenvalues of the tridiagonal matrix in a given range.

The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1977). The bisection routine is based on the EISPACK routine BISECT; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3BSB/DE3BSB. The reference is:
```
CALL E3BSB (N, MXEVAL, A, LDA, NCODA, ELOW, EHIGH, NEVAL, EVAL, ACOPY, WK)
```

The additional arguments are as follows:
ACOPY - Work matrix of size NCODA +1 by n. A and ACOPY may be the same, in which case the first $\mathrm{N}(\mathrm{NCODA}+1)$ elements of A will be destroyed.

WK - Work array of length 5 N
2. Informational error

Type Code
31 The number of eigenvalues in the specified interval exceeds MXEVAL. NEVAL contains the number of eigenvalues in the interval. No eigenvalues will be returned.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 77 ). The eigenvalues in the range $(-2.5,1.5)$ are computed and printed. As a test, this example uses MXEVAL $=5$. The routine EVBSB computes NEVAL, the number of eigenvalues in the given range, has the value 3 .

```
USE EVBSB INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRN_INT
```

IMPLICIT NONE
! Declare variables
INTEGER LDA, MXEVAL, N, NCODA
PARAMETER (MXEVAL=5, $\mathrm{N}=5, \mathrm{NCODA}=2, ~ L D A=N C O D A+1$ )
!
INTEGER NEVAL, NOUT
REAL A (LDA, N), EHIGH, ELOW, EVAL (MXEVAL)
Define values of $A$ :
$A=\left(\begin{array}{rrrrl}-1 & 2 & 1 & & ) \\ (2 & 0 & 2 & 1\end{array}\right)$

```
!
!
!
!
!
!
```



```
DATA A/0.0, 0.0, -1.0, 0.0, 2.0, 0.0, 1.0, 2.0, 0.0, 1.0, 2.0, &
    0.0, 1.0, 2.0, -1.0/
!
ELOW = -2.5
EHIGH = 1.5
CALL EVBSB (MXEVAL, A, NCODA, ELOW, EHIGH, NEVAL, EVAL)
                                    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,'(/,A,I1)') ' NEVAL = ', NEVAL
CALL WRRRN ('EVAL', EVAL, 1, NEVAl, 1)
END
```


## Output

NEVAL $=3$

| EVAL |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -2.464 | -2.000 | 1.000 |

## EVFSB

Computes the eigenvalues in a given interval and the corresponding eigenvectors of a real symmetric matrix stored in band symmetric storage mode.

## Required Arguments

MXEVAL - Maximum number of eigenvalues to be computed. (Input)
A - Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
$\boldsymbol{E L O W}$ - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)
NEVAL - Number of eigenvalues found. (Output)
EVAL - Real vector of length MXEVAL containing the eigenvalues of A in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are significant.

EVEC - Real matrix containing in its first NEVAL columns the eigenvectors associated with the eigenvalues found and stored in EVAL. Eigenvector $J$ corresponds to eigenvalue $J$ for $J=1$ to NEVAL. Each vector is normalized to have Euclidean length equal to the value one. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: N $=\operatorname{SIZE}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL EVFSB (MXEVEL, A, NCODA, ELOW, EHIGH, NEVAL, EVAL, EVEC [, ...])
Specific: The specific interface names are S_EVFSB and D_EVFSB.

## FORTRAN 77 Interface

Single: CALL EVFSB (N, MXEVAL, A, LDA, NCODA, ELOW, EHIGH, NEVAL, EVAL, EVEC, LDEVEC)

Double: The double precision name is DEVFSB.

## Description

Routine EVFSB computes the eigenvalues in a given range and the corresponding eigenvectors of a real band symmetric matrix. Orthogonal similarity transformations are used to reduce the matrix to an equivalent tridiagonal matrix. A bisection algorithm is used to compute the eigenvalues of the tridiagonal matrix in the required range. Inverse iteration and orthogonalization are used to compute the eigenvectors of the given band symmetric matrix.

The reduction routine is based on the EISPACK routine BANDR; see Garbow et al. (1977). The bisection routine is based on the EISPACK routine BISECT; see Smith et al. (1976). The inverse iteration and orthogonalization steps are based on the EISPACK routine BANDV using remarks from Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3FSB/DE3FSB. The reference is:
```
CALL E3FSB (N, MXEVAL, A, LDA, NCODA, ELOW, EHIGH, NEVAL, EVAL, EVEC,
LDEVEC, ACOPY, WK1, WK2, IWK)
```

The additional arguments are as follows:
ACOPY - Work matrix of size NCODA +1 by N.
WK1 - Work array of length 6 N .
WK2 - Work array of length $2 \mathrm{~N} *$ NCODA +N
IWK - Integer work array of length N .
2. Informational errors

Type Code
31 The number of eigenvalues in the specified interval exceeds MXEVAL. NEVAL contains the number of eigenvalues in the interval. No eigenvalues will be returned.
32 Inverse iteration did not converge. Eigenvector is not correct for the specified eigenvalue.
33 The eigenvectors have lost orthogonality.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 75 ). The eigenvalues in the range [1, 6] and their corresponding eigenvectors are computed and printed. As a test, this example uses MXEVAL $=4$. The routine EVFSB computes NEVAL, the number of eigenvalues in the given range has the value 2 . As a check on the computations, the performance index is also computed and printed. For more details, see IMSL routine EPISB.

```
USE EVFSB_INT
USE EPISB INT
USE WRRRN - INT
USE UMACH_INT
IMPLICIT NONE
INTEGER LDA, LDEVEC, MXEVAL, N, NCODA
PARAMETER (MXEVAL=4, N=6, NCODA=2, LDA=NCODA+1, LDEVEC=N)
INTEGER NEVAL, NOUT
REAL A(LDA,N), EHIGH, ELOW, EVAL (MXEVAL), &
    EVEC(LDEVEC,MXEVAL), PI
                                    Define values of A:
                                    A =( }\begin{array}{llll}{5}&{-4}&{1}\end{array}
                                    ( llll}-4 6 -4 1 )
                            ( 1rrrrrl
                            (
                            (
Represented in band symmetric
form this is:
```

!


```
DATA A/0.0, 0.0, 5.0, 0.0, -4.0, 6.0, 1.0, -4.0, 6.0, 1.0, -4.0, &
    6.0, 1.0, -4.0, 6.0, 1.0, -4.0, 5.0/
    Find eigenvalues and vectors
ELOW = 1.0
EHIGH = 6.0
CALL EVFSB (MXEVAL, A, NCODA, ELOW, EHIGH, NEVAL, EVAL, EVEC)
    Compute performance index
PI = EPISB(NEVAL,A,NCODA,EVAL,EVEC)
    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,'(/,A,I1)') ' NEVAL = ', NEVAL
CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
CALL WRRRN ('EVEC', EVEC, N, NEVAL, LDEVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output

```
NEVAL = 2
        EVAL
        1 2
    5.978 2.418
            EVEC
    10.5211 0.5211
    -0.2319 0.2319
    3-0.4179 -0.4179
    4 0.4179 -0.4179
    5 0.2319 0.2319
    6 -0.5211 0.5211
    Performance index = 0.083
```


## EPISB

This function computes the performance index for a real symmetric eigensystem in band symmetric storage mode.

## Required Arguments

EPISB - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs on which the performance is based. (Input)
$\boldsymbol{A}-$ Band symmetric matrix of order N. (Input)
NCODA - Number of codiagonals in A. (Input)
$\boldsymbol{E V A L}$ - Vector of length NEVAL containing eigenvalues of A. (Input)
$\boldsymbol{E V E C}$ - N by NEVAL array containing eigenvectors of A. (Input)
The eigenvector corresponding to the eigenvalue EVAL(J) must be in the $J$-th column of EVEC.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: EPISB (NEVAL, A, NCODA, EVAL, EVEC [, ...])

Specific: The specific interface names are S_EPISB and D_EPISB.

## FORTRAN 77 Interface

Single: EPISB (N, NEVAL, A, LDA, NCODA, EVAL, EVEC, LDEVEC)
Double: The double precision function name is DEPISB.

## Description

Let $M=\operatorname{NEVAL}, \lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, \mathrm{~J})$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision, given by AMACH(4), see the Reference chapter of the manual. The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} x_{j}\right\|_{1}}{10 N \varepsilon\|A\|_{1}\left\|x_{j}\right\|_{1}}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first
developed by the EISPACK project at Argonne National Laboratory; see Smith et al. (1976, pages 124-125).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E2ISB/DE2ISB. The reference is:
```
E2ISB (N, NEVAL, A, LDA, NCODA, EVAL, EVEC, LDEVEC, WK)
```

The additional argument is:
$\boldsymbol{W K}$ - Work array of length N.
2. Informational errors

Type Code
$31 \quad$ Performance index is greater than 100.
32 An eigenvector is zero.
33 The matrix is zero.

## Example

For an example of EPISB, see IMSL routine EVCSB.

## EVLHF

Computes all of the eigenvalues of a complex Hermitian matrix.

## Required Arguments

$\boldsymbol{A}$ - Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.
$\boldsymbol{E V A L}$ - Real vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

FORTRAN 90 Interface
Generic: CALL EVLHF (A, EVAL [,...])

Specific: The specific interface names are S_EVLHF and D_EVLHF.

## FORTRAN 77 Interface

Single: CALL EVLHF (N, A, LDA, EVAL)
Double: The double precision name is DEVLHF.

## Description

Routine EVLHF computes the eigenvalues of a complex Hermitian matrix. Unitary similarity transformations are used to reduce the matrix to an equivalent real symmetric tridiagonal matrix. The implicit QL algorithm is used to compute the eigenvalues of this tridiagonal matrix.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LHF/DE3LHF. The reference is:

CALL E3LHF (N, A, LDA, EVAL, ACOPY, RWK, CWK, IWK)
The additional arguments are as follows:

ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same in which case A will be destroyed.
$\boldsymbol{R W K}$ - Work array of length N .

CWK - Complex work array of length 2 N .

IWK - Integer work array of length N.
2. Informational errors

Type Code
31 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
$4 \quad 1 \quad$ The iteration for an eigenvalue failed to converge.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E3LHF, the internal or working leading dimensions of ACOPY and ECOPY are both increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced
by IVAL(1) and IVAL(2), respectively, in routine EVLHF. Additional memory allocation and option value restoration are automatically done in EVLHF. There is no requirement that users change existing applications that use EVLHF or E3LHF. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$. Items $5-8$ in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVLHF.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 114). The eigenvalues of this complex Hermitian matrix are computed and printed.

```
USE EVLHF_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (N=2, LDA=N)
REAL EVAL (N)
COMPLEX A(LDA,N)
    Set values of A
    A=(\begin{array}{lll}{1}&{-i}\end{array})
    ( i 1 )
    DATA A/ (1.0,0.0), (0.0,1.0),(0.0,-1.0), (1.0,0.0)/
    CALL EVLHF (A, EVAL)
    CALL WRRRN ('EVAL', EVAL, 1, N, 1)
    END
```


## Output

```
    EVAL
```

    \(1 \quad 2\)
    
## EVCHF

Computes all of the eigenvalues and eigenvectors of a complex Hermitian matrix.

## Required Arguments

A-Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.

EVAL - Real vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Complex matrix of order n. (Output)
The $J$-th eigenvector, corresponding to EVAL(J), is stored in the $J$-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=$ SIZE (EVEC, 1 ).

## FORTRAN 90 Interface

Generic: CALL EVCHF (A, EVAL, EVEC [,...])

Specific: The specific interface names are S_EVCHF and D_EVCHF.

## FORTRAN 77 Interface

Single: CALL EVCHF (N, A, LDA, EVAL, EVEC, LDEVEC)

Double: The double precision name is DEVCHF.

## Description

Routine EVCHF computes the eigenvalues and eigenvectors of a complex Hermitian matrix.
Unitary similarity transformations are used to reduce the matrix to an equivalent real symmetric tridiagonal matrix. The implicit QL algorithm is used to compute the eigenvalues and eigenvectors of this tridiagonal matrix. These eigenvectors and the transformations used to reduce the matrix to tridiagonal form are combined to obtain the eigenvectors for the user's problem. The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E5CHF/DE5CHF. The reference is:
```
CALL E5CHF (N, A, LDA, EVAL, EVEC, LDEVEC, ACOPY, RWK, CWK, IWK)
```

The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same, in which case A will be destroyed.
$\boldsymbol{R} \boldsymbol{W} \boldsymbol{K}$ - Work array of length $\mathrm{N}^{2}+\mathrm{N}$.
CWK - Complex work array of length 2 N .
IWK - Integer work array of length N.
2. Informational error

Type Code
31 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
41 The iteration for an eigenvalue failed to converge.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.
3. The success of this routine can be checked using EPIHF.
4. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine E 5 CHF , the internal or working leading dimensions of ACOPY and ECOPY are both increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values $\operatorname{IVAL}(3)$ and $\operatorname{IVAL}(4)$ are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine EVCHF. Additional memory allocation and option value restoration are automatically done in EVCHF. There is no requirement that users change existing applications that use EVCHF or E5CHF. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16,0,1$. Items $5-8$ in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in EVCHF.

## Example

In this example, a DATA statement is used to set $A$ to a complex Hermitian matrix. The eigenvalues and eigenvectors of this matrix are computed and printed. The performance index is also computed and printed. This serves as a check on the computations, for more details, see routine EPIHF.

USE IMSL_libraries
IMPLICIT NONE

INTEGER LDA, LDEVEC, N
PARAMETER $(N=3, L D A=N, L D E V E C=N)$

```
!
    INTEGER NOUT
    REAL EVAL (N), PI
    COMPLEX A(LDA,N), EVEC(LDEVEC,N)
    Set values of A
    A = ((1, 0) ( 1,-7i) ( 0,- i))
        ((1,7i) ( 5, 0) (10,-3i))
        ((0, i) ( 10, 3i) (-2, 0))
DATA A/ (1.0,0.0), (1.0,7.0), (0.0,1.0), (1.0,-7.0), (5.0,0.0), &
    (10.0, 3.0), (0.0,-1.0), (10.0,-3.0), (-2.0,0.0)/
    Find eigenvalues and vectors of A
CALL EVCHF (A, EVAL, EVEC)
PI = EPIHF(N,A,EVAL,EVEC)
CALL UMACH (2, NOUT)
CALL WRRRN ('EVAL', EVAL, 1, N, 1)
CALL WRCRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output

            EVAL
    | 1 | 2 | 3 |
| ---: | ---: | ---: |
| 15.38 | -10.63 | -0.75 |

EVEC

|  | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: |
| 1 | ( 0.0631,-0.4075) | (-0.0598, -0.3117) | ( 0.8539, 0.0000) |
| 2 | ( 0.7703, 0.0000) | (-0.5939, 0.1841) | (-0.0313,-0.1380) |
| 3 | ( 0.4668, 0.1366) | ( 0.7160, 0.0000) | ( 0.0808,-0.4942) |

## EVAHF

Computes the largest or smallest eigenvalues of a complex Hermitian matrix.

## Required Arguments

NEVAL - Number of eigenvalues to be calculated. (Input)
A - Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.
SMALL - Logical variable. (Input)
If . TRUE., the smallest NEVAL eigenvalues are computed. If . FALSE., the largest NEVAL eigenvalues are computed.
$\boldsymbol{E V A L}$ - Real vector of length N containing the extreme eigenvalues of A in decreasing order of magnitude in the first NEVAL elements. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVAHF (NEVAL, A, SMALL, EVAL [,...])

Specific: The specific interface names are S_EVAHF and D_EVAHF.

## FORTRAN 77 Interface

Single: CALL EVAHF (N, NEVAL, A, LDA, SMALL, EVAL)
Double: The double precision name is DEVAHF.

## Description

Routine EVAHF computes the largest or smallest eigenvalues of a complex Hermitian matrix. Unitary transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. The rational QR algorithm with Newton corrections is used to compute the extreme eigenvalues of this tridiagonal matrix.

The reduction routine is based on the EISPACK routine HTRIDI. The QR routine is based on the EISPACK routine RATQR. See Smith et al. (1976) for the EISPACK routines.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3AHF/DE3AHF. The reference is

CALL E3AHF (N, NEVAL, A, LDA, SMALL, EVAL, ACOPY, RWK, CWK, IWK)
The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same in which case A will be destroyed.
$\boldsymbol{R W K}$ - Work array of length 2 N .
$\boldsymbol{C W K}$ - Complex work array of length 2 N .
IWK - Work array of length N .
2. Informational errors

Type Code
31 The iteration for an eigenvalue failed to converge. The best estimate will be returned.
32 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 114). Its largest eigenvalue is computed and printed.

```
USE EVAHF_INT
USE WRRRN_-INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (N=2, LDA=N)
INTEGER NEVAL
REAL EVAL (N)
COMPLEX A(LDA,N)
LOGICAL SMALL
DATA A/ (1.0,0.0), (0.0,1.0), (0.0,-1.0), (1.0,0.0)/
Find the largest eigenvalue of A
NEVAL = 1
SMALL = .FALSE.
CALL EVAHF (NEVAL, A, SMALL, EVAL)
    Print results
CALL WRRRN ('EVAL', EVAL, 1, NEVAl, 1)
END
```


## Output

EVAL
2.000

## EVEHF

Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a complex Hermitian matrix.

## Required Arguments

NEVEC - Number of eigenvectors to be computed. (Input)
$\boldsymbol{A}$ - Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.
SMALL - Logical variable. (Input)
If .TRUE., the smallest NEVEC eigenvectors are computed. If .FALSE., the largest NEVEC eigenvectors are computed.

EVAL - Real vector of length NEVEC containing the eigenvalues of A in decreasing order of magnitude. (Output)

EVEC - Complex matrix of dimension n by NEVEC. (Output)
The J-th eigenvector corresponding to EVAL(J), is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVEHF (NEVEC, A, SMALL, EVAL, EVEC [, ...])

Specific: The specific interface names are S_EVEHF and D_EVEHF.

## FORTRAN 77 Interface

Single: CALL EVEHF (N, NEVEC, A, LDA, SMALL, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVEHF.

## Description

Routine EVEHF computes the largest or smallest eigenvalues and the corresponding eigenvectors of a complex Hermitian matrix. Unitary transformations are used to reduce the matrix to an equivalent real symmetric tridiagonal matrix. The rational QR algorithm with Newton corrections is used to compute the extreme eigenvalues of the tridiagonal matrix. Inverse iteration is used to compute the eigenvectors of the tridiagonal matrix. Eigenvectors of the original matrix are found by back transforming the eigenvectors of the tridiagonal matrix.

The reduction routine is based on the EISPACK routine HTRIDI. The QR routine used is based on the EISPACK routine RATQR. The inverse iteration routine is based on the EISPACK routine TINVIT. The back transformation routine is based on the EISPACK routine HTRIBK. See Smith et al. (1976) for the EISPACK routines.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3EHF/DE3EHF. The reference is:
```
CALL E3EHF (N, NEVEC, A, LDA, SMALL, EVAL, EVEC, LDEVEC, ACOPY, RW1, RW2,
CWK,IWK)
```

The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same, in which case A will be destroyed.
$\boldsymbol{R W 1}$ - Work array of length N * NEVEC. Used to store the real eigenvectors of a symmetric tridiagonal matrix.
$\boldsymbol{R W} \mathbf{2}$ - Work array of length 8N.
CWK - Complex work array of length 2 N .
IWK - Work array of length n.
2. Informational errors

Type Code
31 The iteration for an eigenvalue failed to converge. The best estimate will be returned.
32 The iteration for an eigenvector failed to converge. The eigenvector will be set to 0 . imaginary part.
3. The success of this routine can be checked using EPIHF.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 115). The smallest eigenvalue and its corresponding eigenvector is computed and printed.
The performance index is also computed and printed. This serves as a check on the computations.
For more details, see IMSL routine EPIHF.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDEVEC, N, NEVEC
PARAMETER (N=3, NEVEC=1, LDA=N, LDEVEC=N)
INTEGER NOUT
REAL EVAL (N), PI
COMPLEX A(LDA,N), EVEC(LDEVEC,NEVEC)
LOGICAL SMALL
    Set values of A
    A = (\begin{array}{llll}{2}&{-i}&{0}\\{i}&{i}&{2}&{0}\end{array})
DATA A/ (2.0,0.0), (0.0,1.0), (0.0,0.0), (0.0,-1.0), (2.0,0.0), &
    (0.0,0.0), (0.0,0.0), (0.0,0.0), (3.0,0.0)/
    Find smallest eigenvalue and its
    eigenvectors
SMALL = .TRUE.
CALL EVEHF (NEVEC, A, SMALL, EVAL, EVEC)
    Compute performance index
PI = EPIHF(NEVEC,A,EVAL,EVEC)
CALL UMACH (2, NOUT)
CALL WRRRN ('EVAL', EVAL, 1, NEVEC, 1)
CALL WRCRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```

$!$

## Output

```
EVAL
```

1.000
EVEC
$1 \quad(0.0000,0.7071)$

```
2 ( 0.7071, 0.0000)
3 (0.0000, 0.0000)
Performance index = 0.031
```


## EVBHF

Computes the eigenvalues in a given range of a complex Hermitian matrix.

## Required Arguments

MXEVAL — Maximum number of eigenvalues to be computed. (Input)
$\boldsymbol{A}$ - Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.
$\boldsymbol{E L O W}$ - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)
NEVAL - Number of eigenvalues found. (Output)
$\boldsymbol{E V A L}$ - Real vector of length MXEVAL containing the eigenvalues of A in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are significant.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVBHF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL [, ...])
Specific: The specific interface names are S_EVBHF and D_EVBHF.

## FORTRAN 77 Interface

Single: CALL EVBHF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL)
Double: The double precision name is DEVBHF.

## Description

Routine EVBHF computes the eigenvalues in a given range of a complex Hermitian matrix. Unitary transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. A bisection algorithm is used to compute the eigenvalues in the given range of this tridiagonal matrix.

The reduction routine is based on the EISPACK routine HTRIDI. The bisection routine used is based on the EISPACK routine BISECT. See Smith et al. (1976) for the EISPACK routines.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3BHF/DE3BHF. The reference is:

CALL E3BHF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL, ACOPY, RWK, CWK, IWK)

The additional arguments are as follows:
ACOPY - Complex work matrix of size n by N. A and ACOPY may be the same, in which case the first $N^{2}$ elements of $A$ will be destroyed.
$\boldsymbol{R W K}$ - Work array of length 5 N .
CWK — Complex work array of length 2 N .
IWK - Work array of length MXEVAL.
2. Informational errors

Type Code
31 The number of eigenvalues in the specified range exceeds MXEVAL. NEVAL contains the number of eigenvalues in the range. No eigenvalues will be computed.
32 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.

## Example

In this example, a DATA statement is used to set $A$ to a matrix given by Gregory and Karney (1969, page 114). The eigenvalues in the range [1.5, 2.5] are computed and printed. This example allows a maximum number of eigenvalues MXEVAL $=2$. The routine computes that there is one eigenvalue in the given range. This value is returned in NEVAL.

```
USE EVBHF_INT
USE UMACH_INT
USE WRRRN_INT
```

```
IMPLICIT NONE
! Declare variables
    INTEGER LDA, MXEVAL, N
    PARAMETER (MXEVAL=2, N=2, LDA=N)
!
    INTEGER NEVAL, NOUT
    REAL EHIGH, ELOW, EVAL (MXEVAL)
    COMPLEX A(LDA,N)
        Set values of A
        A = (\begin{array}{llr}{1}&{-i}\end{array})
    DATA A/(1.0,0.0), (0.0,1.0), (0.0,-1.0), (1.0,0.0)/
    Find eigenvalue
    ELOW = 1.5
    EHIGH = 2.5
    CALL EVBHF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL)
    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,'(/,A,I3)') ' NEVAL = ', NEVAL
CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
END
```


## Output

NEVAL = 1

EVAL
2.000

## EVFHF

Computes the eigenvalues in a given range and the corresponding eigenvectors of a complex Hermitian matrix.

## Required Arguments

MXEVAL - Maximum number of eigenvalues to be computed. (Input)
A - Complex Hermitian matrix of order N. (Input)
Only the upper triangle is used.
ELOW - Lower limit of the interval in which the eigenvalues are sought. (Input)
EHIGH - Upper limit of the interval in which the eigenvalues are sought. (Input)
NEVAL - Number of eigenvalues found. (Output)

EVAL - Real vector of length MXEVAL containing the eigenvalues of $A$ in the interval (ELOW, EHIGH) in decreasing order of magnitude. (Output)
Only the first NEVAL elements of EVAL are significant.
EVEC - Complex matrix containing in its first NEVAL columns the eigenvectors associated with the eigenvalues found stored in EVAL. Each vector is normalized to have Euclidean length equal to the value one. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=$ SIZE (EVEC, 1 ).

## FORTRAN 90 Interface

Generic: CALL EVFHF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL, EVEC [,...])
Specific: The specific interface names are S_EVFHF and D_EVFHF.

## FORTRAN 77 Interface

Single: CALL EVFHF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL, EVEC, LDEVEC)

Double: The double precision name is DEVHFH.

## Description

Routine EVFHF computes the eigenvalues in a given range and the corresponding eigenvectors of a complex Hermitian matrix. Unitary transformations are used to reduce the matrix to an equivalent symmetric tridiagonal matrix. A bisection algorithm is used to compute the eigenvalues in the given range of this tridiagonal matrix. Inverse iteration is used to compute the eigenvectors of the tridiagonal matrix. The eigenvectors of the original matrix are computed by back transforming the eigenvectors of the tridiagonal matrix.

The reduction routine is based on the EISPACK routine HTRIDI. The bisection routine is based on the EISPACK routine BISECT. The inverse iteration routine is based on the EISPACK routine TINVIT. The back transformation routine is based on the EISPACK routine HTRIBK. See Smith et al. (1976) for the EISPACK routines.

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3FHF/DE3FHF. The reference is:

CALL E3FHF (N, MXEVAL, A, LDA, ELOW, EHIGH, NEVAL, EVAL, EVEC, LDEVEC, ACOPY, ECOPY, RWK, CWK, IWK)

The additional arguments are as follows:
ACOPY - Complex work matrix of size N by N. A and ACOPY may be the same, in which case the first $N^{2}$ elements of $A$ will be destroyed.

ECOPY - Work matrix of size N by MXEVAL. Used to store eigenvectors of a real tridiagonal matrix.
$\boldsymbol{R W K}$ — Work array of length 8 N .

CWK - Complex work array of length 2 N .
IWK - Work array of length MXEVAL.
2. Informational errors
Type Code

31 The number of eigenvalues in the specified range exceeds MXEVAL. NEVAL contains the number of eigenvalues in the range. No eigenvalues will be computed.
32 The iteration for an eigenvector failed to converge. The eigenvector will be set to 0 .
33 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.

## Example

In this example, a DATA statement is used to set $A$ to a complex Hermitian matrix. The eigenvalues in the range $[-15,0]$ and their corresponding eigenvectors are computed and printed. As a test, this example uses MXEVAL $=3$. The routine EVFHF computes the number of eigenvalues in the given range. That value, NEVAL, is two. As a check on the computations, the performance index is also computed and printed. For more details, see routine EPIHF.

```
USE IMSL_LIBRARIES
```

IMPLICIT NONE

```
! Declare variables
INTEGER LDA, LDEVEC, MXEVAL, N
PARAMETER (MXEVAL=3, N=3, LDA=N, LDEVEC=N)
```

!

```
INTEGER NEVAL, NOUT
REAL EHIGH, ELOW, EVAL (MXEVAL), PI
COMPLEX A(LDA,N), EVEC(LDEVEC,MXEVAL)
        Set values of A
            A =((1, 0) ( 1,-7i) ( 0,- i))
            ((1,7i) ( 5, 0) (10,-3i))
    ((0, i) ( 10, 3i) (-2, 0))
DATA A/ (1.0,0.0), (1.0,7.0), (0.0,1.0), (1.0,-7.0), (5.0,0.0), &
    (10.0,3.0), (0.0,-1.0), (10.0,-3.0), (-2.0,0.0)/
!
ELOW = -15.0
EHIGH = 0.0
CALL EVFHF (MXEVAL, A, ELOW, EHIGH, NEVAL, EVAL, EVEC)
                            Compute performance index
PI = EPIHF(NEVAL,A,EVAL,EVEC)
                                    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,'(/,A,I3)') ' NEVAL = ', NEVAL
CALL WRRRN ('EVAL', EVAL, 1, NEVAL, 1)
CALL WRCRN ('EVEC', EVEC, N, NEVAL, LDEVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output

```
NEVAL = 2
    EVAL
    1
```


## EVEC

1 ..... 2
$1(-0.0598,-0.3117)(0.8539,0.0000)$
$2(-0.5939,0.1841)(-0.0313,-0.1380)$
$3(0.7160,0.0000)(0.0808,-0.4942)$
Performance index $=0.057$

## EPIHF

This function computes the performance index for a complex Hermitian eigensystem.

## Function Return Value

EPIHF - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs on which the performance index computation is based. (Input)
$\boldsymbol{A}$ - Complex Hermitian matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Vector of length NEVAL containing eigenvalues of A. (Input)
EVEC - Complex n by NEVAL array containing eigenvectors of A. (Input)
The eigenvector corresponding to the eigenvalue $\operatorname{EVAL}(\mathrm{J})$ must be in the $J$-th column of EVEC.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: $\operatorname{LDEVEC}=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: EPIHF (NEVAL, A, EVAL, EVEC [, ...])

Specific: The specific interface names are S_EPIHF and D_EPIHF.

## FORTRAN 77 Interface

Single: EPIHF (N, NEVAL, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision function name is DEPIHF.

## Description

Let $M=$ NEVAL, $\lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, J)$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision, given by $\operatorname{AMACH}(4)$, see the Reference chapter of this manual. The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} x_{j}\right\|_{1}}{10 N \varepsilon\|A\|_{1}\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Smith et al. (1976, pages 124-125).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E2IHF / DE2IHF. The reference is:

E2IHF ( $\mathrm{N}, \mathrm{NEVAL}, \mathrm{A}, ~ L D A, ~ E V A L, ~ E V E C, ~ L D E V E C, ~ W K) ~$
The additional argument is

WK - Complex work array of length N .
2. Informational errors

Type Code
31 Performance index is greater than 100.
32 An eigenvector is zero.
33 The matrix is zero.

## Example

For an example of EPIHF, see IMSL routine EVCHF.

## EVLRH

Computes all of the eigenvalues of a real upper Hessenberg matrix.

## Required Arguments

A-Real upper Hessenberg matrix of order N. (Input)
EVAL - Complex vector of length N containing the eigenvalues in decreasing order of magnitude. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLRH (A, EVAL [,$\ldots]$ )
Specific: The specific interface names are S_EVLRH and D_EVLRH.

## FORTRAN 77 Interface

Single: CALL EVLRH (N, A, LDA, EVAL)
Double: The double precision name is DEVLRH.

## Description

Routine EVLRH computes the eigenvalues of a real upper Hessenberg matrix by using the QR algorithm. The QR Algorithm routine is based on the EISPACK routine HQR , Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LRH/DE3LRH. The reference is:
```
CALL E3LRH (N, A, LDA, EVAL, ACOPY, WK, IWK)
```

The additional arguments are as follows:
ACOPY-Real N by N work matrix.

WK - Real vector of length $3 n$.
$I W K$ - Integer vector of length $n$.
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge.

## Example

In this example, a DATA statement is used to set A to an upper Hessenberg matrix of integers. The eigenvalues of this matrix are computed and printed.

```
USE EVLRH INT
USE UMACH_INT
USE WRCRN_INT
```

```
    IMPLICIT NONE
    INTEGER LDA, N
    PARAMETER (N=4, LDA=N)
    INTEGER NOUT
    REAL A (LDA,N)
    COMPLEX EVAL (N)
    Set values of A
        A =(\begin{array}{llll}{2.0}&{1.0}&{3.0}&{4.0}\end{array})
        (\begin{array}{llll}{1.0}&{0.0}&{0.0}&{0.0}\end{array})
        ( 1.0 }\begin{array}{llll}{(}&{0.0}&{0.0}\\{(}&{1.0}&{0.0}\end{array}
    DATA A/2.0, 1.0, 0.0, 0.0, 1.0, 0.0, 1.0, 0.0, 3.0, 0.0, 0.0, &
        1.0, 4.0, 0.0, 0.0, 0.0/
    CALL EVLRH (A, EVAL)
    CALL UMACH (2, NOUT)
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
    END
```


## Output

|  |  |
| :---: | :---: |
| $(2.878, ~ 0.000)^{1}$ | $(0.011,1.243)^{2}$ |$(0.011,-1.243)^{3} \quad(-0.900,0.000)^{4}$

## EVCRH

Computes all of the eigenvalues and eigenvectors of a real upper Hessenberg matrix.

## Required Arguments

$\boldsymbol{A}$ - Real upper Hessenberg matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Complex vector of length N containing the eigenvalues in decreasing order of magnitude. (Output)
$\boldsymbol{E V E C}$ - Complex matrix of order N. (Output)
The J-th eigenvector, corresponding to $\operatorname{EVAL}(J)$, is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL EVCRH (A, EVAL, EVEC [,...])
Specific: The specific interface names are S_EVCRH and D_EVCRH.

## FORTRAN 77 Interface

Single: CALL EVCRH (N, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCRH.

## Description

Routine EVCRH computes the eigenvalues and eigenvectors of a real upper Hessenberg matrix by using the QR algorithm. The QR algorithm routine is based on the EISPACK routine HQR2; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E6CRH/DE6CRH. The reference is:
```
CALL E6CRH (N, A, LDA, EVAL, EVEC, LDEVEC, ACOPY, ECOPY, RWK, IWK)
```

The additional arguments are as follows:
ACOPY - Real N by N work matrix.

ECOPY - Real N by N work matrix.
$\boldsymbol{R W K}$ - Real array of length 3N.
IWK - Integer array of length N.
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge.

## Example

In this example, a DATA statement is used to set $A$ to a Hessenberg matrix with integer entries. The values are returned in decreasing order of magnitude. The eigenvalues, eigenvectors and performance index of this matrix are computed and printed. See routine EPIRG for details

```
USE EVCRH INT
USE EPIRG }\mp@subsup{}{}{-}\mathrm{ INT
USE UMACH INT
USE WRCRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDEVEC, N
PARAMETER (N=4, LDA=N, LDEVEC=N)
    INTEGER NOUT
    REAL A(LDA,N), PI
    COMPLEX EVAL(N), EVEC(LDEVEC,N)
                                    Define values of A:
                    A=( -1.0 -1.0 -1.0 -1.0 )
\(\left(\begin{array}{cccc}1.0 & 0.0 & 0.0 & 0.0\end{array}\right)\)
            (1.0 0.0 0.0)
    DATA A/-1.0, 1.0, 0.0, 0.0, -1.0, 0.0, 1.0, 0.0, -1.0, 0.0, 0.0, &
        1.0, -1.0, 0.0, 0.0, 0.0/
            Find eigenvalues and vectors of A
    CALL EVCRH (A, EVAL, EVEC)
                            Compute performance index
    PI = EPIRG(N,A,EVAL,EVEC)
    CALL UMACH (2, NOUT)
CALL WRCRN ('EVAL', EVAL, 1, N, 1)
CALL WRCRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```

!

## Output

| EVAL |  |  |  |
| :---: | :---: | :---: | :---: |
| $(-0.8090,0.5878)$ | $.8090,-0.5878)$ | $0.3090,0.9511)$ | $0.3090,-0.9511)$ |
| EVEC |  |  |  |
| 1 | 2 | 3 |  |
| 4 |  |  |  |
| $1(-0.4045,0.2939)$ | (-0.4045,-0.2939) | (-0.4045,-0.2939) | (-0.4045, |
| 0.2939) |  |  |  |
| $2(0.5000,0.0000)$ | $(0.5000,0.0000)$ | (-0.4045, 0.2939) | (-0.4045,- |
| $0.2939)$ |  |  |  |
| 3 (-0.4045,-0.2939) | $(-0.4045,0.2939)$ | ( 0.1545, 0.4755) | ( 0.1545, - |
| $0.4755)$ |  |  |  |

```
    4 (0.1545, 0.4755) ( 0.1545,-0.4755) ( 0.5000, 0.0000) ( 0.5000,
0.0000)
    Performance index = 0.098
```


## EVLCH

Computes all of the eigenvalues of a complex upper Hessenberg matrix.

## Required Arguments

A-Complex upper Hessenberg matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Complex vector of length N containing the eigenvalues of A in decreasing order of magnitude. (Output)

## Required Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVLCH (A, EVAL [,...])
Specific: The specific interface names are S_EVLCH and D_EVLCH.

## FORTRAN 77 Interface

Single: CALL EVLCH (N, A, LDA, EVAL)
Double: The double precision name is DEVLCH.

## Description

Routine EVLCH computes the eigenvalues of a complex upper Hessenberg matrix using the QR algorithm. This routine is based on the EISPACK routine COMQR2; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of E3LCH/DE3LCH. The reference is:

CALL E3LCH (N, A, LDA, EVAL, ACOPY, RWK, IWK)

The additional arguments are as follows:
ACOPY - Complex n by n work array. A and ACOPY may be the same, in which case A is destroyed.
$\boldsymbol{R W K}$ - Real work array of length N .
IWK - Integer work array of length N.
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge.

## Example

In this example, a DATA statement is used to set the matrix A. The program computes and prints the eigenvalues of this matrix.

```
USE EVLCH_INT
USE WRCRN_INT
IMPLICIT NONE
Declare variables
INTEGER LDA, N
PARAMETER (N=4, LDA=N)
COMPLEX A(LDA,N), EVAL(N)
            Set values of A
            A = (5+9i 5+5i -6-6i -7-7i)
                                    (3+3i 6+10i -5-5i -6-6i)
                                    ( 0 3+3i -1+3i -5-5i)
                                    ( 0 0 -3-3i 
DATA A / (5.0,9.0), (3.0,3.0), (0.0,0.0), (0.0,0.0), &
    (5.0,5.0), (6.0,10.0), (3.0,3.0), (0.0,0.0), &
    (-6.0,-6.0), (-5.0,-5.0), (-1.0,3.0), (-3.0,-3.0), &
    (-7.0,-7.0), (-6.0,-6.0), (-5.0,-5.0), (0.0,4.0)/
CALL EVLCH (A, EVAL)
                                    Print results
CALL WRCRN ('EVAL', EVAL, 1, N, 1)
END
```


## Output

$(8.22,12.22)^{1}(3.40,7.40)^{2}(1.60,5.60)^{3}(-3.22,0.78)^{4}$

## EVCCH

Computes all of the eigenvalues and eigenvectors of a complex upper Hessenberg matrix.

## Required Arguments

A - Complex upper Hessenberg matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Complex vector of length N containing the eigenvalues of $A$ in decreasing order of magnitude. (Output)
$\boldsymbol{E V E C}$ - Complex matrix of order N. (Output)
The $J$-th eigenvector, corresponding to $\operatorname{EVAL}(J)$, is stored in the $J$-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: CALL EVCCH (A, EVAL, EVEC [, ...])
Specific: The specific interface names are S_EVCCH and D_EVCCH.

## FORTRAN 77 Interface

Single: CALL EVCCH (N, A, LDA, EVAL, EVEC, LDEVEC)
Double: The double precision name is DEVCCH.

## Description

Routine EVCCH computes the eigenvalues and eigenvectors of a complex upper Hessenberg matrix using the QR algorithm. This routine is based on the EISPACK routine COMQR2; see Smith et al. (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of $\mathrm{E} 4 \mathrm{CCH} / \mathrm{DE} 4 \mathrm{CCH}$. The reference is:

CALL E4CCH (N, A, LDA, EVAL, EVEC, LDEVEC, ACOPY, CWORK, RWK, IWK)
The additional arguments are as follows:
ACOPY - Complex n by n work array. A and ACOPY may be the same, in which case $A$ is destroyed.

CWORK - Complex work array of length 2 N .
$\boldsymbol{R} W \boldsymbol{K}$ - Real work array of length N.
IWK - Integer work array of length N.
2 Informational error
Type Code
41 The iteration for the eigenvalues failed to converge.
3. The results of EVCCH can be checked using EPICG. This requires that the matrix A explicitly contains the zeros in $A(I, J)$ for $(I-1)>J$ which are assumed by EVCCH.

## Example

In this example, a DATA statement is used to set the matrix $A$. The program computes the eigenvalues and eigenvectors of this matrix. The performance index is also computed and printed. This serves as a check on the computations; for more details, see IMSL routine EPICG. The zeros in the lower part of the matrix are not referenced by EVCCH, but they are required by EPICG.

```
USE EVCCH_INT
USE EPICG_INT
USE UMACH INT
USE WRCRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDEVEC, N
PARAMETER (N=4, LDA=N, LDEVEC=N)
INTEGER NOUT
REAL PI
COMPLEX A(LDA,N), EVAL(N), EVEC(LDEVEC,N)
                                Set values of A
    A = (5+9i 5+5i -6-6i -7-7i)
                                    (3+3i 6+10i -5-5i -6-6i)
                                    ( 0
            ( 0 0 -3-3i 4i)
```

!

```
DATA A/ (5.0,9.0), (3.0,3.0), (0.0,0.0), (0.0,0.0), (5.0,5.0), &
    (6.0,10.0), (3.0,3.0), (0.0,0.0), (-6.0,-6.0), (-5.0,-5.0), &
    (-1.0,3.0), (-3.0,-3.0), (-7.0,-7.0), (-6.0,-6.0), &
    (-5.0,-5.0), (0.0,4.0)/
    CALL EVCCH (A, EVAL, EVEC)
    PI = EPICG(N,A,EVAL,EVEC)
    CALL UMACH (2, NOUT)
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
CALL WRCRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output

```
\((8.22,12.22)^{1} \quad(3.40,7.40)^{\text {EVAL }} \quad(1.60,5.60)^{3}(-3.22,0.78)^{4}\)
    1
4
    1 (0.7167, 0.0000) (-0.0704, 0.0000) (-0.3678, 0.0000) ( 0.5429,
0.0000)
    2 ( 0.6402,-0.0000) (-0.0046,-0.0000) ( 0.6767, 0.0000) ( 0.4298,-
0.0000)
    3 (0.2598, 0.0000) ( 0.7477, 0.0000) (-0.3005, 0.0000) ( 0.5277,-
0.0000)
    4(-0.0948,-0.0000) (-0.6603,-0.0000) ( 0.5625, 0.0000) (0.4920,-
0.0000)
    Performance index = 0.020
```


## GVLRG

Computes all of the eigenvalues of a generalized real eigensystem $A z=\lambda B z$.

## Required Arguments

$\boldsymbol{A}-$ Real matrix of order N. (Input)
$\boldsymbol{B}$ - Real matrix of order N. (Input)

ALPHA - Complex vector of size N containing scalars $\alpha_{i}, i=1, \ldots, n$. If $\beta_{i} \neq 0, \lambda_{i}=\alpha_{i} / \beta_{i}$ the eigenvalues of the system in decreasing order of magnitude. (Output)
$\boldsymbol{B E T A V}$ - Vector of size N containing scalars $\beta_{i}$. (Output)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB = SIZE $(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL GVLRG (A, B, ALPHA, BETAV [,...] )
Specific: The specific interface names are s_Gvirg and d_Gvirg.

## FORTRAN 77 Interface

Single: CALL GVLRG (N, A, LDA, B, LDB, ALPHA, BETAV)
Double: The double precision name is DGVLRG.

## Description

Routine GVLRG computes the eigenvalues of the generalized eigensystem $A x=\lambda B x$ where $A$ and $B$ are real matrices of order N . The eigenvalues for this problem can be infinite; so instead of returning $\lambda$, gVLRG returns $\alpha$ and $\beta$. If $\beta$ is nonzero, then $\lambda=\alpha / \beta$.
The first step of the QZ algorithm is to simultaneously reduce $A$ to upper Hessenberg form and $B$ to upper triangular form. Then, orthogonal transformations are used to reduce $A$ to quasi-uppertriangular form while keeping $B$ upper triangular. The generalized eigenvalues are then computed.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of $\operatorname{G3LRG} / \mathrm{DG} 3 \mathrm{LRG}$. The reference is:

CALL G3LRG ( $\mathrm{N}, \mathrm{A}, \mathrm{LDA}, \mathrm{B}, \mathrm{LDB}, \mathrm{ALPHA}, \mathrm{BETAV}, \mathrm{ACOPY}, \mathrm{BCOPY}, \mathrm{RWK}, \mathrm{CWK}, \mathrm{IWK})$
The additional arguments are as follows:

ACOPY - Work array of size $\mathrm{N}^{2}$. The arrays A and ACOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of A will be destroyed.
$\boldsymbol{B C O P Y}$ - Work array of size $\mathrm{N}^{2}$. The arrays B and BCOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of B will be destroyed.
$\boldsymbol{R} W \boldsymbol{K}$ - Real work array of size N .
$\boldsymbol{C W K}$ — Complex work array of size N .

IWK - Integer work array of size N .
2. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine G3LRG, the internal or working leading dimension of ACOPY is increased by $\operatorname{IVAL}(3)$ when $N$ is a multiple of $\operatorname{IVAL}(4)$. The values $\operatorname{IVAL}(3)$ and IVAL (4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine GVLRG. Analogous comments hold for BCOPY and the values IVAL(5) IVAL(8). Additional memory allocation and option value restoration are automatically done in GVLRG. There is no requirement that users change existing applications that use GVLRG or G3LRG. Default values for the option are IVAL $(*)=$ $1,16,0,1,1,16,0,1$.

## Example

In this example, DATA statements are used to set $A$ and $B$. The eigenvalues are computed and printed.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDA, LDB, N
PARAMETER ( N=3, LDA=N, LDB=N)
INTEGER I
REAL A(LDA,N), B(LDB,N), BETAV(N)
COMPLEX ALPHA(N), EVAL(N)
Set values of A and B
A = ( 1.0 0.5 0.0 )
    (-10.0 2.0 0.0 )
    ( 5.0 1.0 0.5 )
B =( ( 0.5 0.0 0.0 )
    ( 3.0 3.0 0.0)
    ( 4.0 0.5 1.0 )
Declare variables
DATA A/1.0, -10.0, 5.0, 0.5, 2.0, 1.0, 0.0, 0.0, 0.5/
DATA B/0.5, 3.0, 4.0, 0.0, 3.0, 0.5, 0.0, 0.0, 1.0/
```

$!$
$!$
CALL GVLRG (A, B, ALPHA, BETAV)
! Compute eigenvalues
DO $10 \quad \mathrm{I}=1, \mathrm{~N}$
$\operatorname{EVAL}(I)=A L P H A(I) / B E T A V(I)$
10 CONTINUE
! Print results
CALL WRCRN ('EVAL', EVAL, 1, N, 1)
END

## Output

EVAL
$(0.833,1.993)^{1} \quad(0.833,-1.993)^{2} \quad(0.500,0.000)^{3}$

## GVCRG

Computes all of the eigenvalues and eigenvectors of a generalized real eigensystem $A z=\lambda B z$.

## Required Arguments

A-Real matrix of order N . (Input)
B - Real matrix of order N. (Input)
ALPHA - Complex vector of size N containing scalars $\alpha_{i}$. If $\beta_{i} \neq 0, \lambda_{i}=\alpha_{i} / \beta_{i}, i=1, \ldots, n$ are the eigenvalues of the system.
$\boldsymbol{B E T A V}$ - Vector of size N containing scalars $\beta_{i}$. (Output)
EVEC - Complex matrix of order N . (Output)
The $J$-th eigenvector, corresponding to $\lambda_{J}$, is stored in the $J$-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDB - Leading dimension of B exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=$ SIZE (EVEC, 1 ).

## FORTRAN 90 Interface

Generic: CALL GVCRG (A, B, ALPHA, BETAV, EVEC [, ...])

Specific: The specific interface names are S_GVCRG and D_GVCRG.

## FORTRAN 77 Interface

Single: CALL GVCRG (N, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC)
Double: The double precision name is DGVCRG.

## Description

Routine GVCRG computes the complex eigenvalues and eigenvectors of the generalized eigensystem $A x=\lambda B x$ where $A$ and $B$ are real matrices of order $N$. The eigenvalues for this problem can be infinite; so instead of returning $\lambda$, GVCRG returns complex numbers $\alpha$ and real numbers $\beta$. If $\beta$ is nonzero, then $\lambda=\alpha / \beta$. For problems with small $|\beta|$ users can choose to solve the mathematically equivalent problem $B x=\mu A x$ where $\mu=\lambda^{-1}$.

The first step of the QZ algorithm is to simultaneously reduce $A$ to upper Hessenberg form and $B$ to upper triangular form. Then, orthogonal transformations are used to reduce $A$ to quasi-uppertriangular form while keeping $B$ upper triangular. The generalized eigenvalues and eigenvectors for the reduced problem are then computed.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual.

## Comments

1. Workspace may be explicitly provided, if desired, by use of $G 8 C R G / D G 8 C R G$. The reference is:

CALL G8CRG (N, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC, ACOPY, BCOPY, ECOPY, RWK, CWK, IWK)

The additional arguments are as follows:
ACOPY - Work array of size $\mathrm{N}^{2}$. The arrays A and ACOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of A will be destroyed.

BCOPY - Work array of size $\mathrm{N}^{2}$. The arrays B and BCOPY may be the same, in which case the first $\mathrm{N}^{2}$ elements of B will be destroyed.

ECOPY - Work array of size $\mathrm{N}^{2}$.
$\boldsymbol{R} W \boldsymbol{K}$ - Work array of size N .
$\boldsymbol{C W K}$ — Complex work array of size N .

IWK - Integer work array of size N.
2. Integer Options with Chapter 11 Options Manager

1 This option uses eight values to solve memory bank conflict (access inefficiency) problems. In routine G 8 CRG , the internal or working leading dimensions of ACOPY and ECOPY are both increased by IVAL(3) when $N$ is a multiple of IVAL(4). The values IVAL(3) and IVAL(4) are temporarily replaced by IVAL(1) and IVAL(2), respectively, in routine GVCRG. Analogous comments hold for the array BCOPY and the option values IVAL(5) - IVAL(8). Additional memory allocation and option value restoration are automatically done in GVCRG. There is no requirement that users change existing applications that use GVCRG or $\operatorname{G8CRG}$. Default values for the option are $\operatorname{IVAL}(*)=1,16,0,1,1,16$, 0 , 1. Items $5-8$ in $\operatorname{IVAL}(*)$ are for the generalized eigenvalue problem and are not used in GVCRG.

## Example

In this example, DATA statements are used to set $A$ and $B$. The eigenvalues, eigenvectors and performance index are computed and printed for the systems $A x=\lambda B x$ and $B x=\mu A x$ where $\mu=$ $\lambda^{-1}$. For more details about the performance index, see routine GPIRG.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDA, LDB, LDEVEC, N
PARAMETER ( $N=3, L D A=N, L D B=N, L D E V E C=N)$

INTEGER I, NOUT
REAL A $(L D A, N), B(L D B, N), \operatorname{BETAV}(N), \quad P I$
COMPLEX ALPHA (N), EVAL (N), EVEC (LDEVEC, N)
Define values of $A$ and $B$ :
$A=\left(\begin{array}{lll}1.0 & 0.5 & 0.0\end{array}\right)$
$(-10.0 \quad 2.0 \quad 0.0)$
( 5.0 1.0 0.5 )
$B=\left(\begin{array}{lll}0.5 & 0.0 & 0.0\end{array}\right)$
$\left(\begin{array}{lll}3.0 & 3.0 & 0.0\end{array}\right)$
( 4.0 0.5 1.0 ) Declare variables
DATA A/1.0, -10.0, 5.0, 0.5, 2.0, 1.0, 0.0, 0.0, 0.5/
DATA B/0.5, 3.0, 4.0, 0.0, 3.0, 0.5, 0.0, 0.0, 1.0/

```
    CALL GVCRG (A, B, ALPHA, BETAV, EVEC)
! Compute eigenvalues
    DO 10 I=1, N
                EVAL(I) = ALPHA(I)/BETAV(I)
    10 CONTINUE
! Compute performance index
    PI = GPIRG(N,A,B,ALPHA,BETAV,EVEC)
    CALL UMACH (2, NOUT)
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
    CALL WRCRN ('EVEC', EVEC)
    WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
                                    Solve for reciprocals of values
    CALL GVCRG (B, A, ALPHA, BETAV, EVEC)
! Compute reciprocals
    DO 20 I=1, N
        EVAL(I) = ALPHA(I)/BETAV(I)
    20 CONTINUE
    PI = GPIRG(N,B,A,ALPHA, BETAV,EVEC)
    Print results
    CALL WRCRN ('EVAL reciprocals', EVAL, 1, N, 1)
    CALL WRCRN ('EVEC', EVEC)
    WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
    END
```


## Output

| EVAL |  |  |  |
| :---: | :---: | :---: | :---: |
|  | 1 | 2 | 3 |
|  | 0.833, 1.993) ( | 0.833,-1.993) | 0.500, 0.000) |
| EVEC |  |  |  |
|  | 1 | 2 | 3 |
| 1 | (-0.197, 0.150) | (-0.197, -0.150) | (-0.000, 0.000) |
| 2 | (-0.069,-0.568) | (-0.069, 0.568) | (-0.000, 0.000) |
| 3 | ( 0.782, 0.000) | ( 0.782, 0.000) | ( 1.000, 0.000) |
| Performance index $=0.384$ |  |  |  |
| EVAL reciprocals |  |  |  |
|  | 1 | 2 | 3 |
| $(2.000,0.000)$ |  | 0.179, 0.427) | 0.179,-0.427) |
| EVEC |  |  |  |
|  | 1 | 2 | 3 |
| 1 | $(0.000,0.000)$ | (-0.197, -0.150) | (-0.197, 0.150) |
| 2 | ( 0.000, 0.000) | (-0.069, 0.568) | (-0.069,-0.568) |
| 3 | ( $1.000,0.000)$ | ( 0.782, 0.000) | ( 0.782, 0.000) |
| erformance index $=0.283$ |  |  |  |

## GPIRG

This function computes the performance index for a generalized real eigensystem $A z=\lambda B z$.

## Function Return Value

GPIRG - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs performance index computation is based on. (Input)
$\boldsymbol{A}$ - Real matrix of order N. (Input)
$\boldsymbol{B}$ - Real matrix of order N. (Input)
ALPHA - Complex vector of length NEVAL containing the numerators of eigenvalues. (Input)

BETAV - Real vector of length NEVAL containing the denominators of eigenvalues. (Input)
$\boldsymbol{E V E C}$ - Complex N by NEVAL array containing the eigenvectors. (Input)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\mathbf{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC, 1 ).

## FORTRAN 90 Interface

Generic: GPIRG (NEVAL, A, B, ALPHA, BETAV, EVEC, GPIRG [, ...])
Specific: The specific interface names are S_GPIRG and D_GPIRG.

## FORTRAN 77 Interface

Single: GPIRG (N, NEVAL, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC)
Double: The double precision function name is DGPIRG.

## Description

Let $M=\operatorname{NEVAL}, x_{j}=\operatorname{EVEC}(*, \mathrm{~J})$, the $j-t h$ column of $\operatorname{EVEC}$. Also, let $\varepsilon$ be the machine precision given by $\operatorname{AMACH}(4)$, see the Reference chapter of this manual. The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|\beta_{j} A x_{j}-\alpha_{j} B x_{j}\right\|_{1}}{\varepsilon\left(\left|\beta_{j}\right|\|A\|_{1}+\left|\alpha_{j}\right|\|B\|_{1}\right)\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Garbow et al. (1977, pages 77-79).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G2IRG/DG2IRG. The reference is:

G2IRG (N, NEVAL, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC, WK)
The additional argument is:
WK - Complex work array of length 2N.
2. Informational errors

Type Code
$31 \quad$ Performance index is greater than 100.
$3 \quad 2$ An eigenvector is zero.
33 The matrix A is zero.
34 The matrix B is zero.
3. The $J$-th eigenvalue should be $\operatorname{ALPHA}(J) / \operatorname{BETAV}(J)$, its eigenvector should be in the $J$ th column of EVEC.

## Example

For an example of GPIRG, see routine GVCRG.

## GVLCG

Computes all of the eigenvalues of a generalized complex eigensystem $A z=\lambda B z$.

## Required Arguments

$\boldsymbol{A}$ - Complex matrix of order N . (Input)
B - Complex matrix of order N . (Input)
ALPHA - Complex vector of length N . Ultimately, alpha $(i) /$ betav $(i)$ (for $i=1, n$ ), will be the eigenvalues of the system in decreasing order of magnitude. (Output)

BETAV - Complex vector of length N . (Output)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDB - Leading dimension of B exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=$ SIZE $(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL GVLCG (A, B, ALPhA, BETAV [,...])
Specific: The specific interface names are S_GVLCG and D_GVLCG.

## FORTRAN 77 Interface

Single: CALL GVLCG (N, A, LDA, B, LDB, ALPHA, BETAV)
Double: The double precision name is DGVLCG.

## Description

Routine GVLCG computes the eigenvalues of the generalized eigensystem $A x=\lambda B x$, where $A$ and B are complex matrices of order $n$. The eigenvalues for this problem can be infinite; so instead of returning $\lambda$, GVLCG returns $\alpha$ and $\beta$. If $\beta$ is nonzero, then $\lambda=\alpha / \beta$. If the eigenvectors are needed, then use gvccg.

The underlying code is based on either EISPACK or LAPACK code depending upon which supporting libraries are used during linking. For a detailed explanation, see "Using ScaLAPACK, LAPACK, LINPACK, and EISPACK" in the Introduction section of this manual. Some timing information is given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G3LCG/DG3LCG. The reference is:

CALL G3LCG (N, A, LDA, B, LDB, ALPHA, BETAV, ACOPY, BCOPY, CWK, WK, IWK)
The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same, in which case A will be destroyed.

BCOPY - Complex work array of length $\mathrm{N}^{2}$. B and BCOPY may be the same, in which case B will be destroyed.
$\boldsymbol{C W K}$ - Complex work array of length N .
$\boldsymbol{W} \boldsymbol{K}$ - Real work array of length N.
IWK - Integer work array of length N .
2. Informational error

Type Code
41 The iteration for the eigenvalues failed to converge.

## Example

In this example, DATA statements are used to set $A$ and $B$. Then, the eigenvalues are computed and printed.

```
USE GVLCG INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N
PARAMETER (N=5, LDA=N, LDB=N)
INTEGER I
COMPLEX A (LDA,N), ALPHA (N), B(LDB,N), BETAV (N), EVAL (N)
    Define values of A and B
DATA A/(-238.0,-344.0), (76.0,152.0), (118.0,284.0), &
    (-314.0,-160.0), (-54.0,-24.0), (86.0,178.0), &
    (-96.0,-128.0), (55.0,-182.0), (132.0,78.0), &
```

!

```
        (-205.0,-400.0), (164.0,240.0), (40.0,-32.0), &
        (-13.0,460.0), (114.0,296.0), (109.0,148.0), &
        (-166.0,-308.0), (60.0,184.0), (34.0,-192.0), &
        (-90.0,-164.0), (158.0,312.0), (56.0,158.0), &
        (-60.0,-136.0), (-176.0,-214.0), (-424.0,-374.0), &
        (-38.0,-96.0) /
    DATA B/ (388.0,94.0), (-304.0,-76.0), (-658.0,-136.0), &
    (-640.0,-10.0), (-162.0,-72.0), (-386.0,-122.0), &
    (384.0,64.0), (-73.0,100.0), (204.0,-42.0), (631.0,158.0), &
    (-250.0,-14.0), (-160.0,16.0), (-109.0,-250.0), &
    (-692.0,-90.0), (131.0,52.0), (556.0,130.0), &
    (-240.0,-92.0), (-118.0,100.0), (288.0,66.0), &
    (-758.0,-184.0), (-396.0,-62.0), (240.0,68.0), &
    (406.0,96.0), (-192.0,154.0), (278.0,76.0)/
!
    CALL GVLCG (A, B, ALPHA, BETAV)
                                    Compute eigenvalues
        EVAL = ALPHA/BETAV
! Print results
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
    STOP
    END
```

    Output
    |  |  |
| :---: | :---: |
| $(-1.000,-1.333)^{1}$ | $(0.765,0.941)^{2}$ |$(-0.353,0.412)^{3} \quad(-0.353,-0.412)^{4}$

$(-0.353,-0.412)$

## GVCCG

Computes all of the eigenvalues and eigenvectors of a generalized complex eigensystem $A z=\lambda B z$.

## Required Arguments

A-Complex matrix of order N . (Input)
B - Complex matrix of order N. (Input)
ALPHA - Complex vector of length N . Ultimately, alpha $(i) / \operatorname{betav}(i)($ for $i=1, \ldots, n)$, will be the eigenvalues of the system in decreasing order of magnitude. (Output)

BETAV - Complex vector of length N . (Output)
$\boldsymbol{E V E C}$ - Complex matrix of order N. (Output)
The J-th eigenvector, corresponding to $\operatorname{ALPHA}(J) / \operatorname{BETAV}(J)$, is stored in the J-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: CALL GVCCG (A, B, ALPHA, BETAV, EVEC [, ...])
Specific: The specific interface names are S_GVCCG and D_GVCCG.

## FORTRAN 77 Interface

Single: CALL GVCCG (N, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC)
Double: The double precision name is DGVCCG.

## Description

Routine GVCCG computes the eigenvalues and eigenvectors of the generalized eigensystem $A x=$ $\lambda B x$. Here, $A$ and $B$, are complex matrices of order $n$. The eigenvalues for this problem can be infinite; so instead of returning $\lambda$, GVCCG returns $\alpha$ and $\beta$. If $\beta$ is nonzero, then $\lambda=\alpha / \beta$.

The routine GVCCG uses the QZ algorithm described by Moler and Stewart (1973). The implementation is based on routines of Garbow (1978). Some timing results are given in Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G6CCG/DG6CCG. The reference is:

CALL G6CCG (N, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC, ACOPY, BCOPY, CWK, WK, IWK)

The additional arguments are as follows:
ACOPY - Complex work array of length $\mathrm{N}^{2}$. A and ACOPY may be the same in which case the first $\mathrm{N}^{2}$ elements of A will be destroyed.

BCOPY - Complex work array of length $\mathrm{N}^{2}$. B and BCOPY may be the same in which case the first $\mathrm{N}^{2}$ elements of B will be destroyed.

CWK - Complex work array of length N .

WK - Real work array of length n.

IWK - Integer work array of length N.
2. Informational error

Type Code
41 The iteration for an eigenvalue failed to converge.
3. The success of this routine can be checked using GPICG.

## Example

In this example, DATA statements are used to set $A$ and $B$. The eigenvalues and eigenvectors are computed and printed. The performance index is also computed and printed. This serves as a check on the computations. For more details, see routine GPICG.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDA, LDB, LDEVEC, N
PARAMETER (N=3, LDA=N, LDB=N, LDEVEC=N)
INTEGER I, NOUT
REAL PI
COMPLEX A(LDA,N), ALPHA(N), B(LDB,N), BETAV(N), EVAL(N), &
        EVEC(LDEVEC,N)
                Define values of A and B
                A = ( 1+0i 0.5+i 0+5i )
            (-10+0i 2+i 0+0i )
            ( 5+i 1+0i 0.5+3i )
                B =( 0.5+0i 0+0i 0+0i )
            (}\begin{array}{rrrl}{3+3i}&{3+3i}&{0+i}\end{array}
                Declare variables
DATA A/ (1.0,0.0), (-10.0,0.0), (5.0,1.0), (0.5,1.0), (2.0,1.0), &
```

```
        (1.0,0.0), (0.0,5.0), (0.0,0.0), (0.5,3.0)/
        DATA B/ (0.5,0.0), (3.0,3.0), (4.0,2.0), (0.0,0.0), (3.0,3.0), &
        (0.5,1.0), (0.0,0.0), (0.0,1.0), (1.0,1.0)/
                                    Compute eigenvalues
    CALL GVCCG (A, B, ALPHA, BETAV, EVEC)
                                    EVAL = ALPHA/BETAV
                                    Compute performance index
    PI = GPICG(N,A,B,ALPHA,BETAV,EVEC)
    CALL UMACH (2, NOUT)
    CALL WRCRN ('EVAL', EVAL, 1, N, 1)
    CALL WRCRN ('EVEC', EVEC)
    WRITE (NOUT, '(/,A,F6.3)') ' Performance index = ', PI
    END
```


## Output

EVAL
$(-8.18,-25.38)^{1} \quad(2.18,0.61)^{2}(0.12,-0.39)^{3}$

EVEC

|  |  | 1 | 2 |
| ---: | ---: | ---: | ---: | ---: |
| 1 | $(-0.3267,-0.1245)$ | $(-0.3007,-0.2444)$ | $(0.0371,0.1518)$ |
| 2 | $(0.1767,0.0054)$ | $(0.8959,0.0000)$ | $(0.9577,0.0000)$ |
| 3 | $(0.9201,0.0000)$ | $(-0.2019,0.0801)$ | $(-0.2215,0.0968)$ |

Performance index $=0.709$

## GPICG

This function computes the performance index for a generalized complex eigensystem $A z=\lambda B z$.

## Function Return Value

GPICG - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs performance index computation is based on. (Input)
$\boldsymbol{A}$ - Complex matrix of order N. (Input)
B - Complex matrix of order N. (Input)
ALPHA - Complex vector of length NEVAL containing the numerators of eigenvalues. (Input)

BETAV - Complex vector of length NEVAL containing the denominators of eigenvalues. (Input)

EVEC - Complex n by NEVAL array containing the eigenvectors. (Input)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB = SIZE $(\mathrm{B}, 1)$.

LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC $=\operatorname{SIZE}(E V E C, 1)$.

## FORTRAN 90 Interface

Generic: GPICG (NEVAL, A, B, ALPHA, BETAV, EVEC [, ...])
Specific: The specific interface names are S_GPICG and D_GPICG.

## FORTRAN 77 Interface

Single: GPICG (N, NEVAL, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC)
Double: The double precision name is DGPICG.

## Description

Let $M=\operatorname{NEVAL}, x_{j}=\operatorname{EVEC}(*, J)$, the $j$-th column of $\operatorname{EVEC}$. Also, let $\varepsilon$ be the machine precision given by $\operatorname{AMACH}(4)$. The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|\beta_{j} A x_{j}-\alpha_{j} B x_{j}\right\|_{1}}{\varepsilon\left(\left|\beta_{j}\right|\|A\|_{1}+\left|\alpha_{j}\right|\|B\|_{1}\right)\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$.
The performance index was first developed by the EISPACK project at Argonne National Laboratory; see Garbow et al. (1977, pages 77-79).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G2ICG/DG2ICG. The reference is:

G2ICG (N, NEVAL, A, LDA, B, LDB, ALPHA, BETAV, EVEC, LDEVEC, WK)
The additional argument is:
$W K$ - Complex work array of length 2 N .
2. Informational errors

Type Code
$31 \quad$ Performance index is greater than 100.
32 An eigenvector is zero.
33 The matrix A is zero.
34 The matrix B is zero.
3. The $J$-th eigenvalue should be $\operatorname{ALPHA}(\mathrm{J}) / \operatorname{BETAV}(\mathrm{J})$, its eigenvector should be in the J th column of EVEC.

## Example

For an example of GPICG, see routine GVCCG.

## GVLSP

Computes all of the eigenvalues of the generalized real symmetric eigenvalue problem $A z=\lambda B z$, with $B$ symmetric positive definite.

## Required Arguments

A-Real symmetric matrix of order N. (Input)
$\boldsymbol{B}$ - Positive definite symmetric matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Vector of length N containing the eigenvalues in decreasing order of magnitude.
(Output)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

LDB - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(\mathrm{B}, 1)$.

## FORTRAN 90 Interface

Generic: CALL GVLSP (A, B, EVAL [,...])
Specific: The specific interface names are S_GVLSP and D_GVLSP.

## FORTRAN 77 Interface

Single: CALL GVLSP (N, A, LDA, B, LDB, EVAL)
Double: The double precision name is DGVLSP.

## Description

Routine GVLSP computes the eigenvalues of $A x=\lambda B x$ with $A$ symmetric and $B$ symmetric positive definite. The Cholesky factorization $B=R^{T} R$, with $R$ a triangular matrix, is used to transform the equation $A x=\lambda B x$ to

$$
\left(R^{-T} A R^{-1}\right)(R x)=\lambda(R x)
$$

The eigenvalues of $C=R^{-T} A R^{-1}$ are then computed. This development is found in Martin and Wilkinson (1968). The Cholesky factorization of $B$ is computed based on IMSL routine LFTDS, (see Chapter 1, Linear Systems). The eigenvalues of $C$ are computed based on routine EVLSF. Further discussion and some timing results are given Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G3LSP / DG3LSP. The reference is:

CALL G3LSP (N, A, LDA, B, LDB, EVAL, IWK, WK1, WK2)
The additional arguments are as follows:

IWK - Integer work array of length N.
WK1 - Work array of length 2N.
$W K 2$ - Work array of length $\mathrm{N}^{2}+\mathrm{N}$.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 1 | The iteration for an eigenvalue failed to converge. |
| 4 | 2 | Matrix B is not positive definite. |

## Example

In this example, a DATA statement is used to set the matrices $A$ and $B$. The eigenvalues of the system are computed and printed.

```
USE GVLSP INT
USE WRRRN__INT
IMPLICIT NONE
INTEGER LDA, LDB, N
PARAMETER (N=3, LDA=N, LDB=N)
REAL A(LDA,N), B (LDB,N), EVAL (N)
    Define values of A:
        A=(\begin{array}{llll}{2}&{3}&{5}\end{array})
        ( (\begin{array}{lll}{3}&{2}&{4}\\{(}&{5}&{4}\end{array})
DATA A/2.0, 3.0, 5.0, 3.0, 2.0, 4.0, 5.0, 4.0, 2.0/
    Define values of B:
        B =(\begin{array}{llll}{3}&{1}&{0}\end{array})
            ( 1
DATA B/3.0, 1.0, 0.0, 1.0, 2.0, 1.0, 0.0, 1.0, 1.0/
CALL GVLSP (A, B, EVAL)
CALL WRRRN ('EVAL', EVAL, 1, N, 1)
END
```


## Output

| EVAL |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -4.717 | 4.393 | -0.676 |

## GVCSP

Computes all of the eigenvalues and eigenvectors of the generalized real symmetric eigenvalue problem $A z=\lambda B z$, with $B$ symmetric positive definite.

## Required Arguments

A-Real symmetric matrix of order n. (Input)
$\boldsymbol{B}$ - Positive definite symmetric matrix of order N. (Input)
EVAL - Vector of length N containing the eigenvalues in decreasing order of magnitude. (Output)

EVEC - Matrix of order N. (Output)
The $J$-th eigenvector, corresponding to $\operatorname{EVAL}(J)$, is stored in the $J$-th column. Each vector is normalized to have Euclidean length equal to the value one.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: CALL GVCSP (A, B, EVAL, EVEC [, ...])
Specific: The specific interface names are S_GVCSP and D_GVCSP.

## FORTRAN 77 Interface

Single: CALL GVCSP (N, A, LDA, B, LDB, EVAL, EVEC, LDEVEC)
Double: The double precision name is DGVCSP.

## Description

Routine GVLSP computes the eigenvalues and eigenvectors of $A z=\lambda B z$, with $A$ symmetric and $B$ symmetric positive definite. The Cholesky factorization $B=R^{T} R$, with $R$ a triangular matrix, is used to transform the equation $A z=\lambda B z$, to

$$
\left(R^{-\mathrm{T}} A R^{-1}\right)(R z)=\lambda(R z)
$$

The eigenvalues and eigenvectors of $C=R^{-\mathrm{T}} A R^{-1}$ are then computed. The generalized eigenvectors of $A$ are given by $z=R^{-1} x$, where $x$ is an eigenvector of $C$. This development is
found in Martin and Wilkinson (1968). The Cholesky factorization is computed based on IMSL routine LFTDS, see Chapter 1, Linear Systems. The eigenvalues and eigenvectors of $C$ are computed based on routine EVCSF. Further discussion and some timing results are given Hanson et al. (1990).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G3CSP/DG3CSP. The reference is:

CALL G3CSP (N, A, LDA, B, LDB, EVAL, EVEC, LDEVEC, IWK, WK1, WK2)
The additional arguments are as follows:
$I W K$ - Integer work array of length N.
WK1 - Work array of length 3N.
$W K 2$ - Work array of length $\mathrm{N}^{2}+\mathrm{N}$.
2. Informational errors
Type Code
$41 \quad$ The iteration for an eigenvalue failed to converge.
42 Matrix B is not positive definite.
3. The success of this routine can be checked using GPISP.

## Example

In this example, a DATA statement is used to set the matrices $A$ and $B$. The eigenvalues, eigenvectors and performance index are computed and printed. For details on the performance index, see IMSL routine GPISP.

```
USE GVCSP INT
USE GPISP - INT
USE UMACH_INT
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDB, LDEVEC, N
PARAMETER (N=3, LDA=N, LDB=N, LDEVEC=N)
INTEGER NOUT
REAL A(LDA,N), B(LDB,N), EVAL(N), EVEC(LDEVEC,N), PI
    Define values of A:
    A =( (1.1 1.2 1.4 )
        (1.2 1.3 1.5 )
        (1.4 1.5 1.6 )
DATA A/1.1, 1.2, 1.4, 1.2, 1.3, 1.5, 1.4, 1.5, 1.6/
```

```
! Define values of B:
    B = ( 2.0 1.0 0.0 )
    ( 1.0 2.0 1.0 )
DATA B/2.0, 1.0, 0.0, 1.0, 2.0, 1.0, 0.0, 1.0, 2.0/
CALL GVCSP (A, B, EVAL, EVEC)
Compute performance index
PI = GPISP(N,A,B,EVAL,EVEC)
CALL UMACH (2, NOUT)
CALL WRRRN ('EVAL', EVAL)
CALL WRRRN ('EVEC', EVEC)
WRITE (NOUT,'(/,A,F6.3)') ' Performance index = ', PI
END
```


## Output

```
    EVAL
```

$1 \quad 1.386$
-0.058
$3-0.003$

EVEC

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 0.6431 | -0.1147 | -0.6817 |
| 2 | -0.0224 | -0.6872 | 0.7266 |
| 3 | 0.7655 | 0.7174 | -0.0858 |
|  |  |  |  |
| Performance index $=$ | 0.417 |  |  |

## GPISP

This function computes the performance index for a generalized real symmetric eigensystem problem.

## Function Return Value

GPISP - Performance index. (Output)

## Required Arguments

NEVAL - Number of eigenvalue/eigenvector pairs that the performance index computation is based on. (Input)

A - Symmetric matrix of order N. (Input)
B - Symmetric matrix of order N. (Input)
$\boldsymbol{E V A L}$ - Vector of length NEVAL containing eigenvalues. (Input)
$\boldsymbol{E V E C}$ - N by NEVAL array containing the eigenvectors. (Input)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement in the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
LDEVEC - Leading dimension of EVEC exactly as specified in the dimension statement in the calling program. (Input)
Default: LDEVEC = SIZE (EVEC,1).

## FORTRAN 90 Interface

Generic: GPISP (NEVAL, A, B, EVAL, EVEC [, ...])
Specific: The specific interface names are S_GPISP and D_GPISP.

## FORTRAN 77 Interface

Single: GPISP (N, NEVAL, A, LDA, B, LDB, EVAL, EVEC, LDEVEC)
Double: The double precision name is DGPISP.

## Description

Let $M=\operatorname{NEVAL}, \lambda=\operatorname{EVAL}, x_{j}=\operatorname{EVEC}(*, J)$, the $j$-th column of EVEC. Also, let $\varepsilon$ be the machine precision given by АМАСН(4). The performance index, $\tau$, is defined to be

$$
\tau=\max _{1 \leq j \leq M} \frac{\left\|A x_{j}-\lambda_{j} B x_{j}\right\|_{1}}{\varepsilon\left(\|A\|_{1}+\left|\lambda_{j}\right|\|B\|_{1}\right)\left\|x_{j}\right\|_{1}}
$$

The norms used are a modified form of the 1-norm. The norm of the complex vector $v$ is

$$
\|v\|_{1}=\sum_{i=1}^{N}\left\{\left|\mathfrak{R} v_{i}\right|+\left|\mathfrak{J} v_{i}\right|\right\}
$$

While the exact value of $\tau$ is highly machine dependent, the performance of EVCSF is considered excellent if $\tau<1$, good if $1 \leq \tau \leq 100$, and poor if $\tau>100$. The performance index was first
developed by the EISPACK project at Argonne National Laboratory; see Garbow et al. (1977, pages 77-79).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G2ISP/DG2ISP. The reference is:

G2ISP (N, NEVAL, A, LDA, B, LDB, EVAL, EVEC, LDEVEC, WORK)
The additional argument is:
WORK - Work array of length 2 * N .
2. Informational errors

Type Code
31 Performance index is greater than 100.
32 An eigenvector is zero.
33 The matrix A is zero.
34 The matrix B is zero.
3. The $J$-th eigenvalue should be $\operatorname{ALPHA}(J) / \operatorname{BETAV}(J)$, its eigenvector should be in the $J$ th column of EVEC.

## Example

For an example of GPISP, see routine GVCSP.

## Chapter 3: Interpolation and Approximation

## Routines

3.1 Curve and Surface Fitting with Splines
Returns the derived type array result ..... SPLINE_CONSTRAINTS ..... 652
Returns an array result, given an array of input SPLINE_VALUES ..... 653
Weighted least-squares fitting by B -splines to discrete One-Dimensional data is performed SPLINE_FITTING ..... 654
Returns the derived type array result given optional input......................................SURFACE_CONSTRAINTS ..... 664
Returns a tensor product array result, given two arrays of independent variable values SURFACE_VALUES ..... 665
Weighted least-squares fitting by tensor product
B-splines to discrete two-dimensional datais performedSURFACE_FITTING666
3.2. Cubic Spline Interpolation
Easy to use cubic spline routine ..... CSIEZ ..... 677
Not-a-knot ..... CSINT ..... 680
Derivative end conditions. CSDEC ..... 682
Hermite ..... CSHER ..... 687
Akima CSAKM ..... 690
Shape preserving ..... CSCON ..... 692
Periodic .CSPER ..... 696
3.3. Cubic Spline Evaluation and Integration
Evaluation CSVAL ..... 699
Evaluation of the derivative. CSDER ..... 700
Evaluation on a grid .CS1GD ..... 703
Integration CSITG ..... 706
3.4. B-spline Interpolation
SPLEZ708
One-dimensional interpolation ..... BSINT ..... 711
Knot sequence given interpolation data BSNAK ..... 715
Optimal knot sequence given interpolation data BSOPK ..... 718
Two-dimensional tensor product interpolation ..... BS2IN ..... 720
Three-dimensional tensor product interpolation BS3IN ..... 725
3.5. Spline Evaluation, Integration, and Conversion to Piecewise Polynomial Given the B-spline Representation Evaluation..........................................................................BSVAL ..... 731
Evaluation of the derivative BSDER ..... 732
Evaluation on a grid. BS1GD ..... 735
One-dimensional integration ..... BSITG ..... 738
Two-dimensional evaluation ..... 741
Two-dimensional evaluation of the derivative ..... 742
Two-dimensional evaluation on a grid ..... 746
Two-dimensional integration ..... 750
Three-dimensional evaluation ..... 754
Three-dimensional evaluation of the derivative ..... 756
Three-dimensional evaluation on a grid ..... 760
Three-dimensional integration ..... 766
Convert B-spline representation to piecewise polynomial BSCPP ..... 770
3.6. Piecewise Polynomial
Evaluation PPVAL ..... 771
Evaluation of the derivative PPDER ..... 774
Evaluation on a grid PP1GD ..... 776
Integration .PPITG ..... 780
3.7. Quadratic Polynomial Interpolation Routines for Gridded Data
One-dimensional evaluation ..... QDVAL ..... 782
One-dimensional evaluation of the derivative QDDER ..... 784
Two-dimensional evaluation QD2VL ..... 786
Two-dimensional evaluation of the derivative QD2DR ..... 789
Three-dimensional evaluation QD3VL ..... 792
Three-dimensional evaluation of the derivative QD3DR ..... 796
3.8. Scattered Data Interpolation
Akima's surface fitting method ..... SURF ..... 800
3.9. Least-Squares Approximation
Linear polynomial RLINE ..... 803
General polynomial RCURV ..... 806
General functions FNLSQ ..... 811
Splines with fixed knots BSLSQ ..... 815
Splines with variable knot. ..... 819
Splines with linear constraints CONFT ..... 824
Two-dimensional tensor-product splines with fixed knots.... BSLS2 ..... 833
Three-dimensional tensor-product splines with fixed knots. BSLS3 ..... 838
3.10. Cubic Spline Smoothing
Smoothing by error detection ..... CSSED ..... 844
Smoothing spline ..... CSSMH ..... 848
3.11. Rational L $\infty$ Approximation Rational Chebyshev. RATCH

## Usage Notes

The majority of the routines in this chapter produce piecewise polynomial or spline functions that either interpolate or approximate given data, or are support routines for the evaluation, integration, and conversion from one representation to another. Two major subdivisions of routines are provided. The cubic spline routines begin with the letters "CS" and utilize the piecewise polynomial representation described below. The B-spline routines begin with the letters "BS" and utilize the B-spline representation described below. Most of the spline routines are based on routines in the book by de Boor (1978).

## Piecewise Polynomials

A univariate piecewise polynomial (function) $p$ is specified by giving its breakpoint sequence $\xi \in \mathbf{R}^{n}$, the order $k$ (degree $k-1$ ) of its polynomial pieces, and the $k \times(n-1)$ matrix $c$ of its local polynomial coefficients. In terms of this information, the piecewise polynomial (pp) function is given by

$$
p(x)=\sum_{j=1}^{k} c_{j i} \frac{\left(x-\xi_{i}\right)^{j-1}}{(j-1)!} \quad \text { for } \xi_{i} \leq x<\xi_{i+1}
$$

The breakpoint sequence $\xi$ is assumed to be strictly increasing, and we extend the pp function to the entire real axis by extrapolation from the first and last intervals. The subroutines in this chapter will consistently make the following identifications for FORTRAN variables:

$$
\begin{aligned}
& c=\text { PPCOEF } \\
& \xi=\text { BREAK } \\
& k=\text { KORDER } \\
& N=\text { NBREAK }
\end{aligned}
$$

This representation is redundant when the pp function is known to be smooth. For example, if $p$ is known to be continuous, then we can compute $c_{1, i+1}$ from the $c_{j i}$ as follows

$$
c_{1, i+1}=p\left(\xi_{i+1}\right)=c_{1 i}+c_{2 i} \Delta \xi_{i}+\ldots+c_{k i} \frac{\left(\Delta \xi_{i}\right)^{k-1}}{(k-1)!}
$$

where $\Delta \xi_{i}:=\xi_{i+1}-\xi_{i}$. For smooth pp, we prefer to use the irredundant representation in terms of the B-(for 'basis')-splines, at least when such a function is first to be determined. The above pp representation is employed for evaluation of the pp function at many points since it is more efficient.

## Splines and B-splines

B-splines provide a particularly convenient and suitable basis for a given class of smooth pp functions. Such a class is specified by giving its breakpoint sequence, its order, and the required smoothness across each of the interior breakpoints. The corresponding B-spline basis is specified by giving its knot sequence $\mathbf{t} \in \mathbf{R}^{M}$. The specification rule is the following: If the class is to have all derivatives up to and including the $j$-th derivative continuous across the interior breakpoint $\xi_{i}$, then the number $\xi_{i}$ should occur $k-j-1$ times in the knot sequence. Assuming that $\xi_{1}$, and $\xi_{n}$ are the endpoints of the interval of interest, one chooses the first $k$ knots equal to $\xi_{1}$ and the last $k$ knots equal to $\xi_{n}$. This can be done since the B-splines are defined to be right continuous near $\xi_{1}$ and left continuous near $\xi_{n}$.

When the above construction is completed, we will have generated a knot sequence $\boldsymbol{t}$ of length $M$; and there will be $m:=M-k$ B-splines of order $k$, say $B_{1}, \ldots, B_{m}$ that span the pp functions on the interval with the indicated smoothness. That is, each pp function in this class has a unique representation

$$
p=a_{1} B_{1}+a_{2} B_{2}+\ldots+a_{m} B_{m}
$$

as a linear combination of B-splines. The B-spline routines will consistently make use of the following identifiers for FORTRAN variables:

$$
\begin{aligned}
& a=\mathrm{BSCOEF} \\
& \mathbf{t}=\mathrm{XKNOT} \\
& m=\mathrm{NCOEF} \\
& M=\mathrm{NKNOT}
\end{aligned}
$$

A B-spline is a particularly compact pp function. $B_{i}$ is a nonnegative function that is nonzero only on the interval $\left[t_{i}, t_{i+k}\right]$. More precisely, the support of the $i$-th B-spline is $\left[t_{i}, t_{i}+k\right]$. No pp function in the same class (other than the zero function) has smaller support (i.e., vanishes on more intervals) than a B-spline. This makes B-splines particularly attractive basis functions since the influence of any particular B-spline coefficient extends only over a few intervals. When it is necessary to emphasize the dependence of the B-spline on its parameters, we will use the notation

$$
B_{i, k, t}
$$

to denote the $i$-th B-spline of order $k$ for the knot sequence $\mathbf{t}$.


Figure 3-1 Spline Interpolants of the Same Data

## Cubic Splines

Cubic splines are smooth (i.e., $C^{l}$ or $C^{2}$ ) fourth-order pp functions. For historical and other reasons, cubic splines are the most heavily used pp functions. Therefore, we provide special routines for their construction and evaluation. The routines for their determination use yet another representation (in terms of value and slope at all the breakpoints) but output the pp representation as described above for general pp functions.
We provide seven cubic spline interpolation routines: CSIEZ, CSINT, CSDEC, CSHER, CSAKM, CSCON, and CSPER. The first routine, CSIEZ, is an easy-to-use version of CSINT coupled with CSVAL. The routine CSIEZ will compute the value of the cubic spline interpolant (to given data using the 'not-a-knot' criterion) on a grid. The routine CSDEC allows the user to specify various endpoint conditions (such as the value of the first or second derivative at the right and left points). This means that the natural cubic spline can be obtained using this routine by setting the second derivative to zero at both endpoints. If function values and derivatives are available, then the

Hermite cubic interpolant can be computed using CSHER. The two routines CSAKM and CSCON are designed so that the shape of the curve matches the shape of the data. In particular, CSCON preserves the convexity of the data while CSAKM attempts to minimize oscillations. If the data is periodic, then CSPER will produce a periodic interpolant. The routine CONFT allows the user wide latitude in enforcing shapes. This routine returns the B-spline representation.

It is possible that the cubic spline interpolation routines will produce unsatisfactory results. The adventurous user should consider using the B-spline interpolation routine BSINT that allows one to choose the knots and order of the spline interpolant.

In Figure 3-1, we display six spline interpolants to the same data. This data can be found in Example 1 of the IMSL routine CSCON Notice the different characteristics of the interpolants. The interpolation routines CSAKM and CSCON are the only two that attempt to preserve the shape of the data. The other routines tend to have extraneous inflection points, with the piecewise quartic $(k=5)$ exhibiting the most oscillation.

## Tensor Product Splines

The simplest method of obtaining multivariate interpolation and approximation routines is to take univariate methods and form a multivariate method via tensor products. In the case of two-dimensional spline interpolation, the development proceeds as follows: Let $\mathbf{t}_{x}$ be a knot sequence for splines of order $k_{x}$, and $\mathbf{t}_{y}$ be a knot sequence for splines of order $k_{y}$. Let $N_{x}+k_{x}$ be the length of $\mathbf{t}_{x}$, and $N_{y}+k_{y}$ be the length of $\mathbf{t}_{y}$. Then, the tensor product spline has the form

$$
\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}(x) B_{m, k_{y}, \mathbf{t}_{y}}(y)
$$

Given two sets of points

$$
\left\{x_{i}\right\}_{i=1}^{N_{x}} \text { and }\left\{y_{i}\right\}_{i=1}^{N_{y}}
$$

for which the corresponding univariate interpolation problem could be solved, the tensor product interpolation problem becomes: Find the coefficients $c_{n m}$ so that

$$
\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right) B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{i}\right)=f_{i j}
$$

This problem can be solved efficiently by repeatedly solving univariate interpolation problems as described in de Boor (1978, page 347). Three-dimensional interpolation has analogous behavior. In this chapter, we provide routines that compute the two-dimensional tensorproduct spline coefficients given two-dimensional interpolation data (BS2IN), compute the three-dimensional tensor-product spline coefficients given three-dimensional interpolation data (BS3IN) compute the two-dimensional tensor-product spline coefficients for a tensor-product least squares problem (BSLS2), and compute the three-dimensional tensor-product spline coefficients for a tensor-product least squares problem (BSLS3). In addition, we provide evaluation, differentiation, and integration routines for the twoand three-dimensional tensor-product spline functions. The relevant routines are BS2VL, BS3VL, BS2DR, BS3DR, BS2GD, BS3GD, BS2IG, and BS3IG.

## Quadratic Interpolation

The routines that begin with the letters " $Q D$ " in this chapter are designed to interpolate a one-, two-, or three-dimensional (tensor product) table of values and return an approximation to the value of the underlying function or one of its derivatives at a given point. These routines are all based on quadratic polynomial interpolation.

## Scattered Data Interpolation

We have one routine, SURF, that will return values of an interpolant to scattered data in the plane. This routine is based on work by Akima (1978), which utilizes $C^{l}$ piecewise quintics on a triangular mesh.

## Least Squares

Routines are provided to smooth noisy data: regression using linear polynomials (RLINE), regression using arbitrary polynomials (RCURV), and regression using user-supplied functions (FNLSQ). Additional routines compute the least-squares fit using splines with fixed knots (BSLSQ) or free knots (BSVLS). These routines can produce cubic-spline least-squares fit simply by setting the order to 4 . The routine CONFT computes a fixed-knot spline weighted least-squares fit subject to linear constraints. This routine is very general and is recommended if issues of shape are important. The two- and three-dimensional tensor-product spline regression routines are (BSLS2) and (BSLS3).

## Smoothing by Cubic Splines

Two "smoothing spline" routines are provided. The routine CSSMH returns the cubic spline that smooths the data, given a smoothing parameter chosen by the user. Whereas, CSSCV estimates the smoothing parameter by cross-validation and then returns the cubic spline that smooths the data. In this sense, CSSCV is the easier of the two routines to use. The routine CSSED returns a smoothed data vector approximating the values of the underlying function when the data are contaminated by a few random spikes.

## Rational Chebyshev Approximation

The routine RATCH computes a rational Chebyshev approximation to a user-supplied function. Since polynomials are rational functions, this routine can be used to compute best polynomial approximations.

## Using the Univariate Spline Routines

An easy to use spline interpolation routine CSIEZ is provided. This routine computes an interpolant and returns the values of the interpolant on a user-supplied grid. A slightly more advanced routine SPLEZ computes (at the users discretion) one of several interpolants or least-squares fits and returns function values or derivatives on a user-supplied grid.
For more advanced uses of the interpolation (or least squares) spline routines, one first forms an interpolant from interpolation (or least-squares) data. Then it must be evaluated, differentiated, or integrated once the interpolant has been formed. One way to perform these tasks, using cubic
splines with the 'not-a-knot' end condition, is to call CSINT to obtain the local coefficients of the piecewise cubic interpolant and then call CSVAL to evaluate the interpolant. A more complicated situation arises if one wants to compute a quadratic spline interpolant and then evaluate it (efficiently) many times. Typically, the sequence of routines called might be BSNAK (get the knots), BSINT (returns the B-spline coefficients of the interpolant), BSCPP (convert to pp form), and PPVAL (evaluate). The last two calls could be replaced by a call to the B -spline grid evaluator BS1GD, or the last call could be replaced with pp grid evaluator PP1GD. The interconnection of the spline routines is summarized in Figure 3-2.


Figure 3- 2 Interrelation of the Spline Routines

## Choosing an Interpolation Routine

The choice of an interpolation routine depends both on the type of data and on the use of the interpolant. We provide 18 interpolation routines. These routines are depicted in a decision tree in Figure 3-3. This figure provides a guide for selecting an appropriate interpolation routine. For example, if periodic one-dimensional (univariate) data is available, then the path through univariate to periodic leads to the IMSL routine CSPER, which is the proper routine for this setting. The general-purpose univariate interpolation routines can be found in the box beginning with CSINT. Two- and three-dimensional tensor-product interpolation routines are also provided. For two-dimensional scattered data, the appropriate routine is SURF .


Figure 3- 3 Choosing an Interplation Routine

## SPLINE_CONSTRAINTS

This function returns the derived type array result, ?_SPLINE_CONSTRAINTS, given optional input. There are optional arguments for the derivative index, the value applied to the spline, and the periodic point for any periodic constraint.

The function is used, for entry number $j$,

```
? SPLINE_CONSTRAINTS (J) = \&
    SPLINE_CONSTRAINTS ([DERIVATIVE=DERIVATIVE_INDEX, ] \&
    POINT = WHERE_APPLIED, [VALUE=VALUE_APPLIED,], \&
    TYPE = CONSTRĀINT_INDICATOR, \&
    [PERIODIC_POINT = VALUE_APPLIED])
```

The square brackets enclose optional arguments. For each constraint either (but not both) the 'VALUE $=$ ' or the 'PERIODIC_POINT $=$ ' optional arguments must be present.

## Required Arguments

POINT = WHERE_APPLIED (Input)
The point in the data interval where a constraint is to be applied.
TYPE = CONSTRAINT_INDICATOR (Input)
The indicator for the type of constraint the spline function or its derivatives is to satisfy at the point: where_applied. The choices are the character strings ${ }^{\prime}==$ ', $\quad<=$ ', $\quad>=$ ', $\quad .=. '$, and '. $=-$ '. They respectively indicate that the spline value or its derivatives will be equal to, not greater than, not less than, equal to the value of the spline at another point, or equal to the negative of the spline value at another point. These last two constraints are called periodic and negative-periodic, respectively. The alternate independent variable point is value_applied for either periodic constraint. There is a use of periodic constraints in .

## Optional Arguments

DERIVATIVE = DERIVATIVE_INDEX (Input)
This is the number of the derivative for the spline to apply the constraint. The value 0 corresponds to the function, the value 1 to the first derivative, etc. If this argument is not present in the list, the value 0 is substituted automatically. Thus a constraint without the derivative listed applies to the spline function.

PERIODIC_POINT = VALUE_APPLIED
This optional argument improves readability by automatically identifying the second independent variable value for periodic constraints.

## FORTRAN 90 Interface

Generic: CALL SPLINE_CONSTRAINTS (POINT, TYPE [, ...])
Specific: The specific interface names are S_SPLINE_CONSTRAINTS and D_SPLINE_CONSTRAINTS.

## SPLINE_VALUES

This rank-1 array function returns an array result, given an array of input. Use the optional argument for the covariance matrix when the square root of the variance function is required. The result will be a scalar value when the input variable is scalar.

## Required Arguments

DERIVATIVE = DERIVATIVE (Input)
The index of the derivative evaluated. Use non-negative integer values. For the function itself use the value 0 .

VARIABLES = VARIABLES (Input)
The independent variable values where the spline or its derivatives are evaluated. Either a rank-1 array or a scalar can be used as this argument.

KNOTS = KNOTS (Input)
The derived type ?_spline_knots, defined as the array COEFFS was obtained with the function SPLINE_FITTING. This contains the polynomial spline degree and the number of knots and the knots themselves for this spline function.

COEFFS $=\mathrm{C}$ (Input)
The coefficients in the representation for the spline function,

$$
f(x)=\sum_{j=1}^{N} c_{j} B_{j}(x) .
$$

These result from the fitting process or array assignment C=SPLINE FITTING (...), defined below. The value $N=\operatorname{size}(\bar{C})$ satisfies the identity
$N-1+$ spline_degree $=$ size $\left(? \_k n o t s\right)$, where the two right-most quantities refer to components of the argument knots.

## Optional Arguments

COVARIANCE $=G$ (Input)
This argument, when present, results in the evaluation of the square root of the variance function

$$
e(x)=\left(b(x)^{T} G b(x)\right)^{1 / 2}
$$

where

$$
b(x)=\left[B_{1}(x), \ldots, B_{N}(x)\right]^{T}
$$

and $G$ is the covariance matrix associated with the coefficients of the spline

$$
c=\left[c_{1}, \ldots, c_{N}\right]^{T}
$$

The argument G is an optional output parameter from the function SPLINE_FITTING, described below. When the square root of the variance function is computed, the arguments DERIVATIVE and C are not used.

IOPT = IOPT (Input)
This optional argument, of derived type ?_options, is not used in this release.

## FORTRAN 90 Interface

Generic: CALL SPLINE_VALUES (DERIVATIVE, VARAIBLES, KNOTS, COEFFS [, ...])
Specific: The specific interface names are S_SPLINE_VALUES and D_SPLINE_VALUES.

## SPLINE_FITTING

Weighted least-squares fitting by B-splines to discrete One-Dimensional data is performed.
Constraints on the spline or its derivatives are optional. The spline function

$$
f(x)=\sum_{j=1}^{N} c_{j} B_{j}(x)
$$

its derivatives, or the square root of its variance function are evaluated after the fitting.

## Required Arguments

DATA $=\operatorname{DATA}(1: 3,:) \quad$ (Input/Output)
An assumed-shape array with size (data, 1 ) $=3$. The data are placed in the array: $\operatorname{data}(1, i)=x_{i}$, data $(2, i)=y_{i}$, and data $(3, i)=\sigma_{i}, i=1, \ldots, n d a t a$. If the variances are not known but are proportional to an unknown value, users may set data $(3, i)=1, i=1, \ldots$, ndata .

KNOTS = KNOTS (Input)
A derived type, ?_spline_knots, that defines the degree of the spline and the breakpoints for the data fitting interval.

## Optional Arguments

CONSTRAINTS = SPLINE_CONSTRAINTS (Input)
A rank-1 array of derived type ?_spline_constraints that give constraints the output spline is to satisfy.

COVARIANCE = G (Output)
An assumed-shape rank-2 array of the same precision as the data. This output is the covariance matrix of the coefficients. It is optionally used to evaluate the square root of the variance function.

IOPT = IOPT(:) (Input/Output)
Derived type array with the same precision as the input array; used for passing optional data to SPLINE_FITTING. The options are as follows:

| Packaged Options for SPLINE_FITTING |  |  |
| :---: | :--- | :---: |
| Prefix $=$ None | Option Name | Option Value |
|  | SPLINE_FITTING_TOL_EQUAL | 1 |
|  | SPLINE_FITTING_TOL_LEAST | 2 |

IOPT (IO) = ? OPTIONS (SPLINE_FITTING_TOL_EQUAL, ? VALUE)
This resets the value for determining that equality constraint equations are rankdeficient. The default is ? $\quad$ value $=10^{-4}$.

```
IOPT(IO) = ?_OPTIONS(SPLINE_FITTING_TOL_LEAST, ?_VALUE)
```

This resets the value for determining that least-squares equations are rank-deficient.
The default is ? value $=10^{-4}$.

## FORTRAN 90 Interface

Generic: CALL SPLINE_FITTING (DATA, KNOTS [, ...])
Specific: The specific interface names are S_SPLINE_FITTING and D_SPLINE_FITTING.

## Description

This routine has similar scope to CONFT found in IMSL (2003, pp 734-743). We provide the square root of the variance function, but we do not provide for constraints on the integral of the spline. The least-squares matrix problem for the coefficients is banded, with band-width equal to the spline order. This fact is used to obtain an efficient solution algorithm when there are no constraints. When constraints are present the routine solves a linear-least squares problem with equality and inequality constraints. The processed least-squares equations result in a banded and upper triangular matrix, following accumulation of the spline fitting equations. The algorithm used for solving the constrained least-squares system will handle rank-deficient problems. A set of reference are available in Hanson (1995) and Lawson and Hanson (1995). The Conft routine uses $Q P R O G$ (loc cit., p. 959), which requires that the least-squares equations be of full rank.

## Fatal and Terminal Error Messages

See the messages.gls file for error messages for SPLINE_FITTING. These error messages are numbered 1340-1367.

## Example 1: Natural Cubic Spline Interpolation to Data

The function

$$
g(x)=\exp \left(-x^{2} / 2\right)
$$

is interpolated by cubic splines on the grid of points

$$
x_{i}=(i-1) \Delta x, i=1, \ldots, \text { ndata }
$$

Those natural conditions are

$$
f\left(x_{i}\right)=g\left(x_{i}\right), i=0, \ldots, \text { ndata } ; \frac{d^{2} f}{d x^{2}}\left(x_{i}\right)=\frac{d^{2} g}{d x^{2}}\left(x_{i}\right), i=0 \text { and ndata }
$$

Our program checks the term const. appearing in the maximum truncation error term

$$
\text { error } \approx \text { const } . \times \Delta x^{4}
$$

at a finer grid.
USE spline_fitting_int
USE show_int
USE norm_int
implicit none
! This is Example 1 for SPLINE_FITTING, Natural Spline
! Interpolation using cubic splines. Use the function
! $\exp (-x * * 2 / 2)$ to generate samples.
integer : : i
integer, parameter $:$ : ndata=24, nord=4, ndegree=nord-1, \&
nbkpt=ndata+2*ndegree, ncoeff=nbkpt-nord, nvalues=2*ndata
real(kind(1e0)), parameter : : zero=0e0, one=1e0, half=5e-1
real(kind (1e0)), parameter : : delta_x=0.15, delta_xv=0.4*delta_x
real(kind(le0)), target : : xdata(ndāta), ydata(ndāta), \&
spline_data (3, ndata), bkpt(nbkpt), \&
ycheck(nvalues), coeff(ncoeff), \&
xvalues(nvalues), yvalues(nvalues), diffs
real(kind(le0)), pointer : : pointer_bkpt(:)
type (s_spline_knots) break_points
type (s_spline_constraints) constraints (2)
xdata $=\left(/\left((i-1) * d e l t a \_x, \quad i=1\right.\right.$, ndata $\left.) /\right)$
ydata $=\exp (-$ half*xdat $\bar{a} * * 2)$
xvalues $=\left(/\left(0.03+(i-1) * d e l t a \_x v, i=1, n v a l u e s\right) /\right)$
ycheck $=\exp (-h a l f * x v a l u e s * * 2)$
spline_data $(1,:)=x d a t a$
spline_data $(2,:)=y d a t a$
spline_data $(3,:)=$ one
! Define the knots for the interpolation problem.
bkpt(1:ndegree) $=\left(/\left(i * d e l t a \_x, i=-n d e g r e e,-1\right) /\right)$
bkpt (nord:nbkpt-ndegree) $=x \bar{d} a t a$
bkpt (nbkpt-ndegree+1:nbkpt) = \&
(/(xdata(ndata) +i*delta_x, $i=1$, ndegree) /)
! Assign the degree of the polynomial and the knots. pointer_bkpt $=>$ bkpt
break_points=s_spline_knots(ndegree, pointer_bkpt)

```
! These are the natural conditions for interpolating cubic
! splines. The derivatives match those of the interpolating
! function at the ends.
    constraints(1)=spline_constraints &
        (derivative=2, poin}t=bkpt(nord), type='==', value=-one
    constraints(2)=spline_constraints &
        (derivative=2,point=bkpt(nbkpt-ndegree), type= '==', &
        value=(-one+xdata(ndata)**2)*ydata(ndata))
    coeff = spline_fitting(data=spline_data, knots=break_points,&
            constraints=constraints)
    yvalues=spline_values(0, xvalues, break_points, coeff)
    diffs=norm(yvalues-ycheck,huge(1))/delta_x**nord
    if (diffs <= one) then
        write(*,*) 'Example 1 for SPLINE_FITTING is correct.'
    end if
    end
```


## Output

Example 1 for SPLINE FITTING is correct.

## Additional Examples

## Example 2: Shaping a Curve and its Derivatives

The function

$$
g(x)=\exp \left(-x^{2} / 2\right)(1+\text { noise })
$$

is fit by cubic splines on the grid of equally spaced points

$$
x_{i}=(i-1) \Delta x, i=1, \ldots, n d a t a
$$

The term noise is uniform random numbers from the normalized interval $[-\tau, \tau]$, where $\tau=0.01$. The spline curve is constrained to be convex down for for $0 \leq \mathrm{x} \leq 1$ convex upward for $1<x \leq 4$, and have the second derivative exactly equal to the value zero at $x=1$. The first derivative is constrained with the value zero at $x=0$ and is non-negative at the right and of the interval, $x=4$. A sample table of independent variables, second derivatives and square root of variance function values is printed.

```
use spline_fitting_int
use show int
use rand_int
use norm_int
implicit none
! This is Example 2 for SPLINE_FITTING. Use 1st and 2nd derivative
! constraints to shape the splines.
```

```
    integer :: i, icurv
    integer, parameter :: nbkptin=13, nord=4, ndegree=nord-1, &
        nbkpt=nbkptin+2*ndegree, ndata=21, ncoeff=nbkpt-nord
    real(kind(1e0)), parameter :: zero=0e0, one=1e0, half=5e-1
    real(kind(1e0)), parameter :: range=4.0, ratio=0.02, tol=ratio*half
    real(kind(1e0)), parameter :: delta_x=range/(ndata-1), &
    delta_b=range/(nbkptin-1)
    real(kind(1e0)), target :: xdata(ndata), ydata(ndata), ynoise(ndata),&
        sddata(ndata), spline data (3, ndata), bkpt(nbkpt), &
        values(ndata), derivat1(ndata), derivat2(ndata), &
        coeff(ncoeff), root variance(ndata), diffs
    real(kind(1e0)), dimension(ncoeff,ncoeff) :: sigma_squared
    real(kind(1e0)), pointer :: pointer_bkpt(:)
    type (s_spline_knots) break_points
    type (s_spline_constraints) constraints(nbkptin+2)
    xdata = (/((i-1)*delta_x, i=1,ndata)/)
    ydata = exp(-half*xdata**2)
    ynoise = ratio*ydata*(rand(ynoise) -half)
    ydata = ydata+ynoise
    sddata = ynoise
    spline_data(1,:)=xdata
    spline_data(2,:)=ydata
    spline_data(3,:)=sddata
    bkpt=(/((i-nord)*delta_b, i=1,nbkpt)/)
! Assign the degree of the polynomial and the knots.
    pointer bkpt => bkpt
    break_points=s_spline_knots(ndegree, pointer_bkpt)
    icurv=int(one/delta_b)+1
! At first shape the curve to be convex down.
    do i=1,icurv-1
        constraints(i)=spline_constraints &
(derivative=2, point=bkpt(i+ndegree), type='<=', value=zero)
    end do
! Force a curvature change.
    constraints(icurv)=spline_constraints &
(derivative=2, point=bkpt(icurv+ndegree), type='==', value=zero)
! Finally, shape the curve to be convex up.
    do i=icurv+1,nbkptin
        constraints(i)=spline_constraints &
(derivative=2, point=bkpt(i+ndegree), type='>=', value=zero)
    end do
! Make the slope zero and value non-negative at right.
    constraints(nbkptin+1)=spline_constraints &
(derivative=1, point=bkpt(nord), type='==', value=zero)
    constraints(nbkptin+2)=spline_constraints &
(derivative=0, point=bkpt(nbkptin+\overline{n}degree), type='>=', value=zero)
```

```
    coeff = spline_fitting(data=spline_data, knots=break_points, &
    constraints=constraints, covariance=sigma_squared)
! Compute value, first two derivatives and the variance.
    values=spline values(0, xdata, break points, coeff)
    root_variance=spline_values(0, xdata, break_points, coeff, &
    covariance=sigma_squared)
    derivatl=spline_values(1, xdata, break_points, coeff)
    derivat2=spline_values(2, xdata, break_points, coeff)
    call show(reshape((/xdata, derivat2, root_variance/),(/ndata,3/)),&
"The x values, 2-nd derivatives, and square root of variance.")
! See that differences are relatively small and the curve has
! the right shape and signs.
    diffs=norm(values-ydata)/norm(ydata)
    if (all(values > zero) .and. all(derivatl < epsilon(zero))&
        .and. diffs <= tol) then
        write(*,*) 'Example 2 for SPLINE_FITTING is correct.'
    end if
    end
```


## Output

Example 2 for SPLINE FITTING is correct.

## Example 3: Splines Model a Random Number Generator

The function

$$
\begin{aligned}
g(x) & =\exp \left(-x^{2} / 2\right),-1<x<1 \\
& =0,|x| \geq 1
\end{aligned}
$$

is an unnormalized probability distribution. This function is similar to the standard Normal distribution, with specific choices for the mean and variance, except that it is truncated. Our algorithm interpolates $g(x)$ with a natural cubic spline, $f(x)$. The cumulative distribution is approximated by precise evaluation of the function

$$
q(x)=\int_{-1}^{x} f(t) d t
$$

Gauss-Legendre quadrature formulas, $\operatorname{IMSL}$ (1994, pp. 621-626), of order two are used on each polynomial piece of $f(t)$ to evaluate $q(x)$ cheaply. After normalizing the cubic spline so that $q(1)=1$, we may then generate random numbers according to the distribution $f(x) \cong g(x)$. The values of $x$ are evaluated by solving $q(x)=u,-1<x<1$. Here $u$ is a uniform random sample. Newton's method, for a vector of unknowns, is used for the solution algorithm. Recalling the relation

$$
\frac{d}{d x}(q(x)-u)=f(x),-1<x<1
$$

we believe this illustrates a method for generating a vector of random numbers according to a continuous distribution function having finite support.

```
    use spline_fitting_int
    use linear_operators
    use Numerical_Libraries
    implicit none
! This is Example 3 for SPLINE_FITTING. Use splines to
! generate random (almost normal) numbers. The normal distribution
! function has support (-1,+1), and is zero outside this interval.
! The variance is 0.5.
    integer i, niterat
        integer, parameter :: iweight=1, nfix=0, nord=4, ndata=50
        integer, parameter :: nquad=(nord+1)/2, ndegree=nord-1
        integer, parameter :: nbkpt=ndata+2*ndegree, ncoeff=nbkpt-nord
        integer, parameter :: last=nbkpt-ndegree, n_samples=1000
        integer, parameter :: limit=10
real(kind(le0)), dimension(n_samples) :: fn, rn, x, alpha_x, beta_x
    INTEGER LEFT_OF(n_samples)
    real(kind(le0)), pārameter : : one=1e0, half=5e-1, zero=0e0, two=2e0
    real(kind(1e0)), parameter :: delta_x=two/(ndata-1)
    real(kind(le0)), parameter :: qalpha=zero, qbeta=zero, domain=two
    real(kind(le0)) qx(nquad), qxi(nquad), qw(nquad), qxfix(nquad)
    real(kind(1e0)) alpha_, beta_, quad(0:ndata-1)
    real(kind(le0)), target :: xdata(ndata), ydata(ndata),&
    coeff(ncoeff), spline_data(3, ndata), bkpt(nbkpt)
    real(kind(le0)), pointer :: pointer bkpt(:)
    type (s_spline_knots) break_points
    type (s_spline_constraints) constraints(2)
! Approximate the probability density function by splines.
    xdata = (/(-one+(i-1)*delta_x, i=1,ndata)/)
    ydata = exp(-half*xdata**2)
    spline_data(1,:)=xdata
    spline_data(2,:)=ydata
    spline_data(3,:)=one
    bkpt=(/(-one+(i-nord)*delta_x, i=1,nbkpt)/)
! Assign the degree of the polynomial and the knots.
    pointer_bkpt => bkpt
    break_points=s_spline_knots(ndegree, pointer_bkpt)
! Define the natural derivatives constraints:
    constraints(1)=spline_constraints &
        (derivative=2, poin\overline{t}=bkpt(nord), type='==', &
        value=(-one+xdata(1)**2)*ydata(1))
        constraints(2)=spline_constraints &
            (derivative=2, poin\overline{t}=bkpt(last), type='==', &
            value=(-one+xdata(ndata)**2)*ydata(ndata))
```

```
! Obtain the spline coefficients.
    coeff=spline_fitting(data=spline_data, knots=break_points,&
    constraints=\overline{constraints)}
! Compute the evaluation points 'qx(*)' and weights 'qw(*)' for
! the Gauss-Legendre quadrature. This will give a precise
! quadrature for polynomials of degree <= nquad*2.
    call gqrul(nquad, iweight, qalpha, qbeta, nfix, qxfix, qx, qw)
! Compute pieces of the accumulated distribution function:
    quad(0)=zero
    do i=1, ndata-1
            alpha_= (bkpt(nord+i)-bkpt(ndegree+i))*half
            beta_= (bkpt(nord+i)+bkpt(ndegree+i))*half
! Normalized abscissas are stretched to each spline interval.
! Each polynomial piece is integrated and accumulated.
            qxi = alpha_*qx+beta
            quad(i) = sum(qw*splinne_values(0, qxi, break_points,&
        coeff))*alpha_&
                + quad(i-1)
            end do
! Normalize the coefficients and partial integrals so that the
! total integral has the value one.
        coeff=coeff/quad(ndata-1); quad=quad/quad(ndata-1)
        rn=rand(rn)
        x=zero; niterat=0
    solve_equation: do
! Find the intervals where the x values are located.
        LEFT_OF=NDEGREE; I=NDEGREE
            do
                I=I+1; if(I >= LAST) EXIT
                WHERE(x >= BKPT(I)) LEFT_OF = LEFT_OF+1
            end do
! Use Newton's method to solve the nonlinear equation:
! accumulated_distribution_function - random_number = 0.
    alpha_x = (x-b\overline{kpt}(LEFT_OF))*half
    beta_x = (x+bkpt(LEFT_OF))*half
    FN=Q\overline{UAD}(LEFT OF-NORD) - \overline{R}N
        DO I=1,NQUAD
            FN=FN+QW(I)*spline values(0, alpha x*QX(I)+beta x,&
                        break_points,
        END DO
! This is the Newton method update step:
        x=x-fn/spline_values(0, x, break_points, coeff)
        niterat=niterāt+1
! Constrain the values so they fall back into the interval.
! Newton's method may give approximates outside the interval.
```

```
        where(x <= -one .or. x >= one) x=zero
        if(norm(fn,1) <= sqrt(epsilon(one))*norm(x,1))&
    exit solve_equation
end do solve_equation
! Check that Newton's method converges.
    if (niterat <= limit) then
        write (*,*) 'Example 3 for SPLINE_FITTING is correct.'
        end if
end
```


## Output

Example 3 for SPLINE_FITTING is correct.

## Example 4: Represent a Periodic Curve

The curve tracing the edge of a rectangular box, traversed in a counter-clockwise direction, is parameterized with a spline representation for each coordinate function, $(x(t), y(t))$. The functions are constrained to be periodic at the ends of the parameter interval. Since the perimeter arcs are piece-wise linear functions, the degree of the splines is the value one. Some breakpoints are chosen so they correspond to corners of the box, where the derivatives of the coordinate functions are discontinuous. The value of this representation is that for each $t$ the splines representing $(x(t), y(t))$ are points on the perimeter of the box. This "eases" the complexity of evaluating the edge of the box. This example illustrates a method for representing the edge of a domain in two dimensions, bounded by a periodic curve.

```
    use spline_fitting_int
    use norm_int
    implicit none
! This is Example 4 for SPLINE_FITTING. Use piecewise-linear
! splines to represent the perimeter of a rectangular box.
    integer i, j
    integer, parameter :: nbkpt=9, nord=2, ndegree=nord-1, &
        ncoeff=nbkpt-nord, ndata=7, ngrid=100, &
        nvalues=(ndata-1)*ngrid
    real(kind(1e0)), parameter :: zero=0e0, one=1e0
    real(kind(le0)), parameter :: delta_t=one, delta_b=one, delta_v=0.01
    real(kind(1e0)) delta_x, delta_y
    real(kind(le0)), dimension(ndata) :: sddata=one, &
! These are redundant coordinates on the edge of the box.
            xdata=(/0.0, 1.0, 2.0, 2.0, 1.0, 0.0, 0.0/), &
            ydata=(/0.0, 0.0, 0.0, 1.0, 1.0, 1.0, 0.0/)
real(kind(1e0)) tdata(ndata), xspline_data(3, ndata), &
        yspline data(3, ndata), tvalues(nvalues), &
        xvalues(nvalues), yvalues(nvalues), xcoeff(ncoeff), &
        ycoeff(ncoeff), xcheck(nvalues), ycheck(nvalues), diffs
real(kind(le0)), target :: bkpt(nbkpt)
```

```
    real(kind(1e0)), pointer :: pointer_bkpt(:)
    type (s_spline_knots) break_points
    type (s_spline_constraints) constraints(1)
    tdata = (/((i-1)*delta t, i=1,ndata)/)
    xspline data(1,:)=tdatā; yspline data(1,:)=tdata
    xspline_data(2,:)=xdata; yspline_data(2,:)=ydata
    xspline_data(3,:)=sddata; yspline_data(3,:)=sddata
    bkpt(nord:nbkpt-ndegree)=(/((i-nord)*delta_b, &
                            i=nord, nbkpt-ndegree)/)
! Collapse the outside knots.
    bkpt(1:ndegree)=bkpt(nord)
    bkpt (nbkpt-ndegree+1:nbkpt)=bkpt (nbkpt-ndegree)
! Assign the degree of the polynomial and the knots.
    pointer_bkpt => bkpt
    break_points=s_spline_knots(ndegree, pointer_bkpt)
! Make the two parametric curves also periodic.
    constraints(1)=spline_constraints &
        (derivative=0, poin\overline{t}=\textrm{bkpt(nord), type='.=.', &}
        value=bkpt (nbkpt-ndegree))
    xcoeff = spline_fitting(data=xspline_data, knots=break_points, &
                constraints=constraints)
    ycoeff = spline_fitting(data=yspline_data, knots=break_points, &
                                    constraints=constraints)
! Use the splines to compute the coordinates of points along the perimeter.
! Compare them with the coordinates of the edge points.
    tvalues= (/((i-1)*delta_v, i=1,nvalues)/)
    xvalues=spline values(0', tvalues, break points, xcoeff)
    yvalues=spline_values(0, tvalues, break_points, ycoeff)
    do i=1, nvalues
        j=(i-1)/ngrid+1
        delta_x=(xdata(j+1)-xdata(j))/ngrid
    delta_y=(ydata(j+1) -ydata(j))/ngrid
        xcheck(i)=xdata(j)+mod(i+ngrid-1,ngrid)*delta x
        ycheck(i)=ydata(j)+mod(i+ngrid-1,ngrid)*delta_y
    end do
    diffs=norm(xvalues-xcheck,1)/norm(xcheck,1)+&
            norm(yvalues-ycheck,1)/norm(ycheck,1)
    if (diffs <= sqrt(epsilon(one))) then
        write(*,*) 'Example 4 for SPLINE_FITTING is correct.'
    end if
    end
```


## Output

Example 4 for SPLINE_FITTING is correct.

## SURFACE_CONSTRAINTS

To further shape a surface defined by a tensor product of B-splines, the routine SURFACE_FITTING will least squares fit data with equality, inequality and periodic constraints. These can apply to the surface function or its partial derivatives. Each constraint is packaged in the derived type ?_SURFACE_CONSTRAINTS. This function uses the data consisting of: the place where the constraint is to hold, the partial derivative indices, and the type of the constraint. This object is returned as the derived type function result ? SURFACE_CONSTRAINTS. The function itself has two required and two optional arguments. In a list of constraints, the $j$-th item will be:

```
?_SURFACE_CONSTRAINTS(j) = &
SÜRFACE_CO}NSTRAINTS &
    ([DERIVATIVE=DERIVATIVE_INDEX(1:2),] &
    POINT = WHERE_APPLIED(1:2),[VALUE=VALUE_APPLIED,],&
    TYPE = CONSTRAINT_INDICATOR, &
    [PERIODIC_POINT = PERIODIC_POINT (1:2)])
```

The square brackets enclose optional arguments. For each constraint the arguments 'value =' and 'PERIODIC_POINT =' are not used at the same time.

## Required Arguments

POINT $=$ WHERE_APPLIED (Input)
The point in the data domain where a constraint is to be applied. Each point has an $x$ and $y$ coordinate, in that order.

TYPE = CONSTRAINT_INDICATOR (Input)
The indicator for the type of constraint the tensor product spline function or its partial derivatives is to satisfy at the point: where_applied. The choices are the character strings '==', '<=', '>=', '.=.', and '.=-'. They respectively indicate that the spline value or its derivatives will be equal to, not greater than, not less than, equal to the value of the spline at another point, or equal to the negative of the spline value at another point. These last two constraints are called periodic and negative-periodic, respectively.

## Optional Arguments

## DERIVATIVE = DERIVATIVE_INDEX (1:2) (Input)

These are the number of the partial derivatives for the tensor product spline to apply the constraint. The array ( $/ 0,0 /$ ) corresponds to the function, the value $(/ 1,0 /)$ to the first partial derivative with respect to $x$, etc. If this argument is not present in the list, the value $(/ 0,0 /)$ is substituted automatically. Thus a constraint without the derivatives listed applies to the tensor product spline function.

PERIODIC = PERIODIC_POINT (1:2)
This optional argument improves readability by identifying the second pair of independent variable values for periodic constraints.

## FORTRAN 90 Interface

Generic: CALL SURFACE_CONSTRAINTS (POINT, TYPE [, ...])
Specific: The specific interface names are s_SURFACE_CONSTRAINTS and D_SURFACE_CONSTRAINTS.

## SURFACE_VALUES

This rank-2 array function returns a tensor product array result, given two arrays of independent variable values. Use the optional input argument for the covariance matrix when the square root of the variance function is evaluated. The result will be a scalar value when the input independent variable is scalar.

## Required Arguments

DERIVATIVE = DERIVATIVE (1:2) (Input)
The indices of the partial derivative evaluated. Use non-negative integer values. For the function itself use the array ( $/ 0,0 /$ ) .

```
VARIABLESX = VARIABLESX (Input)
```

The independent variable values in the first or $x$ dimension where the spline or its derivatives are evaluated. Either a rank-1 array or a scalar can be used as this argument.

VARIABLESY = VARIABLESY (Input)
The independent variable values in the second or $y$ dimension where the spline or its derivatives are evaluated. Either a rank-1 array or a scalar can be used as this argument.

KNOTSX = KNOTSX (Input)
The derived type?_spline_knots, used when the array coeffs (:, :) was obtained with the function SURFACE_FITTING. This contains the polynomial spline degree and the number of knots and the knots themselves, in the $x$ dimension.

KNOTSY = KNOTSY (Input)
The derived type ?_spline_knots, used when the array $\operatorname{coeffs}(:,:)$ was obtained with the function SURFACE_FITTING. This contains the polynomial spline degree and the number of knots and the knots themselves, in the $y$ dimension.

COEFFS $=\mathrm{C}$ (Input)
The coefficients in the representation for the spline function,

$$
f(x, y)=\sum_{j=1}^{N} \sum_{i=1}^{M} c_{i j} B_{i}(y) B_{j}(x)
$$

These result from the fitting process or array assignment C=SURFACE_FITTING (...), defined below.

The values $M=\operatorname{size}(C, 1)$ and $N=$ size $(C, 2)$ satisfies the respective identities $N-1+$ spline_degree $=$ size $(?$ knotsx $)$, and
$M-1+$ spline_degree $=$ size $\left(? \_\right.$knotsy $)$, where the two right-most quantities in both equations refer to components of the arguments knotsx and knotsy. The same value of spline_degree must be used for both knotsx and knotsy.

## Optional Arguments

COVARIANCE $=G$ (Input)
This argument, when present, results in the evaluation of the square root of the variance function

$$
e(x, y)=\left(b(x, y)^{T} G b(x, y)\right)^{1 / 2}
$$

where

$$
b(x, y)=\left[B_{1}(x) B_{1}(y), \ldots, B_{N}(x) B_{1}(y), \ldots\right]^{T}
$$

and $G$ is the covariance matrix associated with the coefficients of the spline

$$
c=\left[c_{11}, \ldots, c_{N 1}, \ldots\right]^{T}
$$

The argument G is an optional output from SURFACE_FITTING, described below. When the square root of the variance function is computed, the arguments DERIVATIVE and C are not used.

IOPT = IOPT (Input)
This optional argument, of derived type ?_options, is not used in this release.

## FORTRAN 90 Interface

Generic: CALL SURFACE_VALUES (DERIVATIVE, VARIABLESX, VARIABLESY, KNOTSX, KNOTSY, COEFFS [,...])

Specific: The specific interface names are S_SURFACE_VALUES and D_SURFACE_VALUES.

## SURFACE FITTING

Weighted least-squares fitting by tensor product B-splines to discrete two-dimensional data is performed. Constraints on the spline or its partial derivatives are optional. The spline function

$$
f(x, y)=\sum_{j=1}^{N} \sum_{i=1}^{M} c_{i j} B_{i}(y) B_{j}(x),
$$

its derivatives, or the square root of its variance function are evaluated after the fitting.

## Required Arguments

DATA $=$ DATA (1:4,:) (Input/Output)
An assumed-shape array with size $($ data, 1$)=4$. The data are placed in the array:

```
data(1,i) = x ,
data(2,i) = y ,
data(3,i) = zi,
data(4,i) = \sigmai,i=1,\ldots,ndata.
```

If the variances are not known, but are proportional to an unknown value, use

```
data(4,i) = 1, i=1,\ldots,ndata.
```

KNOTSX = KNOTSX (Input)
A derived type, ?_SPLINE_KNOTS, that defines the degree of the spline and the breakpoints for the data fitting domain, in the first dimension.

## KNOTSY = KNOTSY (Input)

A derived type, ?_SPLINE_KNOTS, that defines the degree of the spline and the breakpoints for the data fitting domain, in the second dimension.

## Optional Arguments

CONSTRAINTS = SURFACE_CONSTRAINTS (Input)
A rank-1 array of derived type ?_SURFACE_CONSTRAINTS that defines constraints the tensor product spline is to satisfy.

COVARIANCE $=$ G (Output)
An assumed-shape rank-2 array of the same precision as the data. This output is the covariance matrix of the coefficients. It is optionally used to evaluate the square root of the variance function.

IOPT = IOPT(:) (Input/Output)
Derived type array with the same precision as the input array; used for passing optional data to SURFACE_FITTING. The options are as follows:

| Packaged Options for SURFACE_FITTING |  |  |
| :--- | :--- | :---: |
| Prefix = None | Option Name | Option Value |
|  | SURFACE_FITTING_SMALLNESS | 1 |
|  | SURFACE_FITTING_FLATNESS | 2 |
|  | SURFACE_FITTING_TOL_EQUAL | 3 |
|  | SURFACE_FITTING_TOL_LEAST | 4 |
|  | SURFACE_FITTING_RESIDUALS | 5 |


| Packaged Options for SURFACE_FITTING |  |  |
| :--- | :--- | :--- |
|  | SURFACE_FITTING_PRINT | 6 |
|  | SURFACE_FITTING_THINNESS | 7 |

IOPT (IO) $=$ ?_OPTIONS\&
(surface_fitting_smallnes, ?_value)

This resets the square root of the regularizing parameter multiplying the squared integral of the unknown function. The argument ?_value is replaced by the default value. The default is ? value $=0$.

```
IOPT(IO) = ?_OPTIONS &
    (SURFACE_FITTING_FLATNESS, ? VALUE)
```

This resets the square root of the regularizing parameter multiplying the squared integral of the partial derivatives of the unknown function. The argument ?_VALUE is replaced by the default value. The default is ?_VALUE = SQRT (EPSILON(?_VALUE))*SIZE, where

$$
\text { size }=\sum \mid \text { data }(3,:) / \text { data }(4,:) \mid /(\text { ndata }+1) .
$$

$\operatorname{IOPT}(I O)=? \quad O P T I O N S \&$
(SURFACE_FITTING_TOL_EQUAL, ? VALUE)
This resets the value for determining that equality constraint equations are rankdeficient. The default is ?_VALUE $=10^{-4}$.

IOPT(IO) = ?_OPTIONS\&
(SURFACE_FITTING_TOL_LEAST, ?_VALUE)
This resets the value for determining that least-squares equations are rank-deficient. The default is ?_VALUE $=10^{-4}$.

IOPT(IO) = ?_OPTIONS\&
(SURFACE_FITTING_RESIDUALS, DUMMY)
This option returns the residuals $=$ surface - data, in data ( $4,:$ ). That row of the array is overwritten by the residuals. The data is returned in the order of cell processing order, or left-to-right in $x$ and then increasing in $y$. The allocation of a temporary for data ( $1: 4,:$ ) is avoided, which may be desirable for problems with large amounts of data. The default is to not evaluate the residuals and to leave data ( $1: 4,:$ ) as input.

IOPT(IO) = ?_OPTIONS\&
(SURFACE_FITTING_PRINT, DUMMY)
This option prints the knots or breakpoints for $x$ and $y$, and the count of data points in cell processing order. The default is to not print these arrays.

IOPT(IO) = ?_OPTIONS\&
(SURFACE_FITTING_THINNESS, ? VALUE)
This resets the square root of the regularizing parameter multiplying the squared integral of the second partial derivatives of the unknown function. The argument ? VALUE is replaced by the default value. The default is ?_VALUE $=10^{-3} \times$ SIZE, , where

$$
\text { size }=\sum|\operatorname{data}(3,:) / \operatorname{data}(4,:)| /(\text { ndata }+1) .
$$

## FORTRAN 90 Interface

Generic: CALL SURFACE_FITTING (DATA, KNOTSX, KNOTSX, KNOTSY [, ...])
Specific: The specific interface names are S_SURFACE_FITTING and D_SURFACE_FITTING.

## Description

The coefficients are obtained by solving a least-squares system of linear algebraic equations, subject to linear equality and inequality constraints. The system is the result of the weighted data equations and regularization. If there are no constraints, the solution is computed using a banded least-squares solver. Details are found in Hanson (1995).

## Fatal and Terminal Error Messages

See the messages.gls file for error messages for SURFACE_FITTING. These error messages are numbered 1151-1152, 1161-1162, 1370-1393.

## Example 1: Tensor Product Spline Fitting of Data

The function

$$
g(x, y)=\exp \left(-x^{2}-y^{2}\right)
$$

is least-squares fit by a tensor product of cubic splines on the square

$$
[0,2] \otimes[0,2]
$$

There are ndata random pairs of values for the independent variables. Each datum is given unit uncertainty. The grid of knots in both $x$ and $y$ dimensions are equally spaced, in the interior cells, and identical to each other. After the coefficients are computed a check is made that the surface approximately agrees with $g(x, y)$ at a tensor product grid of equally spaced values.

```
    USE surface_fitting_int
    USE rand_int
    USE norm_int
    implicit none
! This is Example 1 for SURFACE_FITTING, tensor product
```

```
! B-splines approximation. Use the function
! exp(-x**2-y**2) on the square (0, 2) x (0, 2) for samples.
! The spline order is "nord" and the number of cells is
! "(ngrid-1)**2". There are "ndata" data values in the square.
    integer :: i
    integer, parameter :: ngrid=9, nord=4, ndegree=nord-1, &
        nbkpt=ngrid+2*ndegree, ndata = 2000, nvalues=100
    real(kind(1d0)), parameter :: zero=0d0, one=1d0, two=2d0
    real(kind(1d0)), parameter :: TOLERANCE=1d-3
    real(kind(1d0)), target :: spline data (4, ndata), bkpt(nbkpt), &
        coeff(ngrid+ndegree-1,ngri\overline{d}+ndegree-1), delta, sizev, &
        x(nvalues), y(nvalues), values(nvalues, nvalues)
    real(kind(1d0)), pointer :: pointer_bkpt(:)
    type (d_spline_knots) knotsx, knotsy
! Generate random (x,y) pairs and evaluate the
! example exponential function at these values.
    spline_data(1:2,:)=two*rand(spline_data(1:2,:))
    spline_data(3,:)=exp(-sum(spline_data(1:2,:)**2,dim=1))
    spline_data(4,:)=one
! Define the knots for the tensor product data fitting problem.
        delta = two/(ngrid-1)
        bkpt(1:ndegree) = zero
        bkpt(nbkpt-ndegree+1:nbkpt) = two
        bkpt(nord:nbkpt-ndegree) = (/ (i*delta,i=0,ngrid-1) /)
! Assign the degree of the polynomial and the knots.
    pointer_bkpt => bkpt
    knotsx=\overline{d}_spline_knots(ndegree, pointer_bkpt)
    knotsy=knotsx
! Fit the data and obtain the coefficients.
    coeff = surface_fitting(spline_data, knotsx, knotsy)
! Evaluate the residual = spline - function
! at a grid of points inside the square.
    delta=two/(nvalues+1)
    x=(/(i*delta,i=1,nvalues)/); y=x
    values=exp(-spread(x**2,1,nvalues)-spread(y**2,2,nvalues))
    values=surface_values((/0,0/), x, y, knotsx, knotsy, coeff)-&
        values
! Compute the R.M.S. error:
    sizev=norm(pack(values, (values == values)))/nvalues
    if (sizev <= TOLERANCE) then
        write(*,*) 'Example 1 for SURFACE_FITTING is correct.'
    end if
    end
```


## Output

```
Example 1 for SURFACE_FITTING is correct.
```


## Additional Examples

## Example 2: Parametric Representation of a Sphere

From Struik (1961), the parametric representation of points $(x, y, z)$ on the surface of a sphere of radius $a>0$ is expressed in terms of spherical coordinates,

$$
\begin{gathered}
x(u, v)=a \cos (u) \cos (v),-\pi \leq 2 u \leq \pi \\
y(u, v)=a \cos (u) \sin (v),-\pi \leq v \leq \pi \\
z(u, v)=a \sin (u)
\end{gathered}
$$

The parameters are radians of latitude ( $u$ ) and longitude ( $v$ ). The example program fits the same ndata random pairs of latitude and longitude in each coordinate. We have covered the sphere twice by allowing

$$
-\pi \leq u \leq \pi
$$

for latitude. We solve three data fitting problems, one for each coordinate function. Periodic constraints on the value of the spline are used for both $u$ and $v$. We could reduce the computational effort by fitting a spline function in one variable for the $z$ coordinate. To illustrate the representation of more general surfaces than spheres, we did not do this. When the surface is evaluated we compute latitude, moving from the South Pole to the North Pole,

$$
-\pi \leq 2 u \leq \pi
$$

Our surface will approximately satisfy the equality

$$
x^{2}+y^{2}+z^{2}=a^{2}
$$

These residuals are checked at a rectangular mesh of latitude and longitude pairs. To illustrate the use of some options, we have reset the three regularization parameters to the value zero, the leastsquares system tolerance to a smaller value than the default, and obtained the residuals for each parametric coordinate function at the data points.

```
USE surface_fitting_int
USE rand_int
USE norm int
USE Numerical_Libraries
implicit none
! This is Example 2 for SURFACE_FITTING, tensor product
! B-splines approximation. Fit x, y, z parametric functions
! for points on the surface of a sphere of radius "A".
! Random values of latitude and longitude are used to generate
! data. The functions are evaluated at a rectangular grid
! in latitude and longitude and checked to lie on the surface
! of the sphere.
integer :: i, j
```

```
    integer, parameter :: ngrid=6, nord=6, ndegree=nord-1, &
            nbkpt=ngrid+2*ndegree, ndata =1000, nvalues=50, NOPT=5
        real(kind(1d0)), parameter :: zero=0d0, one=1d0, two=2d0
        real(kind(1d0)), parameter :: TOLERANCE=1d-2
        real(kind(1d0)), target :: spline_data (4, ndata, 3), bkpt(nbkpt), &
            coeff(ngrid+ndegree-1,ngri\overline{d}+ndegree-1, 3), delta, sizev, &
            pi, A, x(nvalues), y(nvalues), values(nvalues, nvalues), &
            data(4,ndata)
    real(kind(1d0)), pointer :: pointer_bkpt(:)
    type (d spline knots) knotsx, knotsy
    type (d_option\overline{s}) OPTIONS (NOPT)
! Get the constant "pi" and a random radius, > 1.
    pi = DCONST("pi"); A=one+rand(A)
! Generate random (latitude, longitude) pairs and evaluate the
! surface parameters at these points.
    spline_data(1:2,:,1)=pi*(two*rand(spline_data(1:2,:,1))-one)
    spline_data(1:2,:,2)=spline_data(1:2,:,1)
    spline_data(1:2,:,3)=spline_data(1:2,:,1)
! Evaluate x, y, z parametric points.
    spline_data(3,:,1)=A*cos(spline_data(1,:,1))*cos(spline_data(2,:,1))
    spline_data(3,:,2)=A*cos(spline_data(1,:,2))*sin(spline_data(2,:,2))
    spline_data(3,:,3)=A*sin(spline_data(1,:,3))
! The values are equally uncertain.
    spline_data(4,:,:)=one
! Define the knots for the tensor product data fitting problem.
    delta = two*pi/(ngrid-1)
    bkpt(1:ndegree) = -pi
    bkpt(nbkpt-ndegree+1:nbkpt) = pi
    bkpt(nord:nbkpt-ndegree)=(/(-pi+i*delta,i=0,ngrid-1)/)
! Assign the degree of the polynomial and the knots.
    pointer bkpt => bkpt
    knotsx=\overline{d}_spline_knots(ndegree, pointer_bkpt)
    knotsy=knotsx
! Fit a data surface for each coordinate.
! Set default regularization parameters to zero and compute
! residuals of the individual points. These are returned
! in DATA(4,:).
    do j=1,3
        data=spline data(:,:,j)
    OPTIONS(1)=d_options(surface_fitting_thinness,zero)
    OPTIONS(2)=d options(surface fitting flatness,zero)
    OPTIONS (3) =d_options(surface_fitting_smallness,zero)
    OPTIONS(4)=d_options(surface_fitting_tol_least,1d-5)
    OPTIONS(5)=surface_fitting_residuals
    coeff(:,:,j) = surface_fitting(data, knotsx, knotsy,&
        IOPT=OPTIONS)
    end do
```

```
! Evaluate the function at a grid of points inside the rectangle of
! latitude and longitude covering the sphere just once. Add the
! sum of squares. They should equal "A**2" but will not due to
! truncation and rounding errors.
    delta=pi/(nvalues+1)
    x=(/(-pi/two+i*delta,i=1,nvalues)/); y=two*x
    values=zero
    do j=1,3
        values=values+&
        surface values((/0,0/), x, y, knotsx, knotsy, coeff(:,:,j))**2
    end do
    values=values-A**2
! Compute the R.M.S. error:
    sizev=norm(pack(values, (values == values)))/nvalues
    if (sizev <= TOLERANCE) then
        write(*,*) "Example 2 for SURFACE_FITTING is correct."
    end if
    end
```


## Output

Example 2 for SURFACE_FITTING is correct.

## Example 3: Constraining Some Points using a Spline Surface

This example illustrates the use of discrete constraints to shape the surface. The data fitting problem of Example 1 is modified by requiring that the surface interpolate the value one at $x=y=0$. The shape is constrained so first partial derivatives in both $x$ and $y$ are zero at $x=y=0$. These constraints mimic some properties of the function $g(x, y)$. The size of the residuals at a grid of points and the residuals of the constraints are checked.

```
USE surface_fitting_int
USE rand_int
USE norm_int
implicit none
```

! This is Example 3 for SURFACE_FITTING, tensor product
! B-splines approximation, $f(x, \bar{y})$. Use the function
! $\exp \left(-x^{* *} 2-y * * 2\right)$ on the $\operatorname{square}(0,2) x(0,2)$ for samples.
! The spline order is "nord" and the number of cells is
! "(ngrid-1)**2". There are "ndata" data values in the square.
! Constraints are put on the surface at (0,0). Namely
$!f(0,0)=1, f_{\_} x(0,0)=0, f_{\_} y(0,0)=0$.
integer : : i
integer, parameter : : ngrid=9, nord=4, ndegree=nord-1, \&
nbkpt=ngrid+2*ndegree, ndata $=2000$, nvalues $=100, \mathrm{NC}=3$
real(kind(1d0)), parameter : : zero=0d0, one=1d0, two=2d0
real(kind(1d0)), parameter : : TOLERANCE=1d-3
real(kind(1d0)), target : : spline_data (4, ndata), bkpt(nbkpt), \&

```
        coeff(ngrid+ndegree-1,ngrid+ndegree-1), delta, sizev, &
        x(nvalues), y(nvalues), values(nvalues, nvalues), &
        f_00, f_x00, f_y00
    real(kind(1d0)), pointer :: pointer_bkpt(:)
    type (d_spline_knots) knotsx, knotsy
    type (d_surface_constraints) C(NC)
    LOGICAL PASS
! Generate random (x,y) pairs and evaluate the
! example exponential function at these values.
    spline_data(1:2,:)=two*rand(spline_data(1:2,:))
    spline_data(3,:)=exp(-sum(spline_data(1:2,:)**2,dim=1))
    spline_data(4,:)=one
! Define the knots for the tensor product data fitting problem.
        delta = two/(ngrid-1)
        bkpt(1:ndegree) = zero
        bkpt(nbkpt-ndegree+1:nbkpt) = two
        bkpt(nord:nbkpt-ndegree) = (/ (i*delta,i=0,ngrid-1)/)
! Assign the degree of the polynomial and the knots.
    pointer_bkpt => bkpt
    knotsx=\overline{d_spline_knots(ndegree, pointer_bkpt)}
    knotsy=knotsx
! Define the constraints for the fitted surface.
    C(1)=surface_constraints(point=(/zero,zero/),type='==', value=one)
    C(2)=surface_constraints(derivative=(/1,0/),&
            point=(//zero,zero/),type='==',value=zero)
    C(3)=surface_constraints(derivative=(/0,1/),&
            point=(//zero,zero/),type='==',value=zero)
! Fit the data and obtain the coefficients.
    coeff = surface_fitting(spline_data, knotsx, knotsy,&
                        CONSTRAINTS=C)
! Evaluate the residual = spline - function
! at a grid of points inside the square.
    delta=two/(nvalues+1)
    x=(/(i*delta,i=1,nvalues)/); y=x
    values=exp (-spread (x**2,1,nvalues) -spread (y**2,2, nvalues))
    values=surface_values((/0,0/), x, y, knotsx, knotsy, coeff)-&
            values
    f_00 = surface_values((/0,0/), zero, zero, knotsx, knotsy, coeff)
    f_x00= surface_values((/1,0/), zero, zero, knotsx, knotsy, coeff)
    f_y00= surface values((/0,1/), zero, zero, knotsx, knotsy, coeff)
! Compute the R.M.S. error:
    sizev=norm(pack(values, (values == values)))/nvalues
    PASS = sizev <= TOLERANCE
    PASS = abs (f_00 - one) <= sqrt(epsilon(one)) .and. PASS
    PASS = f_x00 <= sqrt(epsilon(one)) .and. PASS
```

```
PASS = f_y00 <= sqrt(epsilon(one)) .and. PASS
if (PASS) then
    write(*,*) 'Example 3 for SURFACE_FITTING is correct.'
end if
end
```


## Output

Example 3 for SURFACE_FITTING is correct.

## Example 4: Constraining a Spline Surface to be non-Negative

The review of interpolating methods by Franke (1982) uses a test data set originally due to James Ferguson. We use this data set of 25 points, with unit uncertainty for each dependent variable. Our algorithm does not interpolate the data values but approximately fits them in the least-squares sense. We reset the regularization parameter values of flatness and thinness, Hanson (1995). Then the surface is fit to the data and evaluated at a grid of points. Although the surface appears smooth and fits the data, the values are negative near one corner. Our scenario for the application assumes that the surface be non-negative at all points of the rectangle containing the independent variable data pairs. Our algorithm for constraining the surface is simple but effective in this case. The data fitting is repeated one more time but with positive constraints at the grid of points where it was previously negative.

```
    USE surface_fitting_int
    USE rand_int
    USE surface_fitting_int
    USE rand_in'
    USE norm_int
    implicit none
! This is Example 4 for SURFACE_FITTING, tensor product
! B-splines approximation, f(x,y). Use the data set from
! Franke, due to Ferguson. Without constraints the function
! becomes negative in a corner. Constrain the surface
! at a grid of values so it is non-negative.
    integer :: i, j, q
    integer, parameter :: ngrid=9, nord=4, ndegree=nord-1, &
        nbkpt=ngrid+2*ndegree, ndata = 25, nvalues=50
    real(kind(1d0)), parameter :: zero=0d0, one=1d0
    real(kind(1d0)), parameter :: TOLERANCE=1d-3
    real(kind(1d0)), target :: spline_data (4, ndata), bkptx(nbkpt), &
        bkpty(nbkpt),coeff(ngrid+ndegree-1,ngrid+ndegree-1), &
            x(nvalues), y(nvalues), values(nvalues, nvalues), &
            delta
        real(kind(1d0)), pointer :: pointer_bkpt(:)
        type (d_spline_knots) knotsx, knotsy
        type (d_surface_constraints), allocatable :: C(:)
real(kind(1e0)) :: data (3*ndata) = & ! This is Fergu
    2.981 , 0.291, 3.4, 3.471, -7.062, 3.5,&
```

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline 3.961 & -14.418, & 3.5 & 7.45 & 12.003, & 2.5.8 \\
\hline 7.35 & 6.012, & 3.5 & 7.251, & 0.018, & 3.0, \% \\
\hline 7.151 & -5.973, & 2.0 & 7.051, & -11.967, & 2.5.8 \\
\hline 10.901, & 9.015, & 2.0 & 10.751, & 4.536, & 1.925, \& \\
\hline 10.602, & 0.06 & 1.85, & 10.453, & -4.419, & 1.576, \& \\
\hline 10.304, & -8.895, & 1.7 & 14.055, & 10.509, & 1.5, \& \\
\hline 14.194, & 6.783, & 1.3 & 14.331, & 3.054 , & 1.7.8 \\
\hline 14.469, & -0.672, & 2.1 & 14.607, & -4.398, & 1.75, \& \\
\hline 15.0, & 12.0 & 0.5 & 15.729, & 8.067, & 0.5.8 \\
\hline 16.457, & 4.134, & 0.7 & 17.185, & 0.198, & 1.1, \& \\
\hline 17.914, & -3.735, & 1.7/) & & & \\
\hline
\end{tabular}
    spline_data(1:3,:)=reshape(data,(/3,ndata/)); spline_data(4,:)=one
! Define the knots for the tensor product data fitting problem.
! Use the data limits to the knot sequences.
        bkptx(1:ndegree) = minval(spline_data(1,:))
        bkptx(nbkpt-ndegree+1:nbkpt) = maxval(spline_data(1,:))
        delta=(bkptx(nbkpt) -bkptx(ndegree)) / (ngrid-1)
        bkptx(nord:nbkpt-ndegree)=(/(bkptx(1)+i*delta,i=0,ngrid-1)/)
! Assign the degree of the polynomial and the knots for x.
    pointer_bkpt => bkptx
    knotsx=\overline{d_spline_knots(ndegree, pointer_bkpt)}
            bkpty(1:ndegree) = minval(spline_da\overline{ta(2,:))}
            bkpty(nbkpt-ndegree+1:nbkpt) = maxval(spline_data(2,:))
            delta=(bkpty(nbkpt)-bkpty(ndegree))/(ngrid-1)
            bkpty(nord:nbkpt-ndegree)}=(/(bkpty(1)+i*delta,i=0,ngrid-1)/)
! Assign the degree of the polynomial and the knots for y.
    pointer_bkpt => bkpty
    knotsy=\overline{d}spline_knots(ndegree, pointer_bkpt)
! Fit the data and obtain the coefficients.
    coeff = surface_fitting(spline_data, knotsx, knotsy)
    delta=(bkptx(nbkpt)-bkptx(1))/(nvalues+1)
    x=(/(bkptx(1)+i*delta,i=1,nvalues)/)
    delta=(bkpty(nbkpt) -bkpty(1))/(nvalues+1)
    y=(/(bkpty(1)+i*delta,i=1,nvalues)/)
! Evaluate the function at a rectangular grid.
! Use non-positive values to a constraint.
    values=surface_values((/0,0/), x, y, knotsx, knotsy, coeff)
! Count the number of values <= zero. Then constrain the spline
! so that it is >= TOLERANCE at those points where it was <= zero.
    q=count(values <= zero)
    allocate (C(q))
    DO I=1, nvalues
        DO J=1,nvalues
            IF(values(I,J) <= zero) THEN
                C(q)=surface_constraints(point=(/x(i),y(j)/), type='>=',&
                    value=TOLERANCE)
                q=q-1
```

```
                END IF
            END DO
    END DO
! Fit the data with constraints and obtain the coefficients.
    coeff = surface_fitting(spline_data, knotsx, knotsy,&
                CONSTRAINTS=C)
    deallocate(C)
! Evaluate the surface at a grid and check, once again, for
! non-positive values. All values should now be positive.
    values=surface_values((/0,0/), x, y, knotsx, knotsy, coeff)
if (count(values <= zero) == 0) then
    write(*,*) 'Example 4 for SURFACE_FITTING is correct.'
    end if
    end
```


## Output

Example 4 for SURFACE_FITTING is correct.

## CSIEZ

Computes the cubic spline interpolant with the 'not-a-knot' condition and return values of the interpolant at specified points.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
XVEC - Array of length n containing the points at which the spline is to be evaluated. (Input)

VALUE - Array of length N containing the values of the spline at the points in XVEC. (Output)

## Optional Arguments

NDATA - Number of data points. (Input)
NDATA must be at least 2.
Default: NDATA $=\operatorname{size}(X D A T A, 1)$.
$N$ - Length of vector XVEC. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{XVEC}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CSIEZ (XDATA, FDATA, XVEC, VALUE [, ...])

Specific: The specific interface names are S_CSIEZ and D_CSIEZ.

## FORTRAN 77 Interface

Single: CALL CSIEZ (NDATA, XDATA, FDATA, N, XVEC, VALUE)

Double: The double precision name is DCSIEZ.

## Description

This routine is designed to let the user easily compute the values of a cubic spline interpolant. The routine CSIEZ computes a spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA. The output for this routine consists of a vector of values of the computed cubic spline. Specifically, let $n=N, v=\mathrm{XVEC}$, and $y=\mathrm{VALUE}$, then if $s$ is the computed spline we set

$$
y_{j}=s\left(v_{j}\right) \quad j=1, \ldots, n
$$

Additional documentation can be found by referring to the IMSL routines CSINT or SPLEZ.

## Comments

Workspace may be explicitly provided, if desired, by use of C2IEZ/DC2IEZ. The reference is:

CALL C2IEZ (NDATA, XDATA, FDATA, N, XVEC, VALUE, IWK, WK1, WK2)

The additional arguments are as follows:

IWK - Integer work array of length MAX0(N, NDATA) + N.
WK1 - Real work array of length 5 * NDATA.
WK2 - Real work array of length 2 * N .

## Example

In this example, a cubic spline interpolant to a function $F$ is computed. The values of this spline are then compared with the exact function values.

```
USE CSIEZ_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=11)
INTEGER I, NOUT
```

!

```
    REAL F, FDATA(NDATA), FLOAT, SIN, VALUE (2*NDATA-1), X,&
            XDATA (NDATA), XVEC (2*NDATA-1)
    INTRINSIC FLOAT, SIN
!
    F(X) = SIN(15.0*X)
!
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I-1) /FLOAT(NDATA-1)
        FDATA(I) = F(XDATA(I))
    1 0 ~ C O N T I N U E
    DO 20 I=1, 2*NDATA - 1
        XVEC(I) = FLOAT(I-1)/FLOAT (2*NDATA-2)
    2 0 ~ C O N T I N U E ~
! Compute cubic spline interpolant
    CALL CSIEZ (XDATA, FDATA, XVEC, VALUE)
    Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
99998 FORMAT (13X, 'X', 9X, 'INTERPOLANT', 5X, 'ERROR')
    Print the interpolant and the error
    on a finer grid
    DO 30 I=1, 2*NDATA - 1
        WRITE (NOUT,99999) XVEC(I), VALUE(I), F(XVEC(I)) - VALUE(I)
    CONTINUE
99999 FORMAT(' ', 2F15.3, F15.6)
    END
```


## Output

| X | INTERPOLANT | ERROR |
| :---: | :---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.809 | -0.127025 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.723 | 0.055214 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.549 | -0.022789 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.843 | -0.016246 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.441 | 0.009348 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.903 | 0.019947 |
| 0.600 | 0.412 | 0.000000 |
| 0.650 | -0.315 | -0.004895 |
| 0.700 | -0.880 | 0.000000 |
| 0.750 | -0.938 | -0.029541 |
| 0.800 | -0.537 | 0.000000 |
| 0.850 | 0.148 | 0.034693 |
| 0.900 | 0.804 | 0.000000 |
| 0.950 | 1.086 | -0.092559 |
| 1.000 | 0.650 | 0.000000 |

## CSINT

Computes the cubic spline interpolant with the 'not-a-knot' condition.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces.
(Output)

## Optional Arguments

NDATA - Number of data points. (Input)
NDATA must be at least 2.
Default: NDATA $=\operatorname{size}(X D A T A, 1)$.

## FORTRAN 90 Interface

Generic: CALL CSINT (XDATA, FDATA, BREAK, CSCOEF [, ...])
Specific: The specific interface names are S_CSINT and D_CSINT.

## FORTRAN 77 Interface

Single: CALL CSINT (NDATA, XDATA, FDATA, BREAK, CSCOEF)

Double: The double precision name is DCSINT.

## Description

The routine CSINT computes a $C^{2}$ cubic spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA $=N$. The breakpoints of the spline are the abscissas. Endpoint conditions are automatically determined by the program. These conditions correspond to the "not-a-knot" condition (see de Boor 1978), which requires that the third derivative of the spline be continuous at the second and next-to-last breakpoint. If $N$ is 2 or 3 , then the linear or quadratic interpolating polynomial is computed, respectively.

If the data points arise from the values of a smooth (say $\left.C^{4}\right)$ function $f$, i.e. $f_{i}=f\left(x_{i}\right)$, then the error will behave in a predictable fashion. Let $\xi$ be the breakpoint vector for the above spline interpolant. Then, the maximum absolute error satisfies

$$
\|f-s\|_{\left[\xi_{1}, \xi_{N}\right]} \leq C\left\|f^{(4)}\right\|_{\left[\xi_{1}, \xi_{N}\right]}|\xi|^{4}
$$

where

$$
|\xi|:=\max _{i=2, \ldots, N}\left|\xi_{i}-\xi_{i-1}\right|
$$

For more details, see de Boor (1978, pages 55-56).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2INT/DC2INT. The reference is:

CALL C2INT (NDATA, XDATA, FDATA, BREAK, CSCOEF, IWK)
The additional argument is
IWK - Work array of length ndAtA.
2. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.
3. Note that column NDATA of CSCOEF is used as workspace.

## Example

In this example, a cubic spline interpolant to a function $F$ is computed. The values of this spline are then compared with the exact function values.

```
USE CSINT INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE CSVAL_INT
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=11)
INTEGER I, NINTV, NOUT
REAL BREAK(NDATA), CSCOEF (4,NDATA), F,&
    FDATA(NDATA), FLOAT, SIN, X, XDATA(NDATA)
    INTRINSIC FLOAT, SIN
    F(X) = SIN(15.0*X)
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I-1)/FLOAT(NDATA-1)
        FDATA(I) = F(XDATA(I))
    CONTINUE
CALL CSINT (XDATA, FDATA, BREAK, CSCOEF)
    Get output unit number.
```

```
    CALL UMACH (2, NOUT)
! Write heading
    WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
    NINTV = NDATA - 1
    Print the interpolant and the error
    on a finer grid
    DO 20 I=1, 2*NDATA - 1
        X = FLOAT (I-1)/FLOAT (2*NDATA-2)
        WRITE (NOUT,'(2F15.3,F15.6)') X, CSVAL (X,BREAK,CSCOEF) , &
                        F(X) - CSVAL (X,BREAK, &
                        CSCOEF)
    20 CONTINUE
    END
```


## Output

| X | Interpolant | Error |
| :---: | :---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.809 | -0.127025 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.723 | 0.055214 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.549 | -0.022789 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.843 | -0.016246 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.441 | 0.009348 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.903 | 0.019947 |
| 0.600 | -0.412 | 0.000000 |
| 0.650 | -0.880 | -0.004895 |
| 0.700 | -0.938 | 0.000000 |
| 0.750 | -0.537 | -0.029541 |
| 0.800 | 0.148 | 0.000000 |
| 0.850 | 0.804 | 0.034693 |
| 0.900 | 1.086 | 0.000000 |
| 0.950 | 0.650 | 0.092559 |
| 1.000 |  | 0.000000 |

## CSDEC

Computes the cubic spline interpolant with specified derivative endpoint conditions.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
ILEFT - Type of end condition at the left endpoint. (Input)

## ILEFT Condition

0 "Not-a-knot" condition

1 First derivative specified by DLEFT
2 Second derivative specified by DLEFT
DLEFT - Derivative at left endpoint if ILEFT is equal to 1 or 2 . (Input) If ILEFT $=0$, then DLEFT is ignored.

IRIGHT - Type of end condition at the right endpoint. (Input)

## IRIGHT Condition

0 "Not-a-knot" condition

1 First derivative specified by DRIGHT
2 Second derivative specified by DRIGHT

DRIGHT - Derivative at right endpoint if IRIGHT is equal to 1 or 2. (Input) If IRIGHT $=0$ then DRIGHT is ignored.

BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces. (Output)

## Optional Arguments

NDATA - Number of data points. (Input)
Default: NDATA $=$ size $($ XDATA, 1$)$.

## FORTRAN 90 Interface

Generic: CALL CSDEC (XDATA, FDATA, ILEFT, DLEFT, IRIGHT, DRIGHT, BREAK, CSCOEF [, ...])

Specific: The specific interface names are S_CSDEC and D_CSDEC.

## FORTRAN 77 Interface

Single: CALL CSDEC (NDATA, XDATA, FDATA, ILEFT, DLEFT, IRIGHT, DRIGHT, BREAK, CSCOEF)

Double: The double precision name is DCSDEC.

## Description

The routine CSDEC computes a $C^{2}$ cubic spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA $=N$. The breakpoints of the spline are the abscissas. Endpoint conditions are to be selected by the user. The user may specify not-a-knot, first derivative, or second derivative at each endpoint (see de Boor 1978, Chapter 4).

If the data (including the endpoint conditions) arise from the values of a smooth (say $C^{4}$ ) function $f$, i.e. $f_{i}=f\left(x_{i}\right)$, then the error will behave in a predictable fashion. Let $\xi$ be the breakpoint vector for the above spline interpolant. Then, the maximum absolute error satisfies

$$
\|f-s\|_{\left[\xi, 1, \xi_{N}\right]} \leq C\left\|f^{(4)}\right\|_{\left[\xi, \xi_{N}\right]}|\xi|^{4}
$$

where

$$
|\xi|:=_{i=2, \ldots, N}\left|\xi_{i}-\xi_{i-1}\right|
$$

For more details, see de Boor (1978, Chapter 4 and 5).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2DEC/DC2DEC. The reference is:
```
CALL C2DEC (NDATA, XDATA, FDATA, ILEFT, DLEFT, IRIGHT, DRIGHT, BREAK,
CSCOEF,IWK)
```

The additional argument is:
IWK - Work array of length ndATA.
2. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.
3. Note that column NDATA of CSCOEF is used as workspace.

## Example 1

In Example 1, a cubic spline interpolant to a function $f$ is computed. The value of the derivative at the left endpoint and the value of the second derivative at the right endpoint are specified. The values of this spline are then compared with the exact function values.

```
USE CSDEC INT
USE UMACH - INT
USE CSVAL_INT
IMPLICIT NONE
INTEGER ILEFT, IRIGHT, NDATA
```

```
    PARAMETER (ILEFT=1, IRIGHT=2, NDATA=11)
    INTEGER I, NINTV, NOUT
    REAL BREAK(NDATA), COS, CSCOEF(4,NDATA), DLEFT, &
                                DRIGHT, F, FDATA(NDATA), FLOAT, SIN, X, XDATA(NDATA)
    INTRINSIC COS, FLOAT, SIN
    F(X)}=\operatorname{SIN}(15.0*X
    DLEFT = 15.0*COS (15.0*0.0)
    DRIGHT = -15.0*15.0*SIN(15.0*1.0)
                                    Set up a grid
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT(NDATA-1)
        FDATA(I) = F(XDATA(I))
    1 0
    CONTINUE
! Compute cubic spline interpolant
    CALL CSDEC (XDATA, FDATA, ILEFT, DLEFT, IRIGHT, &
        DRIGHT, BREAK, CSCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
99999
    FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
    NINTV = NDATA - 1
! Print the interpolant on a finer grid
    DO 20 I=1, 2*NDATA - 1
        X = FLOAT (I-1)/FLOAT (2*NDATA-2)
        WRITE (NOUT,'(2F15.3,F15.6)') X, CSVAL (X,BREAK,CSCOEF) ,&
                        F(X) - CSVAL (X,BREAK, &
                        CSCOEF)
        2 0 ~ C O N T I N U E ~
        END
```


## Output

| X | Interpolant | Error |
| :---: | :---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.675 | 0.006332 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.759 | 0.019485 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.558 | -0.013227 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.840 | -0.018765 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.440 | 0.009859 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.902 | 0.020420 |
| 0.600 | -0.412 | 0.000000 |
| 0.650 | -0.880 | -0.007301 |
| 0.700 | -0.947 | 0.000000 |
| 0.750 | -0.537 | -0.020391 |
| 0.800 | 0.182 | 0.000000 |
| 0.850 | 0.000497 |  |


| 0.900 | 0.804 | 0.000000 |
| :--- | :--- | :--- |
| 0.950 | 0.959 | 0.035074 |
| 1.000 | 0.650 | 0.000000 |

## Additional Examples

## Example 2

In Example 2, we compute the natural cubic spline interpolant to a function $f$ by forcing the second derivative of the interpolant to be zero at both endpoints. As in the previous example, we compare the exact function values with the values of the spline.

USE CSDEC_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER ILEFT, IRIGHT, NDATA, NOUT
PARAMETER (ILEFT=2, IRIGHT=2, NDATA=11)
!
INTEGER I, NINTV
REAL BREAK (NDATA), CSCOEF (4,NDATA), DLEFT, DRIGHT,\&
F, FDATA (NDATA), FLOAT, SIN, X, XDATA(NDATA), CSVAL
INTRINSIC FLOAT, SIN
! Initialize DLEFT and DRIGHT
DATA DLEFT/0./, DRIGHT/0./
$F(X)=\operatorname{SIN}(15.0 * X)$
!
DO $10 \mathrm{I}=1$, NDATA XDATA $(I)=$ FLOAT $(I-1) / F L O A T($ NDATA-1) FDATA (I) = F(XDATA (I))
10 CONTINUE
!
CALL CSDEC (XDATA, FDATA, ILEFT, DLEFT, IRIGHT, DRIGHT, \& BREAK, CSCOEF)

CALL UMACH (2, NOUT)
!
Define function
Set up a grid

WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
NINTV = NDATA - 1
! Print the interpolant on a finer grid
DO 20 I=1, 2*NDATA - 1
$\mathrm{X}=\mathrm{FLOAT}(\mathrm{I}-1) / \mathrm{FLOAT}(2 *$ NDATA-2)
WRITE (NOUT,'(2F15.3,F15.6)') X, CSVAL (X,BREAK,CSCOEF), \&
F(X) - CSVAL (X,BREAK, \&
CSCOEF)
20 CONTINUE
END

## Output

| X | Interpolant | Error |
| ---: | ---: | ---: |
| 0.000 | 0.000 | 0.000000 |


| 0.050 | 0.667 | 0.015027 |
| :--- | ---: | ---: |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.761 | 0.017156 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.559 | -0.012609 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.840 | -0.018907 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.440 | 0.009812 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.902 | 0.020753 |
| 0.600 | 0.412 | 0.000000 |
| 0.650 | -0.311 | -0.008586 |
| 0.700 | -0.880 | 0.000000 |
| 0.750 | -0.952 | -0.015585 |
| 0.800 | -0.537 | 0.000000 |

## CSHER

Computes the Hermite cubic spline interpolant.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input)
The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
DFDATA - Array of length NDATA containing the values of the derivative. (Input)
BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces.
(Output)

## Optional Arguments

NDATA - Number of data points. (Input)
Default: NDATA $=$ size $($ XDATA, 1$)$.

## FORTRAN 90 Interface

Generic: CALL CSHER (XDATA, FDATA, DFDATA, BREAK, CSCOEF [,...])

Specific: The specific interface names are S_CSHER and D_CSHER.

## FORTRAN 77 Interface

Single: CALL CSHER (NDATA, XDATA, FDATA, BREAK, CSCOEF)
Double: The double precision name is DCSHER.

## Description

The routine CSHER computes a $C^{1}$ cubic spline interpolant to the set of data points

$$
\left(x_{i}, f_{i}\right) \operatorname{and}\left(x_{i}, f_{i}^{\prime}\right)
$$

for $i=1, \ldots$, NDATA $=\mathrm{N}$. The breakpoints of the spline are the abscissas.
If the data points arise from the values of a smooth (say $C^{4}$ ) function $f$, i.e.,

$$
f_{i}=f\left(x_{i}\right) \text { and } f_{i}^{\prime}=f^{\prime}\left(x_{i}\right)
$$

then the error will behave in a predictable fashion. Let $\xi$ be the
breakpoint vector for the above spline interpolant. Then, the maximum absolute error satisfies

$$
\|f-s\|_{\left[\xi 5, \xi_{N}\right]} \leq C\left\|f^{(4)}\right\|_{\left[\xi, \xi_{N}\right]}|\xi|^{4}
$$

where

$$
|\xi|:={ }_{i=2, \ldots, N}\left|\xi_{i}-\xi_{i-1}\right|
$$

For more details, see de Boor (1978, page 51).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2HER/DC2HER. The reference is:

CALL C2HER (NDATA, XDATA, FDATA, DFDATA, BREAK, CSCOEF, IWK)
The additional argument is:
IWK - Work array of length ndATA.
2. Informational error

Type Code
42 The XDATA values must be distinct.
3. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.
4. Note that column NDATA of CSCOEF is used as workspace.

## Example

In this example, a cubic spline interpolant to a function $f$ is computed. The value of the function $f$ and its derivative $f^{\prime}$ are computed on the interpolation nodes and passed to CSHER. The values of this spline are then compared with the exact function values.

```
    USE CSHER INT
    USE UMACH - INT
    USE CSVAL_INT
    IMPLICIT NONE
    INTEGER NDATA
    PARAMETER (NDATA=11)
!
    INTEGER I, NINTV, NOUT
    REAL BREAK (NDATA), COS, CSCOEF(4,NDATA), DF,&
        DFDATA(NDATA), F, FDATA(NDATA), FLOAT, SIN, X, &
        XDATA (NDATA)
    INTRINSIC COS, FLOAT, SIN
    F(X) = SIN(15.0*X)
    DF(X)}=15.0*\operatorname{COS}(15.0*X
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT (NDATA-1)
        FDATA(I) = F(XDATA(I))
        DFDATA(I) = DF (XDATA(I))
    10 CONTINUE
    CALL CSHER (XDATA, FDATA, DFDATA, BREAK, CSCOEF)
                    Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
    NINTV = NDATA - 1
    Print the interpolant on a finer grid
    DO 20 I=1, 2*NDATA - 1
        X = FLOAT (I-1)/FLOAT (2*NDATA-2)
        WRITE (NOUT,'(2F15.3, F15.6)') X, CSVAL (X,BREAK,CSCOEF) &
                        , F(X) - CSVAL(X,BREAK,&
                        CSCOEF)
    2 0 ~ C O N T I N U E ~
        END
```


## Output

| X | Interpolant | Error |
| :---: | :---: | :---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.673 | 0.008654 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.768 | 0.009879 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.564 | -0.007257 |


| 0.300 | -0.978 | 0.000000 |
| ---: | ---: | ---: |
| 0.350 | -0.848 | -0.010906 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.444 | 0.005714 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.911 | 0.011714 |
| 0.600 | 0.412 | 0.000000 |
| 0.650 | -0.315 | -0.004057 |
| 0.700 | -0.880 | 0.000000 |
| 0.750 | -0.956 | -0.012288 |
| 0.800 | -0.537 | 0.000000 |
| 0.850 | 0.180 | 0.002318 |
| 0.900 | 0.804 | 0.00000 |
| 0.950 | 0.981 | 0.012616 |
| 1.000 | 0.650 | 0.000000 |

## CSAKM

Computes the Akima cubic spline interpolant.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)

BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces. (Output)

## Optional Arguments

NDATA - Number of data points. (Input) Default: NDATA $=\operatorname{size}($ XDATA 1$)$.

## FORTRAN 90 Interface

Generic: CALL CSAKM (XDATA, FDATA, BREAK, CSCOEF [, ...])
Specific: The specific interface names are S_CSAKM and D_CSAKM.

## FORTRAN 77 Interface

Single: CALL CSAKM (NDATA, XDATA, FDATA, BREAK, CSCOEF)
Double: The double precision name is DCSAKM.

## Description

The routine CSAKM computes a $C^{1}$ cubic spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA $=N$. The breakpoints of the spline are the abscissas. Endpoint conditions are automatically determined by the program; see Akima (1970) or de Boor (1978).
If the data points arise from the values of a smooth (say $C^{4}$ ) function $f$, i.e. $f_{i}=f\left(x_{i}\right)$, then the error will behave in a predictable fashion. Let $\xi$ be the breakpoint vector for the above spline interpolant. Then, the maximum absolute error satisfies

$$
\|f-s\|_{\left[\xi, 1, \xi_{N}\right]} \leq C\left\|f^{(2)}\right\|_{\left[\xi_{1}, \xi_{N}\right]}|\xi|^{2}
$$

where

$$
|\xi|:=\max _{i=2, \ldots, N}\left|\xi_{i}-\xi_{i-1}\right|
$$

The routine CSAKM is based on a method by Akima (1970) to combat wiggles in the interpolant. The method is nonlinear; and although the interpolant is a piecewise cubic, cubic polynomials are not reproduced. (However, linear polynomials are reproduced.)

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2AKMD/C2AKM. The reference is:

CALL C2AKM (NDATA, XDATA, FDATA, BREAK, CSCOEF, IWK)
The additional argument is:
IWK - Work array of length NDATA.
2. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.
3. Note that column NDATA of CSCOEF is used as workspace.

## Example

In this example, a cubic spline interpolant to a function $f$ is computed. The values of this spline are then compared with the exact function values.

```
USE CSAKM INT
USE UMACH_INT
USE CSVAL_INT
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=11)
INTEGER I, NINTV, NOUT
REAL BREAK (NDATA), CSCOEF (4,NDATA), F,&
```

```
        FDATA(NDATA), FLOAT, SIN, X, XDATA(NDATA)
    INTRINSIC FLOAT, SIN
! Define function
    F(X) = SIN(15.0*X)
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I-1)/FLOAT(NDATA-1)
        FDATA(I) = F(XDATA(I))
    1 0 \text { CONTINUE}
! Compute cubic spline interpolant
    CALL CSAKM (XDATA, FDATA, BREAK, CSCOEF)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
    NINTV = NDATA - 1
! Print the interpolant on a finer grid
    DO 20 I=1, 2*NDATA - 1
        X = FLOAT (I-1)/FLOAT (2*NDATA-2)
        WRITE (NOUT,'(2F15.3,F15.6)') X, CSVAL (X,BREAK,CSCOEF) , &
                                    F(X) - CSVAL (X,BREAK, &
                                    CSCOEF)
    2 0 ~ C O N T I N U E
    END
```


## Output

| X | Interpolant | Error |
| :---: | ---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.818 | -0.135988 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.615 | 0.163487 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.478 | -0.093376 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.812 | -0.046447 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.386 | 0.064491 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.854 | 0.068274 |
| 0.600 | -0.412 | 0.000000 |
| 0.650 | -0.880 | -0.043288 |
| 0.700 | -0.889 | 0.000000 |
| 0.750 | -0.537 | -0.078947 |
| 0.800 | 0.149 | 0.000000 |
| 0.850 | 0.804 | 0.033757 |
| 0.900 | 0.932 | 0.000000 |
| 0.950 | 0.650 | 0.061260 |
| 1.000 |  | 0.000000 |

## CSCON

Computes a cubic spline interpolant that is consistent with the concavity of the data.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)

IBREAK - The number of breakpoints. (Output) It will be less than 2 * NDATA.

BREAK - Array of length IBREAK containing the breakpoints for the piecewise cubic representation in its first IBREAK positions. (Output) The dimension of BREAK must be at least 2 * NDATA.

CSCOEF - Matrix of size 4 by N where N is the dimension of BREAK. (Output) The first IBREAK - 1 columns of CSCOEF contain the local coefficients of the cubic pieces.

## Optional Arguments

NDATA - Number of data points. (Input) NDATA must be at least 3 . Default: NDATA = size (XDATA,1).

## FORTRAN 90 Interface

Generic: CALL CSCON (XDATA, FDATA, IBREAK, BREAK, CSCOEF [,...])

Specific: The specific interface names are S_CSCON and D_CSCON.

## FORTRAN 77 Interface

Single: CALL CSCON (NDATA, XDATA, FDATA, IBREAK, BREAK, CSCOEF)

Double: The double precision name is DCSCON.

## Descritpion

The routine CSCON computes a cubic spline interpolant to $n=$ NDATA data points $\left\{x_{i}, f_{i}\right\}$ for $i=1, \ldots, n$. For ease of explanation, we will assume that $x_{i}<x_{i+1}$, although it is not necessary for the user to sort these data values. If the data are strictly convex, then the computed spline is convex, $C^{2}$, and minimizes the expression

$$
\int_{x_{1}}^{x_{n}}\left(g^{\prime \prime}\right)^{2}
$$

over all convex $C^{1}$ functions that interpolate the data. In the general case when the data have both convex and concave regions, the convexity of the spline is consistent with the data and the above
integral is minimized under the appropriate constraints. For more information on this interpolation scheme, we refer the reader to Micchelli et al. (1985) and Irvine et al. (1986).

One important feature of the splines produced by this subroutine is that it is not possible, a priori, to predict the number of breakpoints of the resulting interpolant. In most cases, there will be breakpoints at places other than data locations. The method is nonlinear; and although the interpolant is a piecewise cubic, cubic polynomials are not reproduced. (However, linear polynomials are reproduced.) This routine should be used when it is important to preserve the convex and concave regions implied by the data.

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2CON/DC2CON. The reference is:

CALL C2CON (NDATA, XDATA, FDATA, IBREAK, BREAK, CSCOEF, ITMAX, XSRT, FSRT, A, Y, DIVD, ID, WK)
The additional arguments are as follows:
ITMAX - Maximum number of iterations of Newton's method. (Input)
$\boldsymbol{X S R T}$ - Work array of length NDATA to hold the sorted XDATA values.
$\boldsymbol{F S R T}$ - Work array of length NDATA to hold the sorted FDATA values.
A - Work array of length NDATA.
$\boldsymbol{Y}-$ Work array of length NDATA -2.
DIVD - Work array of length NDATA - 2.
ID - Integer work array of length NDATA.
WK - Work array of length 5 * (ndATA -2 ).
2 Informational errors
Type Code
316 Maximum number of iterations exceeded, call C2CON/DC2CON to set a larger number for ITMAX.
43 The XDATA values must be distinct.
3. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.
4. The default value for ITMAX is 25 . This can be reset by calling C2CON/DC2CON directly.

## Example

We first compute the shape-preserving interpolant using CSCON, and display the coefficients and breakpoints. Second, we interpolate the same data using CSINT in a program not shown and overlay the two results. The graph of the result from CSINT is represented by the dashed line. Notice the extra inflection points in the curve produced by CSINT.

```
USE CSCON INT
USE UMACH_}\mp@subsup{}{}{-}\mathrm{ INT
USE WRRRL_INT
IMPLICIT NONE
PARAMETER (NDATA=9)
INTEGER IBREAK, NOUT
REAL BREAK (2*NDATA), CSCOEF (4,2*NDATA), FDATA (NDATA),&
    XDATA (NDATA)
CHARACTER CLABEL(14)*2, RLABEL (4)*2
DATA XDATA/0.0, .1, .2, . 3, .4, .5, .6, . 8, 1./
DATA FDATA/0.0, .9, .95, .9, .1, .05, .05, .2, 1./
DATA RLABEL/' 1', ' 2', ' 3', ' 4'/
DATA CLABEL/'' ', ' 1',' '2', ' 3',' ' 4', ' 5', ' 6',&
    ' 7', ' 8', ' 9', '10', '11', '12', '13'/
    Compute cubic spline interpolant
CALL CSCON (XDATA, FDATA, IBREAK, BREAK, CSCOEF)
CALL UMACH (2, NOUT)
    Print the BREAK points and the
    coefficients (CSCOEF) for
    checking purposes.
WRITE (NOUT,'(1X,A,I2)') 'IBREAK = ', IBREAK
CALL WRRRL ('BREAK', BREAK, RLABEL, CLABEL, 1, IBREAK, 1, &
                FMT='(F9.3)')
CALL WRRRL ('CSCOEF', CSCOEF, RLABEL, CLABEL, 4, IBREAK, 4, &
        FMT='(F9.3)')
    END
```


## Output

| IBREAK $=13$ |  | BREAK |  |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 | 5 | 6 |
| 1 | 0.000 | 0.100 | 0.136 | 0.200 | 0.259 | 0.300 |
|  | 7 | 8 | 9 | 10 | 11 | 12 |
| 1 | 0.400 | 0.436 | 0.500 | 0.600 | 0.609 | 0.800 |


|  | 1 |
| :--- | ---: |
| 1 | 0.000 |
| 2 | 11.886 |
| 3 | 0.000 |
| 4 | -1731.699 |
|  | 7 |
| 1 | 0.100 |
| 2 | -4.121 |
| 3 | 226.470 |
| 4 | -6222.348 |
|  |  |
| 1 | 1.000 |
| 2 | 0.000 |
| 3 | 0.000 |
| 4 | 0.000 |


| 2 | 3 | 4 | 5 | 6 |
| ---: | ---: | ---: | ---: | ---: |
| 0.900 | 0.942 | 0.950 | 0.958 | 0.900 |
| 3.228 | 0.131 | 0.131 | 0.131 | -4.434 |
| -173.170 | 0.000 | 0.000 | 0.000 | 220.218 |
| 4841.604 | 0.000 | 0.000 | -5312.082 | 4466.875 |
|  |  |  |  | 11 | 3

0.000
0.000


Figure 3-4 CSCON vs. CSINT

## CSPER

Computes the cubic spline interpolant with periodic boundary conditions.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)

BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces. (Output)

## Optional Arguments

NDATA - Number of data points. (Input)
NDATA must be at least 4 .
Default: NDATA $=\operatorname{size}(X D A T A, 1)$.

## FORTRAN 90 Interface

Generic: CALL CSPER (XDATA, FDATA, BREAK, CSCOEF [, ...])

Specific: The specific interface names are S_CSPER and D_CSPER.

## FORTRAN 77 Interface

Single: CALL CSPER (NDATA, XDATA, FDATA, BREAK, CSCOEF)
Double: The double precision name is DCSPER.

## Description

The routine CSPER computes a $C^{2}$ cubic spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA $=N$. The breakpoints of the spline are the abscissas. The program enforces periodic endpoint conditions. This means that the spline $s$ satisfies $s(a)=s(b), s^{\prime}(a)=s^{\prime}(b)$, and $s^{\prime \prime}(a)=s^{\prime \prime}(b)$, where $a$ is the leftmost abscissa and $b$ is the rightmost abscissa. If the ordinate values corresponding to $a$ and $b$ are not equal, then a warning message is issued. The ordinate value at $b$ is set equal to the ordinate value at $a$ and the interpolant is computed.

If the data points arise from the values of a smooth ( say $C^{4}$ ) periodic function $f$, i.e. $f_{i}=f\left(x_{i}\right)$, then the error will behave in a predictable fashion. Let $\xi$ be the breakpoint vector for the above spline interpolant. Then, the maximum absolute error satisfies

$$
\|f-s\|_{\left[\xi_{1}, \xi_{N}\right]} \leq C\left\|f^{(4)}\right\|_{\left[\xi, 1, \xi_{N}\right]}|\xi|^{4}
$$

where

$$
|\xi|:=\max _{i=2, \ldots, N}\left|\xi_{i}-\xi_{i-1}\right|
$$

For more details, see de Boor (1978, pages 320-322).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2PER/DC2PER. The reference is:

CALL C2PER (NDATA, XDATA, FDATA, BREAK, CSCOEF, WK, IWK)
The additional arguments are as follows:

WK - Work array of length 6 * NDATA.
IWK - Work array of length ndAtA.
2. Informational error

Type Code
31 The data set is not periodic, i.e., the function values at the smallest and largest XDATA points are not equal. The value at the smallest XDATA point is used.
3. The cubic spline can be evaluated using CSVAL and its derivative can be evaluated using CSDER.

## Example

In this example, a cubic spline interpolant to a function $f$ is computed. The values of this spline are then compared with the exact function values.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=11)
$!$
INTEGER I, NINTV, NOUT
REAL BREAK (NDATA), $\operatorname{CSCOEF}(4, N D A T A), F, \&$
FDATA (NDATA), FLOAT, H, PI, SIN, X, XDATA(NDATA)
INTRINSIC FLOAT, SIN
$!$
! Define function
$F(X)=\operatorname{SIN}(15.0 * X)$
$\mathrm{PI}=\mathrm{CONST}\left(\mathrm{P}^{\prime} \mathrm{PI}\right)$
$\mathrm{H}=2.0 * P I / 15.0 / 10.0$
DO 10 I=1, NDATA
XDATA(I) $=H^{*} F L O A T(I-1)$
FDATA(I) $=F(X D A T A(I))$
10 CONTINUE
Round off will cause FDATA(11) to
be nonzero; this would produce a
warning error since FDATA(1) is zero.
Therefore, the value of FDATA(1) is
used rather than the value of

```
! FDATA(11).
    FDATA(NDATA) = FDATA(1)
!
! Compute cubic spline interpolant
    CALL CSPER (XDATA, FDATA, BREAK, CSCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
    NINTV = NDATA - 1
    H = H/2.0
! Print the interpolant on a finer grid
    DO 20 I=1, 2*NDATA - 1
        X = H*FLOAT (I-1)
        WRITE (NOUT,'(2F15.3,F15.6)') X, CSVAL (X,BREAK,CSCOEF),&
                        F(X) - CSVAL (X,BREAK, &
                                    CSCOEF)
    20 CONTINUE
        END
```


## Output

| X | Interpolant | Error |
| :--- | ---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.021 | 0.309 | 0.000138 |
| 0.042 | 0.588 | 0.000000 |
| 0.063 | 0.809 | 0.000362 |
| 0.084 | 0.951 | 0.000000 |
| 0.105 | 1.000 | 0.000447 |
| 0.126 | 0.951 | 0.000000 |
| 0.147 | 0.809 | 0.000362 |
| 0.168 | 0.588 | 0.000000 |
| 0.188 | 0.309 | 0.000138 |
| 0.209 | 0.000 | 0.000000 |
| 0.230 | -0.309 | -0.000138 |
| 0.251 | -0.588 | 0.000000 |
| 0.272 | -0.809 | -0.000362 |
| 0.293 | -0.951 | 0.000000 |
| 0.314 | -1.000 | -0.000447 |
| 0.335 | -0.951 | 0.000000 |
| 0.356 | -0.809 | -0.000362 |
| 0.377 | -0.588 | 0.000000 |
| 0.398 | -0.309 | -0.000138 |
| 0.419 | 0.000 | 0.000000 |

## CSVAL

This function evaluates a cubic spline.

## Function Return Value

CSVAL - Value of the polynomial at x . (Output)

## Required Arguments

$\boldsymbol{X}$ - Point at which the spline is to be evaluated. (Input)

BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise cubic representation. (Input)
BREAK must be strictly increasing.
CSCOEF - Matrix of size 4 by NINTV +1 containing the local coefficients of the cubic pieces. (Input)

## Optional Arguments

NINTV — Number of polynomial pieces. (Input)

## FORTRAN 90 Interface

Generic: CSVAL (X, BREAK, CSCOEF[,$\ldots$.$] )$

Specific: The specific interface names are S_CSVAL and D_CSVAL.

## FORTRAN 77 Interface

Single: CSVAL (X, NINTV, BREAK, CSCOEF)
Double: The double precision function name is DCSVAL.

## Description

The routine CSVAL evaluates a cubic spline at a given point. It is a special case of the routine PPDER, which evaluates the derivative of a piecewise polynomial. (The value of a piecewise polynomial is its zero-th derivative and a cubic spline is a piecewise polynomial of order 4.) The routine PPDER is based on the routine PPVALU in de Boor (1978, page 89).

## Example

For an example of the use of CSVAL, see IMSL routine CSINT.

## CSDER

This function evaluates the derivative of a cubic spline.

## Function Return Value

CSDER - Value of the IDERIV-th derivative of the polynomial at X. (Output)

## Required Arguments

IDERIV - Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the value of the polynomial.
$\boldsymbol{X}$ - Point at which the polynomial is to be evaluated. (Input)

BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise cubic representation. (Input) BREAK must be strictly increasing.

CSCOEF - Matrix of size 4 by nintv +1 containing the local coefficients of the cubic pieces. (Input)

## Optional Arguments

NINTV - Number of polynomial pieces. (Input) Default: NINTV $=\operatorname{size}($ BREAK, $)-1$.

## FORTRAN 90 Interface

Generic: CSDER (IDERIV, X, BREAK, CSCOEF, CSDER [, ...])
Specific: The specific interface names are S_CSDER and D_CSDER.

## FORTRAN 77 Interface

Single: CSDER (IDERIV, X, NINTV, BREAK, CSCOEF)

Double: The double precision function name is DCSDER.

## Description

The function CSDER evaluates the derivative of a cubic spline at a given point. It is a special case of the routine PPDER, which evaluates the derivative of a piecewise polynomial. (A cubic spline is a piecewise polynomial of order 4.) The routine PPDER is based on the routine PPVALU in de Boor (1978, page 89).

## Example

In this example, we compute a cubic spline interpolant to a function $f$ using IMSL routine CSINT. The values of the spline and its first and second derivatives are computed using CSDER. These values can then be compared with the corresponding values of the interpolated function.

```
USE CSDER INT
USE CSINT INT
USE UMACH_INT
IMPLICIT NONE
```

```
    INTEGER NDATA
    PARAMETER (NDATA=10)
!
    INTEGER I, NINTV, NOUT
    REAL BREAK(NDATA), CDDF, CDF, CF, COS, CSCOEF(4,NDATA),&
        DDF, DF, F, FDATA(NDATA), FLOAT, SIN, X,&
        XDATA(NDATA)
    INTRINSIC COS, FLOAT, SIN
!
    F(X) = SIN(15.0*X)
    DF(X) = 15.0*COS (15.0*X)
    DDF(X) = -225.0*SIN(15.0*X)
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT (NDATA-1)
        FDATA(I) = F(XDATA(I))
    10 CONTINUE
! Compute cubic spline interpolant
    CALL CSINT (XDATA, FDATA, BREAK, CSCOEF)
                Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
99999 FORMAT (9X, 'X', 8X, 'S(X)', 5X, 'Error', 6X, 'S''(X)', 5X,&
            'Error', 6X, 'S''''(X)', 4X, 'Error', /)
    NINTV = NDATA - 1
    DO 20 I=1, 2*NDATA
        X = FLOAT (I-1)/FLOAT (2*NDATA-1)
        CF = CSDER (0, X,BREAK,CSCOEF)
        CDF = CSDER (1, X,BREAK,CSCOEF)
        CDDF = CSDER (2,X,BREAK,CSCOEF)
        WRITE (NOUT,'(F11.3, 3(F11.3, F11.6))') X, CF, F(X) - CF,&
                                    CDF, DF(X) - CDF,&
                                    CDDF, DDF(X) - CDDF
    2 0 ~ C O N T I N U E
        END
```


## Output

| X | $S(X)$ | Error | $S^{\prime}(X)$ | Error | $S^{\prime \prime}(X)$ | Error |
| :---: | ---: | ---: | ---: | ---: | ---: | ---: |
|  |  |  |  |  |  |  |
| 0.000 | 0.000 | 0.000000 | 26.285 | -11.284739 | -379.458 | 379.457794 |
| 0.053 | 0.902 | -0.192203 | 8.841 | 1.722460 | -283.411 | 123.664734 |
| 0.105 | 1.019 | -0.019333 | -3.548 | 3.425718 | -187.364 | -37.628586 |
| 0.158 | 0.617 | 0.081009 | -10.882 | 0.146207 | -91.317 | -65.824875 |
| 0.211 | -0.037 | 0.021155 | -13.160 | -1.837700 | 4.730 | -1.062027 |
| 0.263 | -0.674 | -0.046945 | -10.033 | -0.355268 | 117.916 | 44.391640 |
| 0.316 | -0.985 | -0.015060 | -0.719 | 1.086203 | 235.999 | -11.066727 |
| 0.368 | -0.682 | -0.004651 | 11.314 | -0.409097 | 154.861 | -0.365387 |
| 0.421 | 0.045 | -0.011915 | 14.708 | 0.284042 | -25.887 | 18.552732 |
| 0.474 | 0.708 | 0.024292 | 9.508 | 0.702690 | -143.785 | -21.041260 |
| 0.526 | 0.978 | 0.020854 | 0.161 | -0.771948 | -211.402 | -13.411087 |
| 0.579 | 0.673 | 0.001410 | -11.394 | 0.322443 | -163.483 | 11.674103 |
| 0.632 | -0.064 | 0.015118 | -14.937 | -0.045511 | 28.856 | -17.856323 |


| 0.684 | -0.724 | -0.019246 | -8.859 | -1.170871 | 163.866 | 3.435547 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.737 | -0.954 | -0.044143 | 0.301 | 0.554493 | 184.217 | 40.417282 |
| 0.789 | -0.675 | 0.012143 | 10.307 | 0.928152 | 166.021 | -16.939514 |
| 0.842 | 0.027 | 0.038176 | 15.015 | -0.047344 | 12.914 | -27.575521 |
| 0.895 | 0.764 | -0.010112 | 11.666 | -1.819128 | -140.193 | -29.538193 |
| 0.947 | 1.114 | -0.116304 | 0.258 | -1.357680 | -293.301 | 68.905701 |
| 1.000 | 0.650 | 0.000000 | -19.208 | 7.812407 | -446.408 | 300.092896 |

## CS1GD

Evaluates the derivative of a cubic spline on a grid.

## Required Arguments

IDERIV - Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the values of the cubic spline.
XVEC - Array of length N containing the points at which the cubic spline is to be evaluated. (Input)
The points in XVEC should be strictly increasing.
BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise cubic representation. (Input)
BREAK must be strictly increasing.
CSCOEF - Matrix of size 4 by NINTV +1 containing the local coefficients of the cubic pieces. (Input)

VALUE - Array of length $N$ containing the values of the IDERIV-th derivative of the cubic spline at the points in XVEC. (Output)

## Optional Arguments

$N$ - Length of vector XVEC. (Input) Default: $\mathrm{N}=$ size (XVEC,1).

NINTV - Number of polynomial pieces. (Input)
Default: NINTV $=\operatorname{size}($ BREAK, 1$)-1$.

## FORTRAN 90 Interface

Generic: CALL CS1GD (IDERIV, XVEC, BREAK, CSCOEF, VALUE [, ...])

Specific: The specific interface names are S_CS1GD and D_CS1GD.

## FORTRAN 77 Interface

Single: CALL CS1GD (IDERIV, N, XVEC, NINTV, BREAK, CSCOEF, VALUE)

Double: $\quad$ The double precision name is DCS1GD.

## Description

The routine CS1GD evaluates a cubic spline (or its derivative) at a vector of points. That is, given a vector $x$ of length $n$ satisfying $x_{i}<x_{i+1}$ for $i=1, \ldots, n-1$, a derivative value $j$, and a cubic spline $s$ that is represented by a breakpoint sequence and coefficient matrix this routine returns the values

$$
s^{(j)}\left(x_{i}\right) \quad i=1, \ldots, n
$$

in the array VALUE. The functionality of this routine is the same as that of CSDER called in a loop, however CS1GD should be much more efficient.

## Comments

1. Workspace may be explicitly provided, if desired, by use of C21GD/DC21GD. The reference is:
```
CALL C21GD (IDERIV, N, XVEC, NINTV, BREAK, CSCOEF, VALUE, IWK, WORK1,
WORK2)
```

The additional arguments are as follows:
IWK - Array of length N.
WORK1 - Array of length N.
WORK2 - Array of length N.
2. Informational error

Type Code
44 The points in XVEC must be strictly increasing.

## Example

To illustrate the use of CS1GD, we modify the example program for CSINT. In this example, a cubic spline interpolant to $F$ is computed. The values of this spline are then compared with the exact function values. The routine CS1GD is based on the routine PPVALU in de Boor (1978, page 89).

```
USE CSIGD INT
USE CSINT-INT
USE UMACH_INT
USE CSVAL_INT
IMPLICIT NONE
! Specifications
INTEGER NDATA, N, IDERIV, J
PARAMETER (NDATA=11, N=2*NDATA-1)
INTEGER I, NINTV, NOUT
```

```
    REAL BREAK(NDATA), CSCOEF(4,NDATA), F,&
        FDATA(NDATA), FLOAT, SIN, X, XDATA(NDATA) , &
        FVALUE (N), VALUE (N), XVEC (N)
    INTRINSIC FLOAT, SIN
!
    F(X)}=\operatorname{SIN}(15.0*X
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT (NDATA-1)
        FDATA(I) = F(XDATA(I))
    1 0 \text { CONTINUE}
! Compute cubic spline interpolant
    CALL CSINT (XDATA, FDATA, BREAK, CSCOEF)
    DO 20 I=1, N
        XVEC(I) = FLOAT(I-1)/FLOAT (2*NDATA-2)
        FVALUE(I) = F(XVEC(I))
    20 CONTINUE
        IDERIV = 0
        NINTV = NDATA - 1
        CALL CSIGD (IDERIV, XVEC, BREAK, CSCOEF, VALUE)
! Get output unit number.
    CALL UMACH (2, NOUT)
! Write heading
    WRITE (NOUT,99999)
99999 FORMAT (13X, 'X', 9X, 'Interpolant', 5X, 'Error')
! Print the interpolant and the error
! on a finer grid
    DO 30 J=1, N
        WRITE (NOUT,'(2F15.3,F15.6)') XVEC(J), VALUE(J), &
                            FVALUE (J) -VALUE (J)
    3 0 ~ C O N T I N U E
    END
```

    Output
    | X | Interpolant | Error |
| :---: | :---: | ---: |
| 0.000 | 0.000 | 0.000000 |
| 0.050 | 0.809 | -0.127025 |
| 0.100 | 0.997 | 0.000000 |
| 0.150 | 0.723 | 0.055214 |
| 0.200 | 0.141 | 0.000000 |
| 0.250 | -0.549 | -0.022789 |
| 0.300 | -0.978 | 0.000000 |
| 0.350 | -0.843 | -0.016246 |
| 0.400 | -0.279 | 0.000000 |
| 0.450 | 0.441 | 0.009348 |
| 0.500 | 0.938 | 0.000000 |
| 0.550 | 0.903 | 0.019947 |
| 0.600 | 0.412 | 0.000000 |
| 0.650 | -0.315 | -0.004895 |
| 0.700 | -0.880 | 0.000000 |
| 0.750 | -0.537 | -0.029541 |
| 0.800 | 0.148 | 0.000000 |
| 0.850 | 0.804 | 0.034693 |
| 0.900 | 0.000000 |  |


| 0.950 | 1.086 | -0.092559 |
| ---: | ---: | ---: |
| 1.000 | 0.650 | 0.000000 |

## CSITG

This function evaluates the integral of a cubic spline.

## Function Return Value

CSITG - Value of the integral of the spline from A to B. (Output)

## Required Arguments

$\boldsymbol{A}-$ Lower limit of integration. (Input)

B - Upper limit of integration. (Input)

BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise cubic representation. (Input) BREAK must be strictly increasing.

CSCOEF - Matrix of size 4 by NINTV +1 containing the local coefficients of the cubic pieces. (Input)

## Optional Arguments

NINTV - Number of polynomial pieces. (Input)
Default: NINTV $=\operatorname{size}($ BREAK, 1$)-1$.

## FORTRAN 90 Interface

Generic: CSITG (A, B, BREAK, CSCOEF [, ...])
Specific: The specific interface names are S_CSITG and D_CSITG.

## FORTRAN 77 Interface

Single: CSITG (A, B, NINTV, BREAK, CSCOEF)
Double: The double precision function name is DCSITG.

## Description

The function CSITG evaluates the integral of a cubic spline over an interval. It is a special case of the routine PPITG, which evaluates the integral of a piecewise polynomial. (A cubic spline is a piecewise polynomial of order 4.)

## Example

This example computes a cubic spline interpolant to the function $x^{2}$ using CSINT and evaluates its integral over the intervals [0., .5] and [0., 2.]. Since CSINT uses the not-a knot condition, the interpolant reproduces $x^{2}$, hence the integral values are $1 / 24$ and $8 / 3$, respectively.

```
USE CSITG_INT
USE UMACH INT
USE CSINT_INT
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=10)
INTEGER I, NINTV, NOUT
REAL A, B, BREAK(NDATA), CSCOEF(4,NDATA), ERROR,&
            EXACT, F, FDATA(NDATA), FI, FLOAT, VALUE, X, &
            XDATA (NDATA)
INTRINSIC FLOAT
                                    Define function and integral
F(X) = X*X
FI(X) = X*X*X/3.0
DO 10 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT(NDATA-1)
    FDATA(I) = F(XDATA(I))
    1 0 ~ C O N T I N U E
                                    Compute cubic spline interpolant
CALL CSINT (XDATA, FDATA, BREAK, CSCOEF)
                                    Compute the integral of F over
                                    [0.0,0.5]
A = 0.0
B}=0.
NINTV = NDATA - 1
VALUE = CSITG(A,B,BREAK,CSCOEF)
EXACT = FI(B) - FI(A)
ERROR = EXACT - VALUE
                                    Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) A, B, VALUE, EXACT, ERROR
                                    Compute the integral of F over
                                    [0.0,2.0]
A = 0.0
B = 2.0
VALUE = CSITG (A,B,BREAK, CSCOEF)
EXACT = FI(B) - FI(A)
ERROR = EXACT - VALUE
WRITE (NOUT,99999) A, B, VALUE, EXACT, ERROR
FORMAT (' On the closed interval (', F3.1, ',', F3.1,&
    ') we have :', /, 1X, 'Computed Integral = ', F10.5, /,&
    1X, 'Exact Integral = ', F10.5, /, 1X, 'Error '&
    ,' ' = ', F10.6, /, /)
END
```


## Output

```
On the closed interval (0.0,0.5) we have :
Computed Integral = 0.04167
Exact Integral =0.04167
Error = 0.000000
On the closed interval (0.0,2.0) we have :
Computed Integral = 2.66666
Exact Integral = 2.66667
Error = 0.000006
```


## SPLEZ

Computes the values of a spline that either interpolates or fits user-supplied data.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissae. (Input) The data point abscissas must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
XVEC - Array of length N containing the points at which the spline function values are desired. (Input)
The entries of XVEC must be distinct.

VALUE - Array of length N containing the spline values. (Output)
$\operatorname{VALUE}(\mathrm{I})=S(\operatorname{XVEC}(\mathrm{I}))$ if $\operatorname{IDER}=0, \operatorname{VALUE}(\mathrm{I})=S^{\prime}(\operatorname{XVEC}(\mathrm{I}))$ if $\operatorname{IDER}=1$, and so forth, where $S$ is the computed spline.

## Optional Arguments

NDATA - Number of data points. (Input) Default: NDATA $=\operatorname{size}($ XDATA 1$)$.

All choices of ITYPE are valid if NDATA is larger than 6 . More specifically,

| NDATA $>$ ITYPE | or ITYPE $=1$. |
| :--- | :--- |
| NDATA $>3$ | for ITYPE $=2,3$. |
| NDATA $>(I T Y P E-3)$ | for ITYPE $=4,5,6,7,8$. |
| NDATA $>3$ | for ITYPE $=9,10,11,12$. |
| NDATA $>K O R D E R$ | for ITYPE $=13,14,15$. |

ITYPE - Type of interpolant desired. (Input)
Default: ITYPE $=1$.

## ITYPE

1 yields CSINT
2 yields CSAKM
3 yields CSCON
yields BSINT-BSNAK $K=2$
yields BSINT-BSNAK $\mathrm{K}=3$
yields BSINT-BSNAK $K=4$
yields BSINT-BSNAK $K=5$
yields BSINT-BSNAK $K=6$
yields CSSCV
yields BSLSQ K = 2
yields BSLSQ K = 3
yields BSLSQ $\mathrm{K}=4$
yields BSVLS $K=2$
yields BSVLS $K=3$
yields BSVLS $K=4$
IDER - Order of the derivative desired. (Input)
Default: $I D E R=0$.
$N$ - Number of function values desired. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{XVEC}, 1)$.

## FORTRAN 90 Interface

Generic: CALL SPLEZ (XDATA, FDATA, XVEC, VALUE [, ...])
Specific: The specific interface names are S_SPLEZ and D_SPLEZ.

## FORTRAN 77 Interface

Single: CALL SPLEZ (NDATA, XDATA, FDATA, ITYPE, IDER, N, XVEC, VALUE)
Double: The double precision name is DSPLEZ.

## Description

This routine is designed to let the user experiment with various interpolation and smoothing routines in the library.

The routine SPLEZ computes a spline interpolant to a set of data points $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots$, NDATA if ITYPE $=1, \ldots, 8$. If ITYPE $\geq 9$, various smoothing or least squares splines are computed. The output for this routine consists of a vector of values of the computed spline or its derivatives.
Specifically, let $i=\operatorname{IDER}, n=\mathrm{N}, v=\mathrm{XVEC}$, and $y=\mathrm{VALUE}$, then if $s$ is the computed spline we set

$$
y_{j}=s^{(i)}\left(v_{j}\right) \quad j=1, \ldots, n
$$

The routines called are listed above under the Itype heading. Additional documentation can be found by referring to these routines.

## Example

In this example, all the ITYPE parameters are exercised. The values of the spline are then compared with the exact function values and derivatives.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER NDATA, N
    PARAMETER (NDATA=21, N=2*NDATA-1)
                Specifications for local variables
    INTEGER I, IDER, ITYPE, NOUT
    REAL FDATA(NDATA), FPVAL(N), FVALUE (N),&
            VALUE(N), XDATA(NDATA), XVEC (N), EMAX1 (15), &
            EMAX2(15), X
                Specifications for intrinsics
    INTRINSIC FLOAT, SIN, COS
    REAL FLOAT, SIN, COS
    REAL F, FP
!
! Define a function
    F(X) = SIN (X*X)
    FP(X)=2*X*COS (X*X)
!
    CALL UMACH (2, NOUT)
!
    DO 10 I=1, NDATA
        XDATA(I) = 3.0*(FLOAT (I-1)/FLOAT (NDATA-1))
        FDATA(I) = F(XDATA(I))
    1 0 ~ C O N T I N U E
    DO 20 I=1, N
        XVEC(I) = 3.0*(FLOAT (I-1)/FLOAT (2*NDATA-2))
        FVALUE(I) = F(XVEC(I))
        FPVAL(I) = FP(XVEC(I))
    20 CONTINUE
!
    WRITE (NOUT,99999)
    DO 40 ITYPE=1, 15
        DO 30 IDER=0, 1
            CALL SPLEZ (XDATA, FDATA, XVEC, VALUE, ITYPE=ITYPE, &
                IDER=IDER)
                Compute the maximum error
```

```
                IF (IDER .EQ. O) THEN
                    CALL SAXPY (N, -1.0, FVALUE, 1, VALUE, 1)
                    EMAX1(ITYPE) = ABS(VALUE(ISAMAX(N,VALUE,1)))
                ELSE
                    CALL SAXPY (N, -1.0, FPVAL, 1, VALUE, 1)
                    EMAX2(ITYPE) = ABS(VALUE(ISAMAX (N,VALUE,1)))
                            END IF
    3 0
    CONTINUE
        WRITE (NOUT,'(I7,2F20.6)') ITYPE, EMAX1(ITYPE), EMAX2(ITYPE)
    4 0 ~ C O N T I N U E
!
99999 FORMAT (4X, 'ITYPE', 6X, 'Max error for f', 5X,&
    'Max error for f''', /)
END
```


## Output

| ITYPE | Max error for $f$ | Max error for $f^{\prime}$ |
| :--- | :---: | :---: |
| 1 | 0.014082 | 0.658018 |
| 2 | 0.024682 | 0.897757 |
| 3 | 0.020896 | 0.813228 |
| 4 | 0.083615 | 2.168083 |
| 5 | 0.010403 | 0.508043 |
| 6 | 0.014082 | 0.658020 |
| 7 | 0.004756 | 0.228858 |
| 8 | 0.001070 | 0.077159 |
| 9 | 0.020896 | 6.813228 |
| 10 | 0.392603 | 1.983959 |
| 11 | 0.162793 | 1.582624 |
| 12 | 0.045404 | 7.680381 |
| 13 | 0.588370 | 9.673786 |
| 14 | 0.752475 | 1.713031 |

## BSINT

Computes the spline interpolant, returning the B-spline coefficients.

## Required Arguments

NDATA - Number of data points. (Input)
XDATA - Array of length NDATA containing the data point abscissas. (Input)
FDATA - Array of length NDATA containing the data point ordinates. (Input)
KORDER - Order of the spline. (Input)
KORDER must be less than or equal to NDATA.
XKNOT - Array of length NDATA + KORDER containing the knot sequence. (Input) XKNOT must be nondecreasing.

BSCOEF - Array of length NDATA containing the B-spline coefficients. (Output)

## FORTRAN 90 Interface

Generic: CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
Specific: The specific interface names are S_BSINT and D_BSINT.

## FORTRAN 77 Interface

Single: CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
Double: The double precision name is DBSINT.

## Description

Following the notation in de Boor (1978, page 108), let $B_{j}=B_{j, k, t}$ denote the $j$-th B-spline of order $k$ with respect to the knot sequence $\mathbf{t}$. Then, BSINT computes the vector a satisfying

$$
\sum_{j=1}^{N} a_{j} B_{j}\left(x_{i}\right)=f_{i}
$$

and returns the result in $\mathrm{BSCOEF}=a$. This linear system is banded with at most $k-1$ subdiagonals and $k-1$ superdiagonals. The matrix

$$
A=\left(B_{j}\left(x_{i}\right)\right)
$$

is totally positive and is invertible if and only if the diagonal entries are nonzero. The routine BSINT is based on the routine SPLINT by de Boor (1978, page 204).

The routine BSINT produces the coefficients of the B-spline interpolant of order KORDER with knot sequence XKNOT to the data $\left(x_{i}, f_{i}\right)$ for $i=1$ to NDATA, where $x=$ XDATA and $f=$ FDATA. Let $\mathbf{t}=$ XKNOT,$k=$ KORDER, and $N=$ NDATA. First, BSINT sorts the XDATA vector and stores the result in $x$. The elements of the FDATA vector are permuted appropriately and stored in $f$, yielding the equivalent data $\left(x_{i}, f_{i}\right)$ for $i=1$ to $N$. The following preliminary checks are performed on the data. We verify that

$$
\begin{array}{rlrl}
x_{i}<x_{i+1} & i=1, \ldots, N-1 \\
\mathbf{t}_{i}<\mathbf{t}_{i+1} & i=1, \ldots, N \\
\mathbf{t}_{i} \leq \mathbf{t}_{i+k} & i=1, \ldots, N+k-1
\end{array}
$$

The first test checks to see that the abscissas are distinct. The second and third inequalities verify that a valid knot sequence has been specified.

In order for the interpolation matrix to be nonsingular, we also check $\mathbf{t}_{k} \leq x_{i} \leq \mathbf{t}_{N+1}$ for $i=1$ to $N$. This first inequality in the last check is necessary since the method used to generate the entries of the interpolation matrix requires that the $k$ possibly nonzero B -splines at $x_{i}$,

$$
B_{j-k+1}, \ldots, B_{j} \text { where } j \text { satisfies } \mathbf{t}_{j} \leq x_{i}<\mathbf{t}_{j+1}
$$

be well-defined (that is, $j-k+1 \geq 1$ ).
General conditions are not known for the exact behavior of the error in spline interpolation, however, if $\mathbf{t}$ and $x$ are selected properly and the data points arise from the values of a smooth (say $C^{k}$ ) function $f$, i.e. $f_{i}=f\left(x_{i}\right)$, then the error will behave in a predictable fashion. The maximum absolute error satisfies

$$
\|f-s\|_{\left[\mathbf{t}_{k}, \mathbf{t}_{N+1}\right]} \leq C\left\|f^{(k)}\right\|_{\left[\mathbf{t}_{k}, \mathbf{t}_{N+1}\right]}|\mathbf{t}|^{k}
$$

where

$$
|\mathbf{t}|:=\max _{i=k, \ldots, N}\left|\mathbf{t}_{i+1}-\mathbf{t}_{i}\right|
$$

For more information on this problem, see de Boor (1978, Chapter 13) and the references therein. This routine can be used in place of the IMSL routine CSINT by calling BSNAK to obtain the proper knots, then calling BSINT yielding the B-spline coefficients, and finally calling IMSL routine BSCPP to convert to piecewise polynomial form.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2INT/DB2INT. The reference is:
```
CALL B2INT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF, WK1,WK2,WK3,
IWK)
```

The additional arguments are as follows:
WK1 - Work array of length (5 * KORDER - 2) * NDATA.
WK2 - Work array of length NDATA.
WK3 - Work array of length NDATA.
IWK - Work array of length NDATA.
2. Informational errors

Type Code
31 The interpolation matrix is ill-conditioned.
43 The XDATA values must be distinct.
44 Multiplicity of the knots cannot exceed the order of the spline.
45 The knots must be nondecreasing.
415 The I-th smallest element of the data point array must be greater than the Ith knot and less than the (I + KORDER)-th knot.
416 The largest element of the data point array must be greater than the (NDATA)-th knot and less than or equal to the (NDATA + KORDER)-th knot.
417 The smallest element of the data point array must be greater than or equal to the first knot and less than the (KORDER +1 )st knot.
3. The spline can be evaluated using BSVAL, and its derivative can be evaluated using BSDER.

## Example

In this example, a spline interpolant $s$, to

$$
f(x)=\sqrt{x}
$$

is computed. The interpolated values are then compared with the exact function values using the IMSL routine BSVAL.

USE BSINT INT
USE BSNAK ${ }^{-}$INT
USE UMACH ${ }^{-}$INT
USE BSVAL_INT
IMPLICIT NONE
INTEGER KORDER, NDATA, NKNOT
PARAMETER (KORDER=3, NDATA=5, NKNOT=NDATA+KORDER)
!
INTEGER I, NCOEF, NOUT
REAL BSCOEF (NDATA), BT, F, FDATA (NDATA), FLOAT, \&
SQRT, X, XDATA (NDATA), XKNOT (NKNOT), XT
INTRINSIC FLOAT, SQRT
$\mathrm{F}(\mathrm{X})=\operatorname{SQRT}(\mathrm{X})$

DO $10 \mathrm{I}=1$, NDATA XDATA(I) = FLOAT (I-1)/FLOAT (NDATA-1) FDATA $(I)=F(X D A T A(I))$
10 CONTINUE
CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
Interpolate
CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
Get output unit number
CALL UMACH (2, NOUT)
Write heading
WRITE (NOUT, 99999)
NCOEF $=$ NDATA
$\mathrm{XT}=\mathrm{XDATA}(1)$
Print on a finer grid

Evaluate spline
,
WRITE (NOUT, 99998) XT, BT, F(XT) - BT
DO $20 \mathrm{I}=2$, NDATA
$\mathrm{XT}=(\mathrm{XDATA}(\mathrm{I}-1)+\mathrm{XDATA}(\mathrm{I})) / 2.0$
Evaluate spline
$B T=B S V A L(X T, K O R D E R, X K N O T, N C O E F, B S C O E F)$ WRITE (NOUT, 99998) XT, BT, F(XT) - BT $\mathrm{XT}=\mathrm{XDATA}(\mathrm{I})$

Evaluate spline
$B T=B S V A L(X T, K O R D E R, X K N O T, N C O E F, B S C O E F)$

```
        WRITE (NOUT,99998) XT, BT, F(XT) - BT
    20 CONTINUE
99998 FORMAT (' ', F6.4, 15X, F8.4, 12X, F11.6)
99999 FORMAT (/, 6X, 'X', 19X, 'S(X)', 18X, 'Error', /)
END
```


## Output

| $X$ | $S(X)$ | Error |
| ---: | ---: | ---: |
| 0.0000 | 0.0000 | 0.000000 |
| 0.1250 | 0.2918 | 0.061781 |
| 0.2500 | 0.5000 | 0.000000 |
| 0.3750 | 0.6247 | -0.012311 |
| 0.5000 | 0.7071 | 0.000000 |
| 0.6250 | 0.7886 | 0.002013 |
| 0.7500 | 0.8660 | 0.000000 |
| 0.8750 | 0.9365 | -0.001092 |
| 1.0000 | 1.0000 | 0.000000 |

## BSNAK

Computes the "not-a-knot" spline knot sequence.

## Required Arguments

NDATA - Number of data points. (Input)
XDATA - Array of length NDATA containing the location of the data points. (Input)

KORDER - Order of the spline. (Input)

XKNOT - Array of length NDATA + KORDER containing the knot sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)

Specific: The specific interface names are S_BSNAK and D_BSNAK.

## FORTRAN 77 Interface

Single: CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
Double: The double precision name is DBSNAK.

## Description

Given the data points $x=$ XDATA, the order of the spline $k=$ KORDER, and the number $N=$ NDATA of elements in XDATA, the subroutine BSNAK returns in $\mathbf{t}=$ XKNOT a knot sequence that is
appropriate for interpolation of data on $x$ by splines of order $k$. The vector $\mathbf{t}$ contains the knot sequence in its first $N+k$ positions. If $k$ is even and we assume that the entries in the input vector $x$ are increasing, then $\mathbf{t}$ is returned as
$\mathbf{t}_{i}=x_{1} \quad$ for $i=1, \ldots, k$
$\mathbf{t}_{i}=x_{i-k / 2} \quad$ for $i=k+1, \ldots, N$
$\mathbf{t}_{i}=x_{N}+\varepsilon \quad$ for $i=N+1, \ldots, N+k$
where $\varepsilon$ is a small positive constant. There is some discussion concerning this selection of knots in de Boor (1978, page 211). If $k$ is odd, then $\mathbf{t}$ is returned as

$$
\begin{aligned}
& \mathbf{t}_{i}=x_{1} \quad \text { for } i=1, \ldots, k \\
& \mathbf{t}_{i}=\left(x_{i-\frac{k-1}{2}}+x_{i-1-\frac{k-1}{2}}\right) / 2 \text { for } i=k+1, \ldots, N \\
& \mathbf{t}_{i}=x_{N}+\varepsilon \quad \text { for } i=N+1, \ldots, N+k
\end{aligned}
$$

It is not necessary to sort the values in $x$ since this is done in the routine BSNAK.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2NAK/DB2NAK. The reference is:

CALL B2NAK (NDATA, XDATA, KORDER, XKNOT, XSRT, IWK)
The additional arguments are as follows:
$\boldsymbol{X S R T}$ - Work array of length NDATA to hold the sorted XDATA values. If XDATA is not needed, XSRT may be the same as XDATA.

IWK - Work array of length NDATA to hold the permutation of XDATA.
2. Informational error

Type Code
$4 \quad 4 \quad$ The XDATA values must be distinct.
3. The first knot is at the left endpoint and the last knot is slightly beyond the last endpoint. Both endpoints have multiplicity KORDER.
4. Interior knots have multiplicity one.

## Example

In this example, we compute (for $k=3, \ldots, 8)$ six spline interpolants $s_{k}$ to $F(x)=\sin \left(10 x^{3}\right)$ on the interval $[0,1]$. The routine BSNAK is used to generate the knot sequences for $s_{k}$ and then BSINT is called to obtain the interpolant. We evaluate the absolute error

$$
\left|s_{k}-F\right|
$$

at 100 equally spaced points and print the maximum error for each $k$.
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER KMAX, KMIN, NDATA
PARAMETER (KMAX=8, KMIN=3, NDATA=20)
$!$
INTEGER I, K, KORDER, NOUT
REAL ABS, AMAX1, BSCOEF (NDATA), DIF, DIFMAX, F,\&
FDATA (NDATA), FLOAT, FT, SIN, ST, T, X, XDATA (NDATA), \& XKNOT (KMAX+NDATA), XT
INTRINSIC ABS, AMAX1, FLOAT, SIN
$F(X)=\operatorname{SIN}(10.0 * X * X * X)$
$T(X)=1.0-X * X$
DO $10 \mathrm{I}=1$, NDATA
XT $=$ FLOAT (I-1)/FLOAT (NDATA-1) XDATA $(I)=T(X T)$
$\operatorname{FDATA}(\mathrm{I})=\mathrm{F}(\mathrm{XDATA}(\mathrm{I}))$
10 CONTINUE
! Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999)
Write heading
Loop over different orders
DO 30 K=KMIN, KMAX
KORDER $=\mathrm{K}$
!
!
CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
Interpolate CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF) DIFMAX $=0.0$ DO 20 I=1, 100 $\mathrm{XT}=\mathrm{FLOAT}(\mathrm{I}-1) / 99.0$ Evaluate spline ST = BSVAL (XT, KORDER, XKNOT, NDATA, BSCOEF) $\mathrm{FT} \quad=\mathrm{F}(\mathrm{XT})$ DIF $=\mathrm{ABS}(\mathrm{FT}-\mathrm{ST})$ DIFMAX $=$ AMAX1 (DIF, DIFMAX)
CONTINUE
WRITE (NOUT,99998) KORDER, DIFMAX
CONTINUE
!
99998 FORMAT (' ', I3, 5X, F9.4)
99999 FORMAT (' KORDER', 5X, 'Maximum difference', /)
END

## Output

$\begin{array}{cc}\text { KORDER } & \text { Maximum difference } \\ 3 & 0.0080\end{array}$

| 4 | 0.0026 |
| :--- | :--- |
| 5 | 0.0004 |
| 6 | 0.0008 |
| 7 | 0.0010 |
| 8 | 0.0004 |

## BSOPK

Computes the "optimal" spline knot sequence.

## Required Arguments

NDATA - Number of data points. (Input)
XDATA - Array of length NDATA containing the location of the data points. (Input)
KORDER - Order of the spline. (Input)
$\boldsymbol{X K N O T}$ - Array of length NDATA + KORDER containing the knot sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL BSOPK (NDATA, XDATA, KORDER, XKNOT)

Specific: The specific interface names are S_BSOPK and D_BSOPK.

## FORTRAN 77 Interface

Single: CALL BSOPK (NDATA, XDATA, KORDER, XKNOT)
Double: The double precision name is DBSOPK.

## Description

Given the abscissas $x=$ XDATA for an interpolation problem and the order of the spline interpolant $k=$ KORDER, BSOPK returns the knot sequence $\mathbf{t}=$ XKNOT that minimizes the constant in the error estimate

$$
\|f-s\| \leq c\left\|f^{(k)}\right\|
$$

In the above formula, $f$ is any function in $C^{k}$ and $s$ is the spline interpolant to $f$ at the abscissas $x$ with knot sequence $\mathbf{t}$.

The algorithm is based on a routine described in de Boor (1978, page 204), which in turn is based on a theorem of Micchelli, Rivlin and Winograd (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2OPK/DB2OPK. The reference is:
```
CALL B2OPK (NDATA, XDATA, KORDER, XKNOT, MAXIT, WK, IWK)
```

The additional arguments are as follows:

MAXIT - Maximum number of iterations of Newton's Method. (Input) A suggested value is 10 .

WK - Work array of length (NDATA - KORDER) * (3 * KORDER - 2) + 6 * NDATA + 2 * KORDER + 5.

IWK - Work array of length NDATA.
2. Informational errors

Type Code

| 3 | 6 | Newton's method iteration did not converge. |
| :--- | :--- | :--- |
| 4 | 3 | The XDATA values must be distinct. |
| 4 | 4 | Interpolation matrix is singular. The XDATA values may be too close |
|  |  | together. |

3. The default value for MAXIT is 10 , this can be overridden by calling B2OPK/DB2OPK directly with a larger value.

## Example

In this example, we compute (for $k=3, \ldots, 8)$ six spline interpolants $s_{k}$ to $F(x)=\sin \left(10 x^{3}\right.$ ) on the interval $[0,1]$. The routine BSOPK is used to generate the knot sequences for $s_{k}$ and then BSINT is called to obtain the interpolant. We evaluate the absolute error

$$
\left|s_{k}-F\right|
$$

at 100 equally spaced points and print the maximum error for each $k$.

```
USE BSOPK INT
USE BSINT }\mp@subsup{}{}{-}\mathrm{ INT
USE UMACH INT
USE BSVAL_INT
IMPLICIT NONE
INTEGER KMAX, KMIN, NDATA
PARAMETER (KMAX=8, KMIN=3, NDATA=20)
INTEGER I, K, KORDER, NOUT
REAL ABS, AMAX1, BSCOEF(NDATA), DIF, DIFMAX, F,\delta
    FDATA(NDATA), FLOAT, FT, SIN, ST, T, X, XDATA(NDATA),&
    XKNOT (KMAX+NDATA), XT
INTRINSIC ABS, AMAX1, FLOAT, SIN
F(X) = SIN(10.0*X*X*X)
T(X) = 1.0 - X*X
Set up data
DO 10 I=1, NDATA
    XT = FLOAT (I-1) /FLOAT (NDATA-1)
```

!

```
        XDATA(I) = T(XT)
        FDATA(I) = F(XDATA(I))
    1 0 \text { CONTINUE}
!
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
    DO 30 K=KMIN, KMAX
        KORDER = K
    Generate knots
    CALL BSOPK (NDATA, XDATA, KORDER, XKNOT)
Interpolate
    CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    DIFMAX = 0.0
    DO 20 I=1, 100
        XT = FLOAT (I-1)/99.0
                            Evaluate spline
        ST = BSVAL(XT,KORDER,XKNOT,NDATA,BSCOEF)
        FT = F(XT)
        DIF = ABS (FT-ST)
                                    Compute maximum difference
                DIFMAX = AMAX1(DIF,DIFMAX)
    CONTINUE
        WRITE (NOUT,99998) KORDER, DIFMAX
    3 0 ~ C O N T I N U E ~
!
99998 FORMAT (' ', I3, 5X, F9.4)
99999 FORMAT (' KORDER', 5X, 'Maximum difference', /)
    END
```


## Output

KORDER Maximum difference
30.0096
$4 \quad 0.0018$
50.0005
$6 \quad 0.0004$
$7 \quad 0.0007$
$8 \quad 0.0035$

## BS2IN

Computes a two-dimensional tensor-product spline interpolant, returning the tensor-product Bspline coefficients.

## Required Arguments

XDATA - Array of length NXDATA containing the data points in the x -direction. (Input) XDATA must be strictly increasing.

YDATA - Array of length NYDATA containing the data points in the Y -direction. (Input) YDATA must be strictly increasing.

FDATA - Array of size NXDATA by NYDATA containing the values to be interpolated. (Input)
FDATA ( $I, \mathrm{~J}$ ) is the value at (XDATA (I), YDATA(J)).
KXORD - Order of the spline in the x -direction. (Input) KXORD must be less than or equal to NXDATA.

KYORD - Order of the spline in the Y-direction. (Input) KYORD must be less than or equal to NYDATA.

XKNOT - Array of length NXDATA + KXORD containing the knot sequence in the X -direction. (Input) XKNOT must be nondecreasing.

YKNOT - Array of length NYDATA + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.
BSCOEF - Array of length NXDATA * NYDATA containing the tensor-product B-spline coefficients. (Output)
BSCOEF is treated internally as a matrix of size NXDATA by NYDATA.

## Optional Arguments

NXDATA - Number of data points in the X -direction. (Input)
Default: NXDATA $=\operatorname{size}(X D A T A, 1)$.
NYDATA - Number of data points in the Y-direction. (Input) Default: NYDATA $=\operatorname{size}($ YDATA, 1$)$.
$\boldsymbol{L D F}$ - The leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input) Default: LDF $=\operatorname{size}($ FDATA, 1$)$.

## FORTRAN 90 Interface

Generic: CALL BS2IN (XDATA, YDATA, FDATA, KXORD, KYORD, XKNOT, YKNOT, BSCOEF [, ...])

Specific: The specific interface names are S_BS2IN and D_BS2IN.

## FORTRAN 77 Interface

Single: CALL BS2IN (NXDATA, XDATA, NYDATA, YDATA, FDATA, LDF, KXORD, KYORD, XKNOT, YKNOT, BSCOEF)

Double: The double precision name is DBS2IN.

## Description

The routine BS2 IN computes a tensor product spline interpolant. The tensor product spline interpolant to data $\left\{\left(x_{i}, y_{j}, f_{i j}\right)\right\}$, where $1 \leq i \leq N_{x}$ and $1 \leq j \leq N_{y}$, has the form

$$
\sum_{m=1}^{N_{y}} B_{n, k_{x}, \mathbf{t}_{x}}(x) B_{m, k_{y}, \mathbf{t}_{y}}(y)
$$

where $k_{x}$ and $k_{y}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD and KYORD, respectively.) Likewise, $\mathbf{t}_{x}$ and $\mathbf{t}_{y}$ are the corresponding knot sequences (XKNOT and YKNOT). The algorithm requires that

$$
\begin{array}{ll}
\mathbf{t}_{x}\left(k_{x}\right) \leq x_{i} \leq \mathbf{t}_{x}\left(N_{x}+1\right) & 1 \leq i \leq N_{x} \\
\mathbf{t}_{y}\left(k_{y}\right) \leq y_{j} \leq \mathbf{t}_{y}\left(N_{y}+1\right) & 1 \leq j \leq N_{y}
\end{array}
$$

Tensor product spline interpolants in two dimensions can be computed quite efficiently by solving (repeatedly) two univariate interpolation problems. The computation is motivated by the following observations. It is necessary to solve the system of equations

$$
\sum_{m=1}^{N_{v}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right) B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{j}\right)=f_{i j}
$$

Setting

$$
h_{m i}=\sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t x}}\left(x_{i}\right)
$$

we note that for each fixed $i$ from 1 to $N_{x}$, we have $N_{y}$ linear equations in the same number of unknowns as can be seen below:

$$
\sum_{m=1}^{N_{y}} h_{m i} B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{j}\right)=f_{i j}
$$

The same matrix appears in all of the equations above:

$$
\left[B_{m, k_{y}, t_{y}}\left(y_{j}\right)\right] \quad 1 \leq m, j \leq N_{y}
$$

Thus, we need only factor this matrix once and then apply this factorization to the $N_{x}$ righthand sides. Once this is done and we have computed $h_{m i}$, then we must solve for the coefficients $c_{n m}$ using the relation

$$
\sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right)=h_{m i}
$$

for $m$ from 1 to $N_{y}$, which again involves one factorization and $N_{y}$ solutions to the different righthand sides. The routine BS2IN is based on the routine SPLI2D by de Boor (1978, page 347).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B22IN/DB22IN. The reference is:
```
CALL B22IN (NXDATA, XDATA, NYDATA, YDATA, FDATA, LDF, KXORD, KYORD, XKNOT,
YKNOT, BSCOEF, WK, IWK)
```

The additional arguments are as follows:

```
WK - Work array of length NXDATA * NYDATA + MAX ((2 * KXORD -1)
    NXDATA. (2 * KYORD - 1) * NYDATA) + MAX((3 * KXORD - 2) *
    NXDATA. (3 * KYORD - 2) * NYDATA) + 2 * MAX(NXDATA. NYDATA).
```

$\boldsymbol{I W K}$ - Work array of length MAX (NXDATA . NYDATA) .
2. Informational errors

Type Code
31 Interpolation matrix is nearly singular. LU factorization failed.
32 Interpolation matrix is nearly singular. Iterative refinement failed.

46 The XDATA values must be strictly increasing.
$4 \quad 7$ The YDATA values must be strictly increasing.
413 Multiplicity of the knots cannot exceed the order of the spline.
$4 \quad 14$ The knots must be nondecreasing.
415 The I-th smallest element of the data point array must be greater than the I-th knot and less than the (I + K_ORD)-th knot.
416 The largest element of the data point array must be greater than the ( $\mathrm{N} \_$DATA)-th knot and less than or equal to the ( $\mathrm{N} \_$DATA $+\mathrm{K} \_$ORD)-th knot.
417 The smallest element of the data point array must be greater than or equal to the first knot and less than the ( K _ORD +1 )st knot.

## Example

In this example, a tensor product spline interpolant to a function $f$ is computed. The values of the interpolant and the error on a $4 \times 4$ grid are displayed.

```
USE BS2IN_INT
USE BSNAK_INT
USE BS2VL INT
USE UMACH_INT
IMPLICIT NONE
INTEGER KXORD, KYORD, LDF, NXDATA, NXKNOT, NXVEC, NYDATA, &
    NYKNOT, NYVEC
PARAMETER (KXORD=5, KYORD=2, NXDATA=21, NXVEC=4, NYDATA=6,&
```

```
                NYVEC=4, LDF=NXDATA, NXKNOT=NXDATA+KXORD,&
                NYKNOT=NYDATA+KYORD)
!
    INTEGER I, J, NOUT, NXCOEF, NYCOEF
    REAL BSCOEF (NXDATA,NYDATA), F, FDATA(LDF,NYDATA), FLOAT,&
            X, XDATA(NXDATA), XKNOT (NXKNOT), XVEC(NXVEC), Y,&
                YDATA(NYDATA), YKNOT (NYKNOT), YVEC(NYVEC),VL
    INTRINSIC FLOAT
    Define function
    F(X,Y) = X*X*X + X*Y
    DO 10 I=1, NXDATA
        XDATA(I) = FLOAT(I-11)/10.0
    CONTINUE
    Generate knot sequence
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
        Set up interpolation points
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/5.0
    CONTINUE
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
    Generate FDATA
    DO 40 I=1, NYDATA
        DO 30 J=1, NXDATA
            FDATA(J,I) = F(XDATA(J),YDATA(I))
        CONTINUE
    CONTINUE
        CALL BS2IN (XDATA, YDATA, FDATA, KXORD, KYORD, XKNOT, YKNOT,&
                BSCOEF)
        NXCOEF = NXDATA
        NYCOEF = NYDATA
            CALL UMACH (2, NOUT)
                Get output unit number
            WRITE (NOUT,99999)
            Write heading
                Print over a grid of
                    [0.0,1.0] x [0.0,1.0] at 16 points.
    DO 50 I=1, NXVEC
        XVEC(I) = FLOAT(I-1)/3.0
    5 0 ~ C O N T I N U E
    DO 60 I=1, NYVEC
        YVEC(I) = FLOAT(I-1)/3.0
    6 0 ~ C O N T I N U E ~
    DO }80\mathrm{ I=1, NXVEC
        DO 70 J=1, NYVEC
            VL = BS2VL (XVEC(I), YVEC(J), KXORD, KYORD, XKNOT,&
                YKNOT, NXCOEF, NYCOEF, BSCOEF)
            WRITE (NOUT,'(3F15.4,F15.6)') XVEC(I), YVEC(J),&
                VL, (F(XVEC(I),YVEC(J))-VL)
    7 0 ~ C O N T I N U E ~
    8 0 ~ C O N T I N U E ~
```

```
99999 FORMAT (13X, 'X', 14X, 'Y', 10X, 'S(X,Y)', 9X, 'Error')
```

    END
    
## Output

| $X$ | $Y$ | $S(X, Y)$ | Error |
| :---: | :---: | :---: | :---: |
| 0.0000 | 0.0000 | 0.0000 | 0.000000 |
| 0.0000 | 0.3333 | 0.0000 | 0.000000 |
| 0.0000 | 0.6667 | 0.0000 | 0.000000 |
| 0.0000 | 1.0000 | 0.0000 | 0.000000 |
| 0.3333 | 0.0000 | 0.0370 | 0.000000 |
| 0.3333 | 0.3333 | 0.1481 | 0.000000 |
| 0.3333 | 0.6667 | 0.2593 | 0.000000 |
| 0.3333 | 1.0000 | 0.3704 | 0.000000 |
| 0.6667 | 0.0000 | 0.2963 | 0.000000 |
| 0.6667 | 0.3333 | 0.5185 | 0.000000 |
| 0.6667 | 0.6667 | 0.7407 | 0.000000 |
| 0.6667 | 1.0000 | 0.9630 | 0.000000 |
| 1.0000 | 0.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.3333 | 1.3333 | 0.000000 |
| 1.0000 | 0.6667 | 1.6667 | 0.000000 |
| 1.0000 | 1.0000 | 2.0000 | 0.000000 |

## BS3IN

Computes a three-dimensional tensor-product spline interpolant, returning the tensor-product Bspline coefficients.

## Required Arguments

XDATA - Array of length NXDATA containing the data points in the $x$-direction. (Input) XDATA must be increasing.

YDATA - Array of length NYDATA containing the data points in the $y$-direction. (Input) YDATA must be increasing.

ZDATA - Array of length NZDATA containing the data points in the $z$-direction. (Input) ZDATA must be increasing.

FDATA - Array of size NXDATA by NYDATA by NZDATA containing the values to be interpolated. (Input)
FDATA (I, J, K) contains the value at (XDATA (I), YDATA(J), ZDATA(K)).
KXORD - Order of the spline in the $x$-direction. (Input) KXORD must be less than or equal to NXDATA.

KYORD - Order of the spline in the $y$-direction. (Input) KYORD must be less than or equal to NYDATA.

KZORD - Order of the spline in the $z$-direction. (Input) KZORD must be less than or equal to NZDATA.

XKNOT - Array of length NXDATA + KXORD containing the knot sequence in the $x$-direction. (Input)
XKNOT must be nondecreasing.
YKNOT - Array of length NYDATA + KYORD containing the knot sequence in the $y$-direction. (Input)
YKNOT must be nondecreasing.
ZKNOT - Array of length NZDATA + KZORD containing the knot sequence in the z-direction. (Input)
ZKNOT must be nondecreasing.
BSCOEF - Array of length NXDATA * NYDATA * NZDATA containing the tensor-product Bspline coefficients. (Output)
BSCOEF is treated internally as a matrix of size NXDATA by NYDATA by NZDATA.

## Optional Arguments

NXDATA - Number of data points in the $x$-direction. (Input) Default: NXDATA $=\operatorname{size}(\mathrm{XDATA}, 1)$.

NYDATA - Number of data points in the $y$-direction. (Input) Default: NYDATA $=\operatorname{size}($ YDATA, 1$)$.

NZDATA - Number of data points in the $z$-direction. (Input) Default: NZDATA $=\operatorname{size}($ ZDATA, 1$)$.
$\boldsymbol{L D F}$ - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
Default: LDF $=\operatorname{size}($ FDATA, 1$)$.
MDF - Middle dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
Default: MDF $=$ size (FDATA,2).

## FORTRAN 90 Interface

Generic: CALL BS3IN (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF [, ...])

Specific: The specific interface names are S_BS3IN and D_BS3IN.

## FORTRAN 77 Interface

Single: CALL BS3IN (NXDATA, XDATA, NYDATA, YDATA, NZDATA, ZDATA, FDATA, LDF, MDF, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF)

Double: The double precision name is DBS3IN.

## Description

The routine BS3IN computes a tensor-product spline interpolant. The tensor-product spline interpolant to data $\left\{\left(x_{i}, y_{j}, z_{k}, f_{i j k}\right)\right\}$, where $1 \leq i \leq N_{x}, 1 \leq j \leq N_{y}$, and $1 \leq k \leq N_{z}$ has the form

$$
\sum_{l=1}^{N_{z}} \sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}(x) B_{m, k_{y}, \mathbf{t}_{y}}(y) B_{l, k_{z}, \mathbf{t}_{z}}(z)
$$

where $k_{x}, k_{y}$, and $k_{z}$ are the orders of the splines (these numbers are passed to the subroutine in KXORD, KYORD, and KZORD, respectively). Likewise, $\mathbf{t}_{x}, \mathbf{t}_{y}$, and $\mathbf{t}_{z}$ are the corresponding knot sequences (XKNOT, YKNOT, and ZKNOT). The algorithm requires that

$$
\begin{aligned}
& \mathbf{t}_{x}\left(k_{x}\right) \leq x_{i} \leq \mathbf{t}_{x}\left(N_{x}+1\right) \quad 1 \leq i \leq N_{x} \\
& \mathbf{t}_{y}\left(k_{y}\right) \leq y_{j} \leq \mathbf{t}_{y}\left(N_{y}+1\right) \quad 1 \leq j \leq N_{y} \\
& \mathbf{t}_{z}\left(k_{z}\right) \leq z_{k} \leq \mathbf{t}_{z}\left(N_{z}+1\right) \quad 1 \leq k \leq N_{z}
\end{aligned}
$$

Tensor-product spline interpolants can be computed quite efficiently by solving (repeatedly) three univariate interpolation problems. The computation is motivated by the following observations. It is necessary to solve the system of equations

$$
\sum_{l=1}^{N_{z}} \sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right) B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{j}\right) B_{l, k_{z}, \mathbf{t}_{z}}\left(z_{k}\right)=f_{i j k}
$$

Setting

$$
h_{l i j}=\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right) B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{j}\right)
$$

we note that for each fixed pair $i j$ we have $N_{z}$ linear equations in the same number of unknowns as can be seen below:

$$
\sum_{l=1}^{N_{z}} h_{l i j} B_{l, k_{z}, t_{z}}\left(z_{k}\right)=f_{i j k}
$$

The same interpolation matrix appears in all of the equations above:

$$
\left[B_{l, k_{z}, \mathbf{t}_{z}}\left(z_{k}\right)\right] \quad 1 \leq l, k \leq N_{z}
$$

Thus, we need only factor this matrix once and then apply it to the $N_{x} N_{y}$ right-hand sides. Once this is done and we have computed $h_{l i j}$, then we must solve for the coefficients $c_{n m l}$ using the relation

$$
\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}\left(x_{i}\right) B_{m, k_{y}, \mathbf{t}_{y}}\left(y_{j}\right)=h_{l i j}
$$

that is the bivariate tensor-product problem addressed by the IMSL routine BS2IN. The interested reader should consult the algorithm description in the two-dimensional routine if more detail is desired. The routine BS3IN is based on the routine SPLI2D by de Boor (1978, page 347).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B23IN/DB23IN. The reference is:
```
CALL B23IN (NXDATA, XDATA, NYDATA, YDATA, NZDAYA, ZDATA, FDATA, LDF, MDF,
KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF, WK, IWK)
```

The additional arguments are as follows:

```
WK - Work array of length MAX ((2 * KXORD - 1) * NXDATA, (2 * KYORD
    - 1) * NYDATA, (2 * KZORD - 1) * NZDATA) + MAX((3 * KXORD -
    2) * NXDATA, ( 3 * KYORD - 2) * NYDATA + (3 * KZORD - 2) *
    NZDATA) + NXDATA * NYDATA *NZDATA + 2 * MAX(NXDATA, NYDATA,
    NZDATA).
```

IWK - Work array of length MAX (NXDATA, NYDATA, NZDATA) .
2. Informational errors
Type Code

| 3 | 1 | Interpolation matrix is nearly singular. LU factorization failed. |
| :---: | :---: | :---: |
| 3 | 2 | Interpolation matrix is nearly singular. Iterative refinement failed. |
| 4 | 13 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 14 | The knots must be nondecreasing. |
| 4 | 15 | The I-th smallest element of the data point array must be greater than the Ith knot and less than the ( $\mathrm{I}+\mathrm{K} \_\mathrm{ORD}$ )-th knot. |
| 4 | 16 | The largest element of the data point array must be greater than the (N_DATA)-th knot and less than or equal to the (N_DATA + K_ORD)-th knot. |
| 4 | 17 | The smallest element of the data point array must be greater than or equal to the first knot and less than the ( $\mathrm{K} \_$ORD +1 )st knot. |
| 4 | 18 | The XDATA values must be strictly increasing. |
| 4 | 19 | The YDATA values must be strictly increasing. |
| 4 | 20 | The zDATA values must be strictly increasing. |

## Example

In this example, a tensor-product spline interpolant to a function $f$ is computed. The values of the interpolant and the error on a $4 \times 4 \times 2$ grid are displayed.

```
USE BS3IN_INT
USE BSNAK_INT
```

```
    USE UMACH_INT
    USE BS3GD_INT
    IMPLICIT NONE
        SPECIFICATIONS FOR PARAMETERS
    INTEGER KXORD, KYORD, KZORD, LDF, MDF, NXDATA, NXKNOT, NXVEC, &
        NYDATA, NYKNOT, NYVEC, NZDATA, NZKNOT, NZVEC
    PARAMETER (KXORD=5, KYORD=2, KZORD=3, NXDATA=21, NXVEC=4,&
                NYDATA=6, NYVEC=4, NZDATA=8, NZVEC=2, LDF=NXDATA,&
                MDF=NYDATA, NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD, &
                NZKNOT=NZDATA+KZORD)
    INTEGER I, J, K, NOUT, NXCOEF, NYCOEF, NZCOEF
    REAL BSCOEF (NXDATA,NYDATA,NZDATA), F,&
        FDATA(LDF,MDF,NZDATA), FLOAT, VALUE (NXVEC,NYVEC,NZVEC) &
        , X, XDATA(NXDATA), XKNOT (NXKNOT), XVEC(NXVEC), Y,&
        YDATA(NYDATA), YKNOT (NYKNOT), YVEC (NYVEC), Z, &
        ZDATA(NZDATA), ZKNOT (NZKNOT), ZVEC(NZVEC)
    INTRINSIC FLOAT
        Define function.
        Set up X-interpolation points
    DO 10 I=1, NXDATA
        XDATA(I) = FLOAT (I-11)/10.0
    CONTINUE
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/FLOAT(NYDATA-1)
    CONTINUE
                                Set up Z-interpolation points
    DO 30 I=1, NZDATA
        ZDATA(I) = FLOAT (I-1)/FLOAT (NZDATA-1)
    CONTINUE
    knots
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
    CALL BSNAK (NZDATA, ZDATA, KZORD, ZKNOT)
                                    Generate FDATA
    DO 50 K=1, NZDATA
        DO 40 I=1, NYDATA
            DO 40 J=1, NXDATA
                FDATA(J,I,K) = F(XDATA(J),YDATA(I),ZDATA (K))
    CONTINUE
    CONTINUE
    CALL UMACH (2, NOUT)
    CALL BS3IN (XDATA, YDATA, ZDATA, FDATA, KXORD, &
        KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF)
    NXCOEF = NXDATA
    NYCOEF = NYDATA
    NZCOEF = NZDATA
        Write heading
    WRITE (NOUT,99999)
```

```
!
!
    Print over a grid of
    [-1.0,1.0] x [0.0,1.0] x [0.0,1.0]
    at 32 points.
    DO 60 I=1, NXVEC
            XVEC(I) = 2.0*(FLOAT(I-1)/3.0) - 1.0
        6 0
            DO 70 I=1, NYVEC
            YVEC(I) = FLOAT(I-1)/3.0
        7 0 ~ C O N T I N U E
        DO 80 I=1, NZVEC
            ZVEC(I) = FLOAT(I-1)
        80 CONTINUE
!
                        Call the evaluation routine.
    CALL BS3GD (0, 0, 0, XVEC, YVEC, ZVEC,&
                                    KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF, VALUE)
            DO 110 I=1, NXVEC
                DO 100 J=1, NYVEC
                DO 90 K=1, NZVEC
                    WRITE (NOUT,'(4F13.4, F13.6)') XVEC(I), YVEC(K),&
                        ZVEC(K), VALUE(I,J,K),&
                        F(XVEC(I),YVEC(J),ZVEC(K))&
                        - VALUE(I,J,K)
        90 CONTINUE
    100 CONTINUE
    1 1 0 ~ C O N T I N U E ~
99999 FORMAT (10X, 'X', 11X, 'Y', 10X, 'Z', 10X, 'S(X,Y,Z)', 7X,&
        'Error')
    END
```


## Output

| $X$ | $Y$ | $Z$ | Error |  |
| :---: | :---: | :---: | ---: | ---: |
| -1.0000 | 0.0000 | 0.0000 | -1.0000 | 0.000000 |
| -1.0000 | 0.3333 | 1.0000 | -1.0000 | 0.000000 |
| -1.0000 | 0.0000 | 0.0000 | -1.0000 | 0.000000 |
| -1.0000 | 0.3333 | 1.0000 | -1.3333 | 0.000000 |
| -1.0000 | 0.0000 | 0.0000 | -1.0000 | 0.000000 |
| -1.0000 | 0.3333 | 1.0000 | -1.6667 | 0.000000 |
| -1.0000 | 0.0000 | 0.0000 | -1.0000 | 0.000000 |
| -1.0000 | 0.3333 | 1.0000 | -2.0000 | 0.000000 |
| -0.3333 | 0.0000 | 0.0000 | -0.0370 | 0.000000 |
| -0.3333 | 0.3333 | 1.0000 | -0.0370 | 0.000000 |
| -0.3333 | 0.0000 | 0.0000 | -0.0370 | 0.000000 |
| -0.3333 | 0.3333 | 1.0000 | -0.1481 | 0.000000 |
| -0.3333 | 0.0000 | 0.0000 | -0.0370 | 0.000000 |
| -0.3333 | 0.3333 | 1.0000 | -0.2593 | 0.000000 |
| -0.3333 | 0.0000 | 0.0000 | -0.0370 | 0.000000 |
| -0.3333 | 0.3333 | 1.0000 | -0.3704 | 0.000000 |
| 0.3333 | 0.0000 | 0.0000 | 0.0370 | 0.000000 |
| 0.3333 | 0.3333 | 1.0000 | 0.0370 | 0.000000 |
| 0.3333 | 0.0000 | 0.0000 | 0.0370 | 0.000000 |
| 0.3333 | 0.3333 | 1.0000 | 0.1481 | 0.000000 |
| 0.3333 | 0.0000 | 0.0000 | 0.0370 | 0.000000 |
| 0.3333 | 0.3333 | 1.0000 | 0.2593 | 0.000000 |
| 0.3333 | 0.0000 | 0.0000 | 0.0370 | 0.000000 |


| 0.3333 | 0.3333 | 1.0000 | 0.3704 | 0.000000 |
| :--- | :--- | :--- | :--- | :--- |
| 1.0000 | 0.0000 | 0.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.3333 | 1.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.0000 | 0.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.3333 | 1.0000 | 1.3333 | 0.000000 |
| 1.0000 | 0.0000 | 0.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.3333 | 1.0000 | 1.6667 | 0.000000 |
| 1.0000 | 0.0000 | 0.0000 | 1.0000 | 0.000000 |
| 1.0000 | 0.3333 | 1.0000 | 2.0000 | 0.000000 |

## BSVAL

This function evaluates a spline, given its B-spline representation.

## Function Return Value

BSVAL - Value of the spline at X . (Output)

## Required Arguments

$\boldsymbol{X}$ - Point at which the spline is to be evaluated. (Input)
KORDER - Order of the spline. (Input)
XKNOT - Array of length KORDER + NCOEF containing the knot sequence. (Input) XKNOT must be nondecreasing.

NCOEF - Number of B-spline coefficients. (Input)
BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Input)

## FORTRAN 90 Interface

Generic: BSVAL (X, KORDER, XKNOT, NCOEF, BSCOEF)

Specific: The specific interface names are S_BSVAL and D_BSVAL.

## FORTRAN 77 Interface

Single: BSVAL (X, KORDER, XKNOT, NCOEF, BSCOEF)
Double: The double precision function name is DBSVAL.

## Description

The function BSVAL evaluates a spline (given its B-spline representation) at a specific point. It is a special case of the routine BSDER, which evaluates the derivative of a spline given its B-spline representation. The routine BSDER is based on the routine BVALUE by de Boor (1978, page 144).

Specifically, given the knot vector $\mathbf{t}$, the number of coefficients $N$, the coefficient vector $a$, and a point $x$, BSVAL returns the number

$$
\sum_{j=1}^{N} a_{j} B_{j, k}(x)
$$

where $B_{j, k}$ is the $j$-th B -spline of order $k$ for the knot sequence $\mathbf{t}$. Note that this function routine arbitrarily treats these functions as if they were right continuous near XKNOT(KORDER) and left continuous near XKNOT (NCOEF + 1). Thus, if we have KORDER knots stacked at the left or right end point, and if we try to evaluate at these end points, then we will get the value of the limit from the interior of the interval.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2VAL/DB2VAL. The reference is:

CALL B2VAL (X, KORDER, XKNOT, NCOEF, BSCOEF, WK1, WK2, WK3)
The additional arguments are as follows:
WK1 - Work array of length KORDER.
WK2 - Work array of length KORDER.

WK3 - Work array of length KORDER.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 4 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 5 | The knots must be nondecreasing. |

## Example

For an example of the use of BSVAL, see IMSL routine BSINT.

## BSDER

This function evaluates the derivative of a spline, given its B-spline representation.

## Function Return Value

BSDER - Value of the IDERIV-th derivative of the spline at X . (Output)

## Required Arguments

IDERIV - Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the value of the spline.
$\boldsymbol{X}$ - Point at which the spline is to be evaluated. (Input)
KORDER - Order of the spline. (Input)
XKNOT - Array of length NCOEF + KORDER containing the knot sequence. (Input) XKNOT must be nondecreasing.

NCOEF - Number of B-spline coefficients. (Input)
BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Input)

## FORTRAN 90 Interface

Generic: BSDER (IDERIV, X, KORDER, XKNOT, NCOEF, BSCOEF)

Specific: The specific interface names are S_BSDER and D_BSDER.

## FORTRAN 77 Interface

Single: BSDER (IDERIV, X, KORDER, XKNOT, NCOEF, BSCOEF)
Double: The double precision function name is DBSDER.

## Description

The function BSDER produces the value of a spline or one of its derivatives (given its B-spline representation) at a specific point. The function BSDER is based on the routine BVALUE by de Boor (1978, page 144).

Specifically, given the knot vector $\mathbf{t}$, the number of coefficients $N$, the coefficient vector $a$, the order of the derivative $i$ and a point $x$, BSDER returns the number

$$
\sum_{j=1}^{N} a_{j} B_{j, k}^{(i)}(x)
$$

where $B_{j, k}$ is the $j$-th B -spline of order $k$ for the knot sequence $\mathbf{t}$. Note that this function routine arbitrarily treats these functions as if they were right continuous near XKNOT(KORDER) and left continuous near XKNOT(NCOEF + 1). Thus, if we have KORDER knots stacked at the left or right end point, and if we try to evaluate at these end points, then we will get the value of the limit from the interior of the interval.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2DER/DB2DER. The reference is:

CALL B2DER (IDERIV, X, KORDER, XKNOT, NCOEF, BSCOEF, WK1, WK2, WK3)
The additional arguments are as follows:

> WK1 - Array of length KORDER.
> WK2 - Array of length KORDER.
> WK3 - Array of length KORDER.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 4 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 5 | The knots must be nondecreasing. |

## Example

A spline interpolant to the function

$$
f(x)=\sqrt{x}
$$

is constructed using BSINT. The B-spline representation, which is returned by the IMSL routine BSINT, is then used by BSDER to compute the value and derivative of the interpolant. The output consists of the interpolation values and the error at the data points and the midpoints. In addition, we display the value of the derivative and the error at these same points.

```
    USE BSDER INT
    USE BSINT 'INT
    USE BSNAK INT
    USE UMACH_}\mp@subsup{}{-}{-
    IMPLICIT NONE
    INTEGER KORDER, NDATA, NKNOT
    PARAMETER (KORDER=3, NDATA=5, NKNOT=NDATA+KORDER)
!
    INTEGER I, NCOEF, NOUT
    REAL BSCOEF (NDATA), BT0, BT1, DF, F, FDATA(NDATA),&
        FLOAT, SQRT, X, XDATA(NDATA), XKNOT (NKNOT), XT
    INTRINSIC FLOAT, SQRT
    F(X) = SQRT(X)
    DF(X) = 0.5/SQRT(X)
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I) /FLOAT(NDATA)
        FDATA(I) = F(XDATA(I))
    10 CONTINUE
    CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
    Interpolate
    CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
    NCOEF = NDATA
```

```
    XT = XDATA(1)
            Evaluate spline
    BT0 = BSDER(0,XT,KORDER,XKNOT,NCOEF,BSCOEF)
    BT1 = BSDER(1,XT,KORDER,XKNOT,NCOEF,BSCOEF)
    WRITE (NOUT,99998) XT, BT0, F(XT) - BT0, BT1, DF(XT) - BT1
    DO 20 I=2, NDATA
        XT = (XDATA (I-1) +XDATA(I))/2.0
                                    Evaluate spline
        BT0 = BSDER(0,XT,KORDER,XKNOT,NCOEF,BSCOEF)
        BT1 = BSDER(1,XT,KORDER,XKNOT,NCOEF,BSCOEF)
        WRITE (NOUT,99998) XT, BT0, F(XT) - BT0, BT1, DF(XT) - BT1
        XT = XDATA(I)
            Evaluate spline
        BT0 = BSDER(0,XT, KORDER,XKNOT,NCOEF,BSCOEF)
        BT1 = BSDER(1,XT,KORDER,XKNOT,NCOEF,BSCOEF)
        WRITE (NOUT,99998) XT, BT0, F(XT) - BT0, BT1, DF(XT) - BT1
    2 0 ~ C O N T I N U E
99998 FORMAT (' ', F6.4, 5X, F7.4, 3X, F10.6, 5X, F8.4, 3X, F10.6)
99999 FORMAT (6X, 'X', 8X, 'S(X)', 7X, 'Error', 8X, 'S''(X)', 8X,&
        'Error', /)
        END
```


## Output

| X | S(X) | Error | $S^{\prime}(X)$ | Error |
| ---: | ---: | ---: | ---: | ---: |
| 0.2000 | 0.4472 | 0.000000 | 1.0423 | 0.075738 |
| 0.3000 | 0.5456 | 0.002084 | 0.9262 | -0.013339 |
| 0.4000 | 0.6325 | 0.000000 | 0.8101 | -0.019553 |
| 0.5000 | 0.7077 | -0.000557 | 0.6940 | 0.013071 |
| 0.6000 | 0.7746 | 0.000000 | 0.6446 | 0.000869 |
| 0.7000 | 0.8366 | 0.000071 | 0.5952 | 0.002394 |
| 0.8000 | 0.8944 | 0.000000 | 0.5615 | -0.002525 |
| 0.9000 | 0.9489 | -0.000214 | 0.5279 | -0.000818 |
| 1.0000 | 1.0000 | 0.000000 | 0.4942 | 0.005814 |

## BS1GD

Evaluates the derivative of a spline on a grid, given its B-spline representation.

## Required Arguments

IDERIV — Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the value of the spline.
$X V E C$ - Array of length N containing the points at which the spline is to be evaluated. (Input)
XVEC should be strictly increasing.
KORDER - Order of the spline. (Input)

XKNOT - Array of length NCOEF + KORDER containing the knot sequence. (Input) XKNOT must be nondecreasing.

BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Input)

VALUE - Array of length N containing the values of the IDERIV-th derivative of the spline at the points in XVEC. (Output)

## Optional Arguments

$N$ - Length of vector XVEC. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{XVEC}, 1)$.
NCOEF - Number of B-spline coefficients. (Input)
Default: $\mathrm{NCOEF}=\operatorname{size}(\mathrm{BSCOEF}, 1)$.

## FORTRAN 90 Interface

Generic: CALL BS1GD (IDERIV, XVEC, KORDER, XKNOT, BSCOEF, VALUE [, ...])
Specific: The specific interface names are S_BS1GD and D_BS1GD.

## FORTRAN 77 Interface

Single: CALL BS1GD (IDERIV, N, XVEC, KORDER, XKNOT, NCOEF, BSCOEF, VALUE)
Double: The double precision name is DBS1GD.

## Description

The routine BSIGD evaluates a B-spline (or its derivative) at a vector of points. That is, given a vector $x$ of length $n$ satisfying $x_{i}<x_{i+1}$ for $i=1, \ldots, n-1$, a derivative value $j$, and a B-spline $s$ that is represented by a knot sequence and coefficient sequence, this routine returns the values

$$
s^{(j)}\left(x_{i}\right) \quad i=1, \ldots, n
$$

in the array VALUE. The functionality of this routine is the same as that of BSDER called in a loop, however BS1GD should be much more efficient. This routine converts the B-spline representation to piecewise polynomial form using the IMSL routine BSCPP, and then uses the IMSL routine PPVAL for evaluation.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B21GD/DB21GD. The reference is:

CALL B21GD (IDERIV, N, XVEC, KORDER, XKNOT, NCOEF, BSCOEF, VALUE, RWK1, RWK2, IWK3, RWK4, RWK5, RWK6)

The additional arguments are as follows:

```
RWK1 - Real array of length KORDER * (NCOEF - KORDER + 1).
RWK2 - Real array of length NCOEF - KORDER + 2.
IWK3 - Integer array of length N.
RWK4 - Real array of length N.
RWK5 - Real array of length N.
RWK6 - Real array of length (KORDER + 3) * KORDER
```

2. Informational error

| Type | Code |
| :---: | :---: | :---: |
| 4 | 5 | The points in XVEC must be strictly increasing.

## Example

To illustrate the use of BS1GD, we modify the example program for BSDER. In this example, a quadratic (order 3) spline interpolant to $F$ is computed. The values and derivatives of this spline are then compared with the exact function and derivative values. The routine BS1GD is based on the routines BSPLPP and PPVALU in de Boor (1978, page 89).

```
USE BS1GD INT
USE BSINT_INT
USE BSNAK_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER KORDER, NDATA, NKNOT, NFGRID
PARAMETER (KORDER=3, NDATA=5, NKNOT=NDATA+KORDER, NFGRID = 9)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, NCOEF, NOUT
REAL ANSO(NFGRID), ANS1 (NFGRID), BSCOEF (NDATA),&
    FDATA (NDATA),&
    X, XDATA(NDATA), XKNOT (NKNOT), XVEC(NFGRID)
                                SPECIFICATIONS FOR INTRINSICS
INTRINSIC FLOAT, SQRT
REAL FLOAT, SQRT
REAL DF, F
F(X) = SQRT(X)
DF(X) = 0.5/SQRT(X)
CALL UMACH (2, NOUT)
DO 10 I=1, NDATA
    XDATA(I) = FLOAT(I)/FLOAT(NDATA)
```

```
        FDATA(I) = F(XDATA(I))
    1 0 ~ C O N T I N U E ~
    CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
    CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    WRITE (NOUT,99999)
!
    NCOEF = NDATA
    XVEC(1) = XDATA(1)
    DO 20 I=2, 2*NDATA - 2, 2
        XVEC(I) = (XDATA(I/2+1)+XDATA (I/2))/2.0
        XVEC(I+1) = XDATA(I/2+1)
    20 CONTINUE
    CALL BS1GD (0, XVEC, KORDER, XKNOT, BSCOEF, ANSO)
    CALL BS1GD (1, XVEC, KORDER, XKNOT, BSCOEF, ANS1)
    DO 30 I=1, 2*NDATA - 1
        WRITE (NOUT,99998) XVEC(I), ANSO(I), F(XVEC(I)) - ANSO(I),&
                                    ANS1(I), DF(XVEC(I)) - ANS1(I)
    30 CONTINUE
99998 FORMAT (' ', F6.4, 5X, F7.4, 5X, F8.4, 5X, F8.4, 5X, F8.4)
99999 FORMAT (6X, 'X', 8X, 'S(X)', 7X, 'Error', 8X, 'S''(X)', 8X,&
            'Error', /)
    END
```


## Output

| $X$ | $S(X)$ | Error | $S^{\prime}(X)$ | Error |
| ---: | ---: | ---: | ---: | ---: |
| 0.2000 | 0.4472 | 0.0000 | 1.0423 | 0.0757 |
| 0.3000 | 0.5456 | 0.0021 | 0.9262 | -0.0133 |
| 0.4000 | 0.6325 | 0.0000 | 0.8101 | -0.0196 |
| 0.5000 | 0.7077 | -0.0006 | 0.6940 | 0.0131 |
| 0.6000 | 0.7746 | 0.0000 | 0.6446 | 0.0009 |
| 0.7000 | 0.8366 | 0.0001 | 0.5952 | 0.0024 |
| 0.8000 | 0.8944 | 0.0000 | 0.5615 | -0.0025 |
| 0.9000 | 0.9489 | -0.0002 | 0.5279 | -0.0008 |
| 1.0000 | 1.0000 | 0.0000 | 0.4942 | 0.0058 |

## BSITG

This function evaluates the integral of a spline, given its B-spline representation.

## Function Return Value

BSITG - Value of the integral of the spline from A to B. (Output)

## Required Arguments

$\boldsymbol{A}-$ Lower limit of integration. (Input)
$\boldsymbol{B}$ - Upper limit of integration. (Input)

KORDER - Order of the spline. (Input)
XKNOT - Array of length KORDER + NCOEF containing the knot sequence. (Input) XKNOT must be nondecreasing.

NCOEF - Number of B-spline coefficients. (Input)
BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Input)

## FORTRAN 90 Interface

Generic: BSITG (A, B, KORDER, XKNOT, NCOEF, BSCOEF)
Specific: The specific interface names are S_BSITG and D_BSITG.

## FORTRAN 77 Interface

Single: BSITG (A, B, KORDER, XKNOT, NCOEF, BSCOEF)
Double: The double precision function name is DBSITG.

## Description

The function BSITG computes the integral of a spline given its B-spline representation. Specifically, given the knot sequence $\mathbf{t}=$ XKNOT, the order $k=$ KORDER, the coefficients $a=$ BSCOEF,$n=$ NCOEF and an interval $[a, b]$, BSITG returns the value

$$
\int_{a}^{b} \sum_{i=1}^{n} a_{i} B_{i, k, \mathbf{t}}(x) d x
$$

This routine uses the identity (22) on page 151 of de Boor (1978), and it assumes that $\mathbf{t}_{1}=\ldots=\mathbf{t}_{k}$ and $\mathbf{t}_{n+1}=\ldots=\mathbf{t}_{n+k}$.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2ITG/DB2ITG. The reference is:

CALL B2ITG (A, B, KORDER, XKNOT, NCOEF, BSCOEF, TCOEF, AJ, DL, DR)
The additional arguments are as follows:
TCOEF - Work array of length KORDER +1.
$\boldsymbol{A} \boldsymbol{J}-$ Work array of length KORDER +1.
$\mathbf{D L}$ - Work array of length KORDER +1 .
$\boldsymbol{D R}$ - Work array of length KORDER +1.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 7 | The upper and lower endpoints of integration are equal. |
| 3 | 8 | The lower limit of integration is less than XKNOT(KORDER). |
| 3 | 9 | The upper limit of integration is greater than XKNOT(NCOEF+1). |
| 4 | 4 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 5 | The knots must be nondecreasing. |

## Example

We integrate the quartic $(k=5)$ spline that interpolates $x^{3}$ at the points $\{i / 10: i=-10, \ldots, 10\}$ over the interval $[0,1]$. The exact answer is $1 / 4$ since the interpolant reproduces cubic polynomials.

```
    USE BSITG INT
    USE BSNAK }\mp@subsup{}{}{-}\mathrm{ INT
    USE BSINT_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER KORDER, NDATA, NKNOT
    PARAMETER (KORDER=5, NDATA=21, NKNOT=NDATA+KORDER)
!
    INTEGER I, NCOEF, NOUT
    REAL A, B, BSCOEF (NDATA), ERROR, EXACT, F,&
            FDATA(NDATA), FI, FLOAT, VAL, X, XDATA(NDATA), &
            XKNOT (NKNOT)
    INTRINSIC FLOAT
    F(X) = X****
    FI(X) = X**4/4.0
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I-11)/10.0
        FDATA(I) = F(XDATA(I))
    10 CONTINUE
    CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
                    Interpolate
    CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    NCOEF = NDATA
    A = 0.0
    B}=1.
! Integrate from A to B
    VAL = BSITG(A,B,KORDER,XKNOT,NCOEF,BSCOEF)
    EXACT = FI(B) - FI(A)
    ERROR = EXACT - VAL
! Print results
    WRITE (NOUT,99999) A, B, VAL, EXACT, ERROR
99999 FORMAT (' On the closed interval (', F3.1, ',', F3.1,&
    ') we have :', /, 1X, 'Computed Integral = ', F10.5, /,&
```

```
                1X, 'Exact Integral = ', F10.5, /, 1X, 'Error '&
                , ' = ', F10.6, /, /)
END
```


## Output

```
On the closed interval (0.0,1.0) we have :
```

On the closed interval (0.0,1.0) we have :
Computed Integral = 0.25000
Computed Integral = 0.25000
Exact Integral = 0.25000
Exact Integral = 0.25000
Error = 0.000000

```
Error = 0.000000
```


## BS2VL

This function evaluates a two-dimensional tensor-product spline, given its tensor-product B-spline representation.

## Function Return Value

BS2VL - Value of the spline at (X, Y). (Output)

## Required Arguments

$\boldsymbol{X}$ - x -coordinate of the point at which the spline is to be evaluated. (Input)
$\boldsymbol{Y}$ - Y-coordinate of the point at which the spline is to be evaluated. (Input)

KXORD - Order of the spline in the x -direction. (Input)

KYORD - Order of the spline in the Y-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X-direction. (Input)
XKNOT must be nondecreasing.

YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.

NXCOEF - Number of B-spline coefficients in the X -direction. (Input)
NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
BSCOEF - Array of length NXCOEF * NYCOEF containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF

## FORTRAN 90 Interface

Generic: BS2VL (X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF)
Specific: The specific interface names are S_BS2VL and D_BS2VL.

## FORTRAN 77 Interface

Single: BS2VL (X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF)
Double: The double precision function name is DBS2VL.

## Description

The function BS2VL evaluates a bivariate tensor product spline (represented as a linear combination of tensor product B-splines) at a given point. This routine is a special case of the routine BS2DR, which evaluates partial derivatives of such a spline. (The value of a spline is its zero-th derivative.) For more information see de Boor (1978, pages 351-353).

This routine returns the value of the function $s$ at a point $(x, y)$ given the coefficients $c$ by computing

$$
s(x, y)=\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}(x) B_{m, k_{y}, t_{y}}(y)
$$

where $k_{x}$ and $k_{y}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD and KYORD, respectively.) Likewise, $\mathbf{t}_{x}$ and $\mathbf{t}_{y}$ are the corresponding knot sequences (XKNOT and YKNOT).

## Comments

Workspace may be explicitly provided, if desired, by use of B22VL/DB22VL. The reference is:

CALL B22VL (X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF, WK)
The additional argument is
WK - Work array of length 3 * MAX (KXORD, KYORD) + KYORD.

## Example

For an example of the use of BS2VL, see IMSL routine BS2IN.

## BS2DR

This function evaluates the derivative of a two-dimensional tensor-product spline, given its tensorproduct B-spline representation.

## Function Return Value

$\boldsymbol{B S} 2 \boldsymbol{D} \boldsymbol{R}$ - Value of the (IXDER, IYDER) derivative of the spline at (X, Y). (Output)

## Required Arguments

IXDER - Order of the derivative in the x -direction. (Input)

IYDER - Order of the derivative in the Y -direction. (Input)
$\boldsymbol{X}$ - X -coordinate of the point at which the spline is to be evaluated. (Input)
$\boldsymbol{Y}$ - Y -coordinate of the point at which the spline is to be evaluated. (Input)
KXORD - Order of the spline in the x -direction. (Input)
KYORD - Order of the spline in the Y -direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the $\mathrm{X}-$ direction. (Input)
XKNOT must be nondecreasing.

YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.
NXCOEF - Number of B-spline coefficients in the x -direction. (Input)
NYCOEF - Number of B-spline coefficients in the Y -direction. (Input)
BSCOEF - Array of length NXCOEF * NYCOEF containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF.

## FORTRAN 90 Interface

Generic: BS2DR (IXDER, IYDER, X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF)

Specific: The specific interface names are S_BS2DR and D_BS2DR.

## FORTRAN 77 Interface

Single: BS2DR (IXDER, IYDER, X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF)

Double: The double precision function name is DBS2DR.

## Description

The routine BS2DR evaluates a partial derivative of a bivariate tensor-product spline (represented as a linear combination of tensor product B-splines) at a given point; see de Boor (1978, pages 351-353).

This routine returns the value of $\mathrm{s}^{(p, q)}$ at a point $(x, y)$ given the coefficients $c$ by computing

$$
s^{(p, q)}(x, y)=\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, t_{x}}^{(p)}(x) B_{m, k_{y}, t_{y}}^{(q)}(y)
$$

where $k_{x}$ and $k_{y}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD and KYORD, respectively.) Likewise, $\mathbf{t}_{x}$ and $\mathbf{t}_{y}$ are the corresponding knot sequences (XKNOT and YKNOT).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B22DR/DB22DR. The reference is:

CALL B22DR(IXDER, IYDER, X, Y, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF, WK)

The additional argument is:
WK - Work array of length 3 * MAX (KXORD, KYORD) + KYORD.
2. Informational errors

Type Code
31 The point X does not satisfy XKNOT(KXORD) .LE. X.LE. XKNOT(NXCOEF + 1).
32 The point $Y$ does not satisfy
YKNOT(KYORD) .LE. Y.LE. YKNOT(NYCOEF + 1).

## Example

In this example, a spline interpolant $s$ to a function $f$ is constructed. We use the IMSL routine BS2IN to compute the interpolant and then BS2DR is employed to compute $s^{(2,1)}(x, y)$. The values of this partial derivative and the error are computed on a $4 \times 4$ grid and then displayed.

```
USE BS2DR INT
USE BSNAK INT
USE UMACH INT
USE BS2IN-INT
IMPLICIT NONE
! SPECIFICATIONS FOR PARAMETERS
INTEGER KXORD, KYORD, LDF, NXDATA, NXKNOT, NYDATA, NYKNOT
PARAMETER (KXORD=5, KYORD=3, NXDATA=21, NYDATA=6, LDF=NXDATA,&
    NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD)
```

!

```
    INTEGER I, J, NOUT, NXCOEF, NYCOEF
    REAL BSCOEF (NXDATA,NYDATA), F, F21,&
        FDATA(LDF,NYDATA), FLOAT, S21, X, XDATA(NXDATA), &
        XKNOT (NXKNOT), Y, YDATA(NYDATA), YKNOT (NYKNOT)
    INTRINSIC FLOAT
! Define function and (2,1) derivative
    F(X,Y) = X*X*X*X + X*X*X*Y*Y
    F21(X,Y) = 12.0*X*Y
    DO 10 I=1, NXDATA
        XDATA(I) = FLOAT(I-11)/10.0
    10 CONTINUE
! Generate knot sequence
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
    Set up interpolation points
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/5.0
    20 CONTINUE
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
        Generate FDATA
    DO 40 I=1, NYDATA
        DO 30 J=1, NXDATA
            FDATA(J,I) = F(XDATA(J),YDATA(I))
        CONTINUE
    CONTINUE
    CALL BS2IN (XDATA, YDATA, FDATA, KXORD, KYORD, XKNOT, &
            YKNOT, BSCOEF)
    NXCOEF = NXDATA
    NYCOEF = NYDATA
! Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
        Write heading
        Print (2,1) derivative over a
        grid of [0.0,1.0] x [0.0,1.0]
        at }16\mathrm{ points.
    DO 60 I=1, 4
        DO 50 J=1, 4
        X = FLOAT (I-1)/3.0
        Y = FLOAT (J-1)/3.0
            Evaluate spline
        S21 = BS2DR (2,1,X,Y,KXORD,KYORD,XKNOT,YKNOT,NXCOEF,NYCOEF, &
                BSCOEF)
        WRITE (NOUT,'(3F15.4, F15.6)') X, Y, S21, F21(X,Y) - S21
    50 CONTINUE
    6 0 ~ C O N T I N U E ~
99999 FORMAT (39X, '(2,1)', /, 13X, 'X', 14X, 'Y', 10X, 'S (X,Y)',&
                5X, 'Error')
    END
```


## Output

|  | $(2,1)$ |  |  |
| ---: | ---: | ---: | ---: |
| X | Y | $(\mathrm{S}, \mathrm{Y})$ | Error |
| 0.0000 | 0.0000 | 0.0000 | 0.000000 |
| 0.0000 | 0.3333 | 0.0000 | 0.000000 |
| 0.0000 | 0.6667 | 0.0000 | 0.000000 |
| 0.0000 | 1.0000 | 0.0000 | 0.000001 |
| 0.3333 | 0.0000 | 0.0000 | 0.000000 |
| 0.3333 | 0.3333 | 1.3333 | 0.000002 |
| 0.3333 | 0.6667 | 2.6667 | -0.000002 |
| 0.3333 | 1.0000 | 4.0000 | 0.000008 |
| 0.6667 | 0.0000 | 0.0000 | 0.000006 |
| 0.6667 | 0.3333 | 2.6667 | -0.000011 |
| 0.6667 | 0.6667 | 5.3333 | 0.000028 |
| 0.6667 | 1.0000 | 8.0001 | -0.000134 |
| 1.0000 | 0.0000 | -0.0004 | 0.000439 |
| 1.0000 | 0.3333 | 4.0003 | -0.000319 |
| 1.0000 | 0.6667 | 7.9996 | 0.000363 |
| 1.0000 | 1.0000 | 12.0005 | -0.000458 |

## BS2GD

Evaluates the derivative of a two-dimensional tensor-product spline, given its tensor-product B -spline representation on a grid.

## Required Arguments

IXDER - Order of the derivative in the x -direction. (Input)
IYDER - Order of the derivative in the Y -direction. (Input)
$\boldsymbol{X V E C}$ - Array of length NX containing the x -coordinates at which the spline is to be evaluated. (Input)
The points in XVEC should be strictly increasing.
YVEC - Array of length NY containing the Y-coordinates at which the spline is to be evaluated. (Input)
The points in YVEC should be strictly increasing.
KXORD - Order of the spline in the X -direction. (Input)
KYORD - Order of the spline in the Y-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X-direction. (Input)
XKNOT must be nondecreasing.
YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.

BSCOEF - Array of length NXCOEF * NYCOEF containing the tensor-product B-spline coefficients. (Input)
bSCOEF is treated internally as a matrix of size nXCOEF by NYCOEF.
VALUE - Value of the (IXDER, IYDER) derivative of the spline on the NX by NY grid. (Output)
VALUE $(I, J)$ contains the derivative of the spline at the point (XVEC(I), YVEC(J)).

## Optional Arguments

$\boldsymbol{N X}$ - Number of grid points in the x -direction. (Input)
Default: NX = size (XVEC,1).
$N Y$ - Number of grid points in the Y-direction. (Input)
Default: NY = size (YVEC,1).
NXCOEF - Number of B-spline coefficients in the X -direction. (Input)
Default: NXCOEF $=$ size $($ XKNOT, 1$)-$ KXORD.
NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
Default: NYCOEF $=$ size $($ YKNOT, 1$)-$ KYORD.
LDVALU - Leading dimension of VALUE exactly as specified in the dimension statement of the calling program. (Input)
Default: LDVALU $=$ SIZE $($ VALUE, 1$)$.

## FORTRAN 90 Interface

Generic: CALL BS2GD (IXDER, IDER, XVEC, YVEC, KXORD, KYORD, XKNOT, YKNOT, BSCOEF, VALUE [, ...])

Specific: The specific interface names are S_BS2GD and D_BS2GD.

## FORTRAN 77 Interface

Single: CALL BS2GD (IXDER, IYDER, NX, XVEC, NY, YVEC, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF, VALUE, LDVALU)

Double: The double precision name is DBS2GD.

## Description

The routine BS2GD evaluates a partial derivative of a bivariate tensor-product spline (represented as a linear combination of tensor-product B-splines) on a grid of points; see de Boor (1978, pages 351-353).

This routine returns the values of $\mathrm{s}^{(p, q)}$ on the $\operatorname{grid}\left(x_{i}, y_{j}\right)$ for $i=1, \ldots, n x$ and $j=1, \ldots, n y$ given the coefficients $c$ by computing (for all $(x, y)$ in the grid)

$$
s^{(p, q)}(x, y)=\sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m} B_{n, k_{x}, \mathbf{t}_{x}}^{(p)}(x) B_{m, k_{y}, \mathbf{t}_{y}}^{(q)}(y)
$$

where $k_{x}$ and $k_{y}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD and KYORD, respectively.) Likewise, $\mathbf{t}_{x}$ and $\mathbf{t}_{y}$ are the corresponding knot sequences (XKNOT and YКNOT). The grid must be ordered in the sense that $x_{i}<x_{i+1}$ and $y_{j}<y_{j+l}$.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B22GD/DB22GD. The reference is:
```
CALL B22GD (IXDER, IYDER, NX, XVEC, NY, YVEC, KXORD, KYORD, XKNOT, YKNOT,
NXCOEF, NYCOEF, BSCOEF, VALUE, LDVALU, LEFTX, LEFTY, A, B, DBIATX, DBIATY,
BX,BY)
```

The additional arguments are as follows:
LEFTX - Integer work array of length NX.
LEFTY - Integer work array of length NY.
A- Work array of length KXORD * KXORD.
$\boldsymbol{B}$ - Work array of length KYORD * KYORD.
DBIATX - Work array of length KXORD * (IXDER + 1).
DBIATY - Work array of length KYORD * (IYDER + 1).
$\boldsymbol{B X}$ - Work array of length KXORD * NX.
BY - Work array of length KYORD * NY.
2 Informational errors
Type Code
$31 \quad \operatorname{XVEC}(\mathrm{I})$ does not satisfy
XKNOT (KXORD) . LE. XVEC(I) .LE. XKNOT(NXCOEF + 1)
$32 \operatorname{YVEC}(\mathrm{I})$ does not satisfy
YKNOT (KYORD) .LE. YVEC(I) .LE. YKNOT(NYCOEF + 1)
43 XVEC is not strictly increasing.
44 YVEC is not strictly increasing.

## Example

In this example, a spline interpolant $s$ to a function $f$ is constructed. We use the IMSL routine BS2IN to compute the interpolant and then BS2GD is employed to compute $s^{(2, l)}(x, y)$ on a grid. The values of this partial derivative and the error are computed on a $4 \times 4$ grid and then displayed.

```
USE BS2GD INT
USE BS2IN }\mp@subsup{}{}{-}\mathrm{ INT
USE BSNAK_INT
USE UMACH_INT
IMPLICIT NONE
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, J, KXORD, KYORD, LDF, NOUT, NXCOEF, NXDATA, &
NYCOEF, NYDATA
REAL DCCFD}(21,6), DOCBSC(21,6), DOCXD(21), DOCXK(26),&
DOCYD(6), DOCYK(9), F, F21, FLOAT, VALUE(4,4),&
X, XVEC(4), Y, YVEC(4)
INTRINSIC FLOAT
F(X,Y) = X*X*X*X + X*X*X*Y*Y
F21(X,Y) = 12.0*X*Y
yj Initialize/Setup
CALL UMACH (2, NOUT)
KXORD = 5
KYORD = 3
NXDATA = 21
NYDATA = 6
LDF = NXDATA
Set up interpolation points
DO 10 I=1, NXDATA
        DOCXD(I) = FLOAT(I-11)/10.0
CONTINUE
Set up interpolation points
DO 20 I=1, NYDATA
        DOCYD(I) = FLOAT(I-1)/5.0
    CONTINUE
CALL BSNAK (NXDATA, DOCXD, KXORD, DOCXK)
Generate knot sequence
CALL BSNAK (NYDATA, DOCYD, KYORD, DOCYK)
                                    Generate FDATA
DO 40 I=1, NYDATA
        DO 30 J=1, NXDATA
            DCCFD(J,I) = F(DOCXD(J),DOCYD(I))
    CONTINUE
CONTINUE
CALL BS2IN (DOCXD, DOCYD, DCCFD, KXORD, KYORD, &
DOCXK, DOCYK, DOCBSC)
Print (2,1) derivative over a
grid of [0.0,1.0] x [0.0,1.0]
at 16 points.
NXCOEF = NXDATA
NYCOEF = NYDATA
```

```
    WRITE (NOUT,99999)
    DO 50 I=1, 4
        XVEC(I) = FLOAT(I-1)/3.0
        YVEC(I) = XVEC(I)
    5 0 ~ C O N T I N U E
    CALL BS2GD (2, 1, XVEC, YVEC, KXORD, KYORD, DOCXK, DOCYK,&
                DOCBSC, VALUE)
    DO 70 I=1, 4
        DO 60 J=1, 4
            WRITE (NOUT,'(3F15.4,F15.6)') XVEC(I), YVEC(J),&
                                    VALUE (I,J),&
                                    F21(XVEC(I), YVEC(J)) -&
                            VALUE(I,J)
    6 0 ~ C O N T I N U E ~
    7 0 ~ C O N T I N U E ~
99999 FORMAT (39X, '(2,1)', /, 13X, 'X', 14X, 'Y', 10X, 'S (X,Y)',&
                5x, 'Error')
    END
```


## Output

|  | $(2,1)$ |  |  |
| :---: | :---: | :---: | :---: |
| X | Y | S (X,Y) | Error |
| 0.0000 | 0.0000 | 0.0000 | 0.000000 |
| 0.0000 | 0.3333 | 0.0000 | 0.000000 |
| 0.0000 | 0.6667 | 0.0000 | 0.000000 |
| 0.0000 | 1.0000 | 0.0000 | 0.000001 |
| 0.3333 | 0.0000 | 0.0000 | -0.000001 |
| 0.3333 | 0.3333 | 1.3333 | 0.000001 |
| 0.3333 | 0.6667 | 2.6667 | -0.000004 |
| 0.3333 | 1.0000 | 4.0000 | 0.000008 |
| 0.6667 | 0.0000 | 0.0000 | -0.000001 |
| 0.6667 | 0.3333 | 2.6667 | -0.000008 |
| 0.6667 | 0.6667 | 5.3333 | 0.000038 |
| 0.6667 | 1.0000 | 8.0001 | -0.000113 |
| 1.0000 | 0.0000 | -0.0005 | 0.000488 |
| 1.0000 | 0.3333 | 4.0004 | -0.000412 |
| 1.0000 | 0.6667 | 7.9995 | 0.000488 |
| 1.0000 | 1.0000 | 12.0002 | -0.000244 |

## BS2IG

This function evaluates the integral of a tensor-product spline on a rectangular domain, given its tensor-product B-spline representation.

## Function Return Value

BS2IG - Integral of the spline over the rectangle (A, B) by (C, D). (Output)

## Required Arguments

A - Lower limit of the x -variable. (Input)
B - Upper limit of the x -variable. (Input)
$\boldsymbol{C}$ - Lower limit of the Y-variable. (Input)

D - Upper limit of the Y-variable. (Input)
KXORD - Order of the spline in the x -direction. (Input)
KYORD - Order of the spline in the Y-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X-direction. (Input) XKNOT must be nondecreasing.

YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.
BSCOEF - Array of length NXCOEF * NYCOEf containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF.

## Optional Arguments

NXCOEF - Number of B-spline coefficients in the x -direction. (Input)
Default: NXCOEF $=$ size $($ XKNOT, 1$)-$ KXORD.

NYCOEF - Number of B-spline coefficients in the Y-direction. (Input) Default: NYCOEF = size (YKNOT,1) - KYORD.

## FORTRAN 90 Interface

Generic: BS2IG (A, B, C, D, KXORD, KYORD, XKNOT, YKNOT, BSCOEF [, ...])

Specific: The specific interface names are S_BS2IG and D_BS2IG.

## FORTRAN 77 Interface

Single: BS2IG (A, B, C, D, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF)
Double: The double precision function name is DBS2IG.

## Description

The function BS2IG computes the integral of a tensor-product two-dimensional spline given its B-spline representation. Specifically, given the knot sequence $\mathbf{t}_{x}=$ XKNOT, $\mathbf{t}_{y}=$ YKNOT, the order $k_{x}=$ KXORD, $k_{y}=$ KYORD, the coefficients $\beta=$ BSCOEF, the number of coefficients $n_{x}=$ NXCOEF, $n_{y}=$ NYCOEF and a rectangle $[a, b]$ by $[c, d]$, BS 2 IG returns the value

$$
\int_{a}^{b} \int_{c}^{d} \sum_{i=1}^{n_{x}} \sum_{j=1}^{n_{y}} \beta_{i j} B_{i j} d y d x
$$

where

$$
B_{i, j}(x, y)=B_{i, k_{x}, \mathbf{t}_{x}}(x) B_{j, k_{y}, t_{y}}(y)
$$

This routine uses the identity (22) on page 151 of de Boor (1978). It assumes (for all knot sequences) that the first and last $k$ knots are stacked, that is, $t_{1}=\ldots=t_{k}$ and $t_{n+l}=\ldots=t_{n+k}$, where $k$ is the order of the spline in the $x$ or $y$ direction.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B22IG/DB22IG. The reference is:

CALL B22IG (A, B, C, D, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, BSCOEF, WK)

The additional argument is:
WK - Work array of length 4 * (MAX (KXORD, KYORD) + 1) + NYCOEF.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | The lower limit of the X-integration is less than XKNOT(KXORD). |
| 3 | 2 | The upper limit of the X-integration is greater than XKNOT(NXCOEF + <br> 1). |
| 3 | 3 | The lower limit of the Y-integration is less than YKNOT(KYORD). |
| 3 | 4 | The upper limit of the Y-integration is greater than YKNOT(NYCOEF + |
|  |  | 1). |
| 4 | 13 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 14 | The knots must be nondecreasing. |

## Example

We integrate the two-dimensional tensor-product quartic $\left(k_{x}=5\right)$ by linear $\left(k_{y}=2\right)$ spline that interpolates $x^{3}+x y$ at the points $\{(i / 10, j / 5): i=-10, \ldots, 10$ and $j=0, \ldots, 5\}$ over the rectangle $[0,1] \times[.5,1]$. The exact answer is $5 / 16$.

```
USE BS2IG INT
USE BSNAK}\mp@subsup{}{}{-}\mathrm{ INT
USE BS2IN_INT
```

```
USE UMACH_INT
    IMPLICIT NONE
    TNTEGER
    PARAMETER (KXORD=5, KYORD=2, NXDATA=21, NYDATA=6, LDF=NXDATA,&
                NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD)
!
    INTEGER I, J, NOUT, NXCOEF, NYCOEF
    REAL A, B, BSCOEF (NXDATA,NYDATA), C , D, F,&
        FDATA(LDF,NYDATA), FI, FLOAT, VAL, X, XDATA(NXDATA),&
        XKNOT (NXKNOT), Y, YDATA(NYDATA), YKNOT (NYKNOT)
    INTRINSIC FLOAT
    Define function and integral
    F(X,Y) = X*X*X + X*Y
    FI(A,B,C ,D) = .25*((B**4-A**4)*(D-C ) +(B*B-A*A)*(D*D-C *C ))
                                    Set up interpolation points
    DO 10 I=1, NXDATA
        XDATA(I) = FLOAT(I-11)/10.0
    CONTINUE
                            Generate knot sequence
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
                            Set up interpolation points
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/5.0
    CONTINUE
                                    Generate knot sequence
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
                            Generate FDATA
    DO 40 I=1, NYDATA
        DO 30 J=1, NXDATA
            FDATA(J,I) = F(XDATA(J),YDATA(I))
    CONTINUE
    CONTINUE
    CALL BS2IN (XDATA, YDATA, FDATA, KXORD, &
        KYORD, XKNOT, YKNOT, BSCOEF)
                                    Integrate over rectangle
                                    [0.0,1.0] x [0.0,0.5]
    NXCOEF = NXDATA
    NYCOEF = NYDATA
    A}=0.
    B}=1.
    C = 0.5
    D = 1.0
    VAL = BS2IG(A,B,C ,D,KXORD,KYORD,XKNOT,YKNOT,BSCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    Print results
    WRITE (NOUT,99999) VAL, FI(A,B,C ,D), FI(A,B,C ,D) - VAL
    FORMAT (' Computed Integral = ', F10.5, /, ' Exact Integral '&
    , '= ', F10.5, /, ' Error '&
    , '= ', F10.6, /)
```

    END
    
## Output

```
Computed Integral = 0.31250
Exact Integral = 0.31250
Error = 0.000000
```

This function Evaluates a three-dimensional tensor-product spline, given its tensor-product Bspline representation.

## Function Return Value

BS3VL - Value of the spline at (X, Y, Z). (Output)

## Required Arguments

$\boldsymbol{X}$ - X -coordinate of the point at which the spline is to be evaluated. (Input)
$\boldsymbol{Y}$ - y-coordinate of the point at which the spline is to be evaluated. (Input)
$\mathbf{Z}$ - z-coordinate of the point at which the spline is to be evaluated. (Input)
KXORD - Order of the spline in the x -direction. (Input)
KYORD - Order of the spline in the Y-direction. (Input)
KZORD - Order of the spline in the z-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X-direction. (Input)
XKNOT must be nondecreasing.
YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.
ZKNOT - Array of length NZCOEF + KZORD containing the knot sequence in the z-direction. (Input) ZKNOT must be nondecreasing.

NXCOEF - Number of B-spline coefficients in the X-direction. (Input)
NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
NZCOEF - Number of B-spline coefficients in the z-direction. (Input)

BSCOEF - Array of length NXCOEF * NYCOEF * NZCOEF containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF by NZCOEF.

## FORTRAN 90 Interface

Generic: BS3VL (X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Specific: The specific interface names are S_BS3VL and D_BS3VL.

## FORTRAN 77 Interface

Single: BS3VL (X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Double: The double precision function name is DBS3VL.

## Description

The function BS3VL evaluates a trivariate tensor-product spline (represented as a linear combination of tensor-product B-splines) at a given point. This routine is a special case of the IMSL routine BS3DR, which evaluates a partial derivative of such a spline. (The value of a spline is its zero-th derivative.) For more information, see de Boor (1978, pages 351-353).

This routine returns the value of the function $s$ at a point $(x, y, z)$ given the coefficients $c$ by computing

$$
s(x, y, z)=\sum_{l=1}^{N_{z}} \sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}(x) B_{m, k_{y}, \mathbf{t}_{y}}(y) B_{l, k_{z}, \mathbf{t}_{z}}(z)
$$

where $k_{x}, k_{y}$, and $k_{z}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD, KYORD, and KZORD, respectively.) Likewise, $\mathbf{t}_{x}, \mathbf{t}_{y}$, and $\mathbf{t}_{z}$ are the corresponding knot sequences (XKNOT, YKNOT, and ZKNOT).

## Comments

Workspace may be explicitly provided, if desired, by use of B23VL/DB23VL. The reference is:

```
CALL B23VL (X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF,
NYCOEF, NZCOEF, BSCOEF,WK)
```

The additional argument is:

WK - Work array of length 3 * MAX(KXORD, KYORD, KZORD) + KYORD * KZORD + KZORD.

## Example

For an example of the use of BS3VL, see IMSL routine BS3IN.

## BS3DR

This function evaluates the derivative of a three-dimensional tensor-product spline, given its tensor-product B-spline representation.

## Function Return Value

BS3DR - Value of the (IXDER, IYDER, IZDER) derivative of the spline at (X, Y, Z). (Output)

## Required Arguments

IXDER - Order of the X-derivative. (Input)
IYDER - Order of the Y-derivative. (Input)
IZDER - Order of the z-derivative. (Input)
$\boldsymbol{X}$ - x -coordinate of the point at which the spline is to be evaluated. (Input)
$\mathbf{Y}$ - Y-coordinate of the point at which the spline is to be evaluated. (Input)
$\mathbf{Z}$ - z-coordinate of the point at which the spline is to be evaluated. (Input)
KXORD - Order of the spline in the x -direction. (Input)
KYORD - Order of the spline in the Y-direction. (Input)
KZORD - Order of the spline in the z-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X -direction. (Input)
KNOT must be nondecreasing.
YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input)
YKNOT must be nondecreasing.
ZKNOT - Array of length NZCOEF + KZORD containing the knot sequence in the z-direction. (Input)
ZKNOT must be nondecreasing.

NXCOEF - Number of B-spline coefficients in the X -direction. (Input)

NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
NZCOEF - Number of B-spline coefficients in the z-direction. (Input)
BSCOEF - Array of length NXCoef * NYcoef * nZcoef containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF by NZCOEF.

## FORTRAN 90 Interface

Generic: BS3DR (IXDER, IYDER, IZDER, X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Specific: The specific interface names are S_BS3DR and D_BS3DR.

## FORTRAN 77 Interface

Single: BS3DR (IXDER, IYDER, IZDER, X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Double: The double precision function name is DBS3DR.

## Description

The function BS3DR evaluates a partial derivative of a trivariate tensor-product spline (represented as a linear combination of tensor-product B-splines) at a given point. For more information, see de Boor (1978, pages 351-353).
This routine returns the value of the function $s^{(p, q, r)}$ at a point $(x, y, z)$ given the coefficients $c$ by computing

$$
s^{(p, q, r)}(x, y, z)=\sum_{l=1}^{N_{z}} \sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}^{(p)}(x) B_{m, k_{y}, \mathbf{t}_{y}}^{(q)}(y) B_{l, k_{z}, \mathbf{t}_{z}}^{(r)}(z)
$$

where $k_{x}, k_{y}$, and $k_{z}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD, KYORD, and KZORD, respectively.) Likewise, $\mathbf{t}_{x}, \mathbf{t}_{y}$, and $\mathbf{t}_{z}$ are the corresponding knot sequences (XKNOT, YKNOT, and ZKNOT).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B23DR/DB23DR. The reference is:

CALL B23DR (IXDER, IYDER, IZDER, X, Y, Z, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF, WK)
The additional argument is:

```
WK - Work array of length 3* MAXO (KXORD, KYORD, KZORD) + KYORD * KZORD
    + KZORD.
```

2. Informational errors

| Type | Code |  |
| :---: | :---: | :---: |
| 3 | 1 | The point X does not satisfy |
|  |  | XKNOT(KXORD) .LE. X .LE. XKNOT(NXCOEF + 1). |
| 3 | 2 | The point $Y$ does not satisfy |
|  |  | YKNOT(KYORD) .LE. Y .LE. YKNOT(NYCOEF + 1). |
| 3 | 3 | The point z does not satisfy |
|  |  | ZKNOT (KZORD) .LE. Z .LE. ZKNOT(NZCOEF + 1) |

## Example

In this example, a spline interpolant $s$ to a function $f(x, y, z)=x^{4}+y(x z)^{3}$ is constructed using BS3IN. Next, BS3DR is used to compute $s^{(2,0,1)}(x, y, z)$. The values of this partial derivative and the error are computed on a $4 \times 4 \times 2$ grid and then displayed.

```
    USE BS3DR INT
    USE BS3IN 'INT
    USE BSNAK_INT
    USE UMACH_INT
    IMPLICIT NONE
! SPECIFICATIONS FOR PARAMETERS
    INTEGER KXORD, KYORD, KZORD, LDF, MDF, NXDATA, NXKNOT,&
        NYDATA, NYKNOT, NZDATA, NZKNOT
    PARAMETER (KXORD=5, KYORD=2, KZORD=3, NXDATA=21, NYDATA=6,&
            NZDATA=8, LDF=NXDATA, MDF=NYDATA, &
            NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD,&
            NZKNOT=NZDATA+KZORD)
            INTEGER I, J, K, L, NOUT, NXCOEF, NYCOEF, NZCOEF
            REAL BSCOEF (NXDATA,NYDATA,NZDATA), F, F201,&
            FDATA(LDF,MDF,NZDATA), FLOAT, S201, X, XDATA(NXDATA),&
            XKNOT (NXKNOT), Y, YDATA(NYDATA), YKNOT(NYKNOT), Z,&
            ZDATA(NZDATA), ZKNOT (NZKNOT)
            INTRINSIC FLOAT
                                    Define function and (2,0,1)
                                    derivative
                            F(X,Y,Z) = X*X*X*X + X*X*X*Y*Z*Z*Z
                            F201(X,Y,Z) = 18.0*X*Y*Z
                            DO 10 I=1, NXDATA
        XDATA(I) = FLOAT(I-11)/10.0
    CONTINUE
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/FLOAT(NYDATA-1)
    20 CONTINUE
DO 30 I=1, NZDATA
```

```
            ZDATA(I) = FLOAT(I-1)/FLOAT(NZDATA-1)
    3 0 ~ C O N T I N U E ~
! Generate knots
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
    CALL BSNAK (NZDATA, ZDATA, KZORD, ZKNOT)
    50 Generate FDATA
        50 K=1, NZDATA
        DO 40 I=1, NYDATA
            DO 40 J=1, NXDATA
                FDATA(J,I,K) = F(XDATA (J),YDATA (I),ZDATA (K))
    CONTINUE
    CONTINUE
! Get output unit number
    CALL UMACH (2, NOUT)
    CALL BS3IN (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD, KZORD, XKNOT, &
                YKNOT, ZKNOT, BSCOEF)
!
    NXCOEF = NXDATA
    NYCOEF = NYDATA
    NZCOEF = NZDATA
! Write heading
    WRITE (NOUT,99999)
! Print over a grid of
    [-1.0,1.0] x [0.0,1.0] x [0.0,1.0]
                            at 32 points.
    DO 80 I=1, 4
        DO 70 J=1, 4
            DO 60 L=1, 2
                X = 2.0*(FLOAT (I-1)/3.0) - 1.0
                Y = FLOAT (J-1)/3.0
                Z = FLOAT (L-1)
            Evaluate spline
                S201 = BS3DR (2,0,1,X,Y,Z,KXORD,KYORD,KZORD,XKNOT,YKNOT, &
                    ZKNOT,NXCOEF,NYCOEF,NZCOEF,BSCOEF)
                WRITE (NOUT,'(3F12.4,2F12.6)') X, Y, Z, S201,&
                    F201(X,Y,Z) - S201
        CONTINUE
    CONTINUE
    CONTINUE
9999 FORMAT (38X, '(2,0,1)', /, 9X, 'X', 11X, &
                        'Y', 11X, 'Z', 4X, 'S (X,Y,Z) Error')
    END
```


## Output

|  | $(2,0,1)$ |  |  |  |
| :---: | :---: | :---: | ---: | ---: |
| X | Y | Z | S |  |
| -1.0000 | 0.0000 | 0.0000 | -0.000107 | 0.000107 |
| -1.0000 | 0.0000 | 1.0000 | 0.000053 | -0.000053 |
| -1.0000 | 0.3333 | 0.0000 | 0.064051 | -0.064051 |
| -1.0000 | 0.3333 | 1.0000 | -5.935941 | -0.064059 |
| -1.0000 | 0.6667 | 0.0000 | 0.127542 | -0.127542 |
| -1.0000 | 0.6667 | 1.0000 | -11.873034 | -0.126966 |


| -1.0000 | 1.0000 | 0.0000 | 0.191166 | -0.191166 |
| ---: | ---: | ---: | ---: | ---: |
| -1.0000 | 1.0000 | 1.0000 | -17.808527 | -0.191473 |
| -0.3333 | 0.0000 | 0.0000 | -0.000002 | 0.000002 |
| -0.3333 | 0.0000 | 1.0000 | 0.000000 | 0.000000 |
| -0.3333 | 0.3333 | 0.0000 | 0.021228 | -0.021228 |
| -0.3333 | 0.3333 | 1.0000 | -1.978768 | -0.021232 |
| -0.3333 | 0.6667 | 0.0000 | 0.042464 | -0.042464 |
| -0.3333 | 0.6667 | 1.0000 | -3.957536 | -0.042464 |
| -0.3333 | 1.0000 | 0.0000 | 0.063700 | -0.063700 |
| -0.3333 | 1.0000 | 1.0000 | -5.936305 | -0.063694 |
| 0.3333 | 0.0000 | 0.0000 | -0.000003 | 0.000003 |
| 0.3333 | 0.0000 | 1.0000 | 0.000000 | 0.000000 |
| 0.3333 | 0.3333 | 0.0000 | -0.021229 | 0.021229 |
| 0.3333 | 0.3333 | 1.0000 | 1.978763 | 0.021238 |
| 0.3333 | 0.6667 | 0.0000 | -0.042465 | 0.042465 |
| 0.3333 | 0.6667 | 1.0000 | 3.957539 | 0.042462 |
| 0.3333 | 1.0000 | 0.0000 | -0.063700 | 0.063700 |
| 0.3333 | 1.0000 | 1.0000 | 5.936304 | 0.063697 |
| 1.0000 | 0.0000 | 0.0000 | -0.000098 | 0.000098 |
| 1.0000 | 0.0000 | 1.0000 | 0.000053 | -0.000053 |
| 1.0000 | 0.3333 | 0.0000 | -0.063855 | 0.063855 |
| 1.0000 | 0.3333 | 1.0000 | 5.936146 | 0.063854 |
| 1.0000 | 0.6667 | 0.0000 | -0.127631 | 0.127631 |
| 1.0000 | 0.6667 | 1.0000 | 11.873067 | 0.126933 |
| 1.0000 | 1.0000 | 0.0000 | -0.191442 | 0.191442 |
| 1.0000 | 1.0000 | 1.0000 | 17.807940 | 0.192060 |

## BS3GD

Evaluates the derivative of a three-dimensional tensor-product spline, given its tensor-product B-spline representation on a grid.

## Required Arguments

IXDER — Order of the X -derivative. (Input)
IYDER - Order of the Y-derivative. (Input)
IZDER — Order of the z-derivative. (Input)

XVEC - Array of length NX containing the $x$-coordinates at which the spline is to be evaluated. (Input)
The points in XVEC should be strictly increasing.
YVEC - Array of length NY containing the $y$-coordinates at which the spline is to be evaluated. (Input)
The points in YVEC should be strictly increasing.
ZVEC - Array of length NY containing the $y$-coordinates at which the spline is to be evaluated. (Input)
The points in YVEC should be strictly increasing.

KXORD - Order of the spline in the $x$-direction. (Input)
KYORD - Order of the spline in the $y$-direction. (Input)
KZORD - Order of the spline in the $z$-direction. (Input)
XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the $x$-direction. (Input)
XKNOT must be nondecreasing.
YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the $y$-direction. (Input)
YKNOT must be nondecreasing.
ZKNOT - Array of length NZCOEF + KZORD containing the knot sequence in the $z$-direction. (Input)
ZKNOT must be nondecreasing.
BSCOEF - Array of length NXCOEF * NYCOEF * NZCOEF containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by nYCOEF by nZCOEF.
VALUE - Array of size NX by NY by NZ containing the values of the (IXDER, IYDER, IZDER) derivative of the spline on the NX by NY by NZ grid. (Output) $\operatorname{VALUE}(\mathrm{I}, \mathrm{J}, \mathrm{K})$ contains the derivative of the spline at the point (XVEC(I), $\operatorname{YVEC}(\mathrm{J})$, ZVEC(K)).

## Optional Arguments

$N X$ - Number of grid points in the $x$-direction. (Input)
Default: NX $=$ size $(\mathrm{XVEC}, 1)$.
$N Y$ - Number of grid points in the $y$-direction. (Input)
Default: NY = size (YVEC,1).
$N Z$ - Number of grid points in the $z$-direction. (Input)
Default: NZ $=$ size (ZVEC,1).
NXCOEF - Number of B-spline coefficients in the $x$-direction. (Input)
Default: $\mathrm{NXCOEF}=\operatorname{size}($ XKNOT, 1$)-$ KXORD.
NYCOEF - Number of B-spline coefficients in the $y$-direction. (Input)
Default: NYCOEF $=$ size $($ YKNOT, 1$)-$ KYORD.
NZCOEF - Number of B-spline coefficients in the $z$-direction. (Input)
Default: NZCOEF $=\operatorname{size}($ ZKNOT, 1$)-$ KZORD.

LDVALU - Leading dimension of VALUE exactly as specified in the dimension statement of the calling program. (Input)
Default: LDVALU $=$ SIZE (VALUE, 1 ).

MDVALU - Middle dimension of VALUE exactly as specified in the dimension statement of the calling program. (Input)
Default: MDVALU $=$ SIZE (VALUE,2).

## FORTRAN 90 Interface

Generic: CALL BS3GD (IXDER, IYDER, IZDER, XVEC, YVEC, ZVEC, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF, VALUE [, ...])

Specific: The specific interface names are S_BS3GD and D_BS3GD.

## FORTRAN 77 Interface

Single: CALL BS3GD (IXDER, IYDER, IZDER, NX, XVEC, NY, YVEC, NZ, ZVEC, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF, VALUE, LDVALU, MDVALU)

Double: The double precision name is DBS3GD.

## Description

The routine BS3GD evaluates a partial derivative of a trivariate tensor-product spline (represented as a linear combination of tensor-product B-splines) on a grid. For more information, see de Boor (1978, pages 351-353).

This routine returns the value of the function $s^{(p, q, r)}$ on the grid $\left(x_{i}, y_{j}, z_{k}\right)$ for $i=1, \ldots, n x$, $j=1, \ldots, n y$, and $k=1, \ldots, n z$ given the coefficients $c$ by computing (for all $(x, y, z)$ on the grid)

$$
s^{(p, q, r)}(x, y, z)=\sum_{l=1}^{N_{z}} \sum_{m=1}^{N_{y}} \sum_{n=1}^{N_{x}} c_{n m l} B_{n, k_{x}, \mathbf{t}_{x}}^{(p)}(x) B_{m, k_{y}, \mathbf{t}_{y}}^{(q)}(y) B_{l, k_{z}, \mathbf{t}_{z}}^{(r)}(z)
$$

where $k_{x}, k_{y}$, and $k_{z}$ are the orders of the splines. (These numbers are passed to the subroutine in KXORD, KYORD, and KZORD, respectively.) Likewise, $\mathbf{t}_{x}, \mathbf{t}_{y}$, and $\mathbf{t}_{z}$ are the corresponding knot sequences (XKNOT, YKNOT, and ZKNOT). The grid must be ordered in the sense that $x_{i}<x_{i+1}, y_{j}<y_{j+1}$, and $z_{k}<z_{k+1}$.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B23GD/DB23GD. The reference is:

$$
\begin{aligned}
& \text { CALL B23GD ( (IXDER, IYDER, IZDER, NX, XVEC, NY, YVEC, NZ, ZVEC, KXORD, } \\
& \text { KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF, VALUE, }
\end{aligned}
$$

```
        LDVALU, MDVALU, LEFTX, LEFTY, LEFTZ, A, B, C, DBIATX, DBIATY, DBIATZ, BX, BY,
        BZ)
The additional arguments are as follows:
LEFTX - Work array of length NX.
LEFTY - Work array of length NY.
LEFTZ - Work array of length NZ.
A-Work array of length KXORD * KXORD.
B — Work array of length KYORD * KYORD.
C - Work array of length KZORD * KzORD.
DBIATX - Work array of length KXORD * (IXDER + 1).
DBIATY - Work array of length KYORD * (IYDER + 1).
DBIATZ - Work array of length KZORD * (IZDER + 1).
\(\boldsymbol{B} \boldsymbol{X}\) - Work array of length KXORD * NX.
\(\boldsymbol{B Y}\) - Work array of length KYORD * NY.
BZ - Work array of length KZORD * NZ.
2. Informational errors
\begin{tabular}{|c|c|c|}
\hline Type & Code & \\
\hline 3 & 1 & XVEC(I) does not satisfy \\
\hline 3 & 2 & XKNOT \((\) KXORD \() \leq \operatorname{XVEC}(\mathrm{I}) \leq \operatorname{XKNOT}(N X C O E F+1)\). YVEC(I) does not satisfy \\
\hline 3 & 3 & YKNOT \((\) KYORD \() \leq \operatorname{YVEC}(\mathrm{I}) \leq\) YKNOT \((N Y C O E F+1)\). ZVEC(I) does not satisfy \\
\hline 4 & 4 & ZKNOT \((\) KZORD \() \leq \operatorname{ZVEC}(\mathrm{I}) \leq\) ZKNOT \((N Z C O E F+1)\). XVEC is not strictly increasing. \\
\hline 4 & 5 & YVEC is not strictly increasing. \\
\hline 4 & 6 & ZVEC is not strictly increasing. \\
\hline
\end{tabular}
```


## Example

In this example, a spline interpolant $s$ to a function $f(x, y, z)=x^{4}+y(x z)^{3}$ is constructed using BS3IN. Next, BS3GD is used to compute $s^{(2,0,1)}(x, y, z)$ on the grid. The values of this partial derivative and the error are computed on a $4 \times 4 \times 2$ grid and then displayed.

USE BS3GD_INT

```
    USE BS3IN_INT
    USE BSNAK INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER KXORD, KYORD, KZORD, LDF, LDVAL, MDF, MDVAL, NXDATA,&
        NXKNOT, NYDATA, NYKNOT, NZ, NZDATA, NZKNOT
    PARAMETER (KXORD=5, KYORD=2, KZORD=3, LDVAL=4, MDVAL=4,&
            NXDATA=21, NYDATA=6, NZ=2, NZDATA=8, LDF=NXDATA,&
            MDF=NYDATA, NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD,&
            NZKNOT=NZDATA+KZORD)
!
    INTEGER I, J, K, L, NOUT, NXCOEF, NYCOEF, NZCOEF
    REAL BSCOEF(NXDATA,NYDATA,NZDATA), F, F201,&
        FDATA(LDF,MDF,NZDATA), FLOAT, VALUE(LDVAL,MDVAL,NZ),&
        X, XDATA(NXDATA), XKNOT(NXKNOT), XVEC(LDVAL), Y,&
        YDATA(NYDATA), YKNOT (NYKNOT), YVEC(MDVAL), Z,&
        ZDATA(NZDATA), ZKNOT(NZKNOT), ZVEC(NZ)
    INTRINSIC FLOAT
!
!
!
    F(X,Y,Z) = X*X*X*X + X*X*X*Y*Z*Z*Z
    F201(X,Y,Z) = 18.0*X*Y*Z
!
! Set up X interpolation points
    DO 10 I=1, NXDATA
        XDATA(I) = 2.0*(FLOAT(I-1)/FLOAT (NXDATA-1)) - 1.0
    10 CONTINUE
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/FLOAT(NYDATA-1)
    20 CONTINUE
    DO 30 I=1, NZDATA
        ZDATA(I) = FLOAT(I-1)/FLOAT(NZDATA-1)
    30 CONTINUE
!
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
    CALL BSNAK (NZDATA, ZDATA, KZORD, ZKNOT)
                            Generate FDATA
    DO 50 K=1, NZDATA
        DO 40 I=1, NYDATA
            DO 40 J=1, NXDATA
                FDATA(J,I,K) = F(XDATA(J),YDATA(I),ZDATA (K))
    4 0 ~ C O N T I N U E ~
    5 0 ~ C O N T I N U E ~
    CALL BS3IN (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD,&
        KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF)
    NXCOEF = NXDATA
    NYCOEF = NYDATA
```

```
        NZCOEF = NZDATA
    Print over a grid of
                                    [-1.0,1.0] x [0.0,1.0] x [0.0,1.0]
                                    at 32 points.
            DO 60 I=1, 4
            XVEC(I) = 2.0*(FLOAT (I-1)/3.0) - 1.0
    6 0 ~ C O N T I N U E
            DO 70 J=1, 4
            YVEC(J) = FLOAT(J-1)/3.0
    7 0 ~ C O N T I N U E
            DO 80 L=1, 2
            ZVEC(L) = FLOAT(L-1)
    8 0 ~ C O N T I N U E
            CALL BS3GD (2, 0, 1, XVEC, YVEC, ZVEC, KXORD, KYORD, &
                KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF, VALUE)
            WRITE (NOUT,99999)
            DO 110 I=1, 4
                DO 100 J=1, 4
                    DO 90 L=1, 2
                    WRITE (NOUT,'(5F13.4)') XVEC(I), YVEC(J), ZVEC(L), &
                                    VALUE (I, J,L), &
                                    F201 (XVEC(I),YVEC (J) , ZVEC(L)) -&
                                    VALUE (I, J, L)
    90 CONTINUE
    100 CONTINUE
    110 CONTINUE
99999 FORMAT (44X, '(2,0,1)', /, 10X, 'X', 11X, 'Y', 10X, 'Z', 10X,&
            'S (X,Y,Z) Error')
            STOP
            END
```

                Output
    |  |  | $(2,0,1)$ |  |  |
| :---: | :---: | :---: | :---: | :---: |
| X | Y | Z | S (X, Y, Z) | Error |
| -1.0000 | 0.0000 | 0.0000 | -0.0005 | 0.0005 |
| -1.0000 | 0.0000 | 1.0000 | 0.0002 | -0.0002 |
| -1.0000 | 0.3333 | 0.0000 | 0.0641 | -0.0641 |
| -1.0000 | 0.3333 | 1.0000 | -5.9360 | -0.0640 |
| -1.0000 | 0.6667 | 0.0000 | 0.1274 | -0.1274 |
| -1.0000 | 0.6667 | 1.0000 | -11.8730 | -0.1270 |
| -1.0000 | 1.0000 | 0.0000 | 0.1911 | -0.1911 |
| -1.0000 | 1.0000 | 1.0000 | -17.8086 | -0.1914 |
| -0.3333 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| -0.3333 | 0.0000 | 1.0000 | 0.0000 | 0.0000 |
| -0.3333 | 0.3333 | 0.0000 | 0.0212 | -0.0212 |
| -0.3333 | 0.3333 | 1.0000 | -1.9788 | -0.0212 |
| -0.3333 | 0.6667 | 0.0000 | 0.0425 | -0.0425 |
| -0.3333 | 0.6667 | 1.0000 | -3.9575 | -0.0425 |
| -0.3333 | 1.0000 | 0.0000 | 0.0637 | -0.0637 |
| -0.3333 | 1.0000 | 1.0000 | -5.9363 | -0.0637 |
| 0.3333 | 0.0000 | 0.0000 | 0.0000 | 0.0000 |
| 0.3333 | 0.0000 | 1.0000 | 0.0000 | 0.0000 |


| 0.3333 | 0.3333 | 0.0000 | -0.0212 | 0.0212 |
| :--- | ---: | ---: | ---: | ---: |
| 0.3333 | 0.3333 | 1.0000 | 1.9788 | 0.0212 |
| 0.3333 | 0.6667 | 0.0000 | -0.0425 | 0.0425 |
| 0.3333 | 0.6667 | 1.0000 | 3.9575 | 0.0425 |
| 0.3333 | 1.0000 | 0.0000 | -0.0637 | 0.0637 |
| 0.3333 | 1.0000 | 1.0000 | 5.9363 | 0.0637 |
| 1.0000 | 0.0000 | 0.0000 | -0.0005 | 0.0005 |
| 1.0000 | 0.0000 | 1.0000 | 0.0000 | 0.0000 |
| 1.0000 | 0.3333 | 0.0000 | -0.0637 | 0.0637 |
| 1.0000 | 0.3333 | 1.0000 | 5.9359 | 0.0641 |
| 1.0000 | 0.6667 | 0.0000 | -0.1273 | 0.1273 |
| 1.0000 | 0.6667 | 1.0000 | 11.8733 | 0.1267 |
| 1.0000 | 1.0000 | 0.0000 | -0.1912 | 0.1912 |
| 1.0000 | 1.0000 | 1.0000 | 17.8096 | 0.1904 |

## BS3IG

This function evaluates the integral of a tensor-product spline in three dimensions over a threedimensional rectangle, given its tensor-product B-spline representation.

## Function Return Value

BS3IG - Integral of the spline over the three-dimensional rectangle (A, B) by (C, D) by (E, F). (Output)

## Required Arguments

A - Lower limit of the X -variable. (Input)
B - Upper limit of the x -variable. (Input)
$\boldsymbol{C}$ - Lower limit of the Y-variable. (Input)
D- Upper limit of the Y-variable. (Input)
$\boldsymbol{E}$ - Lower limit of the z-variable. (Input)
$\boldsymbol{F}$ - Upper limit of the z-variable. (Input)
KXORD - Order of the spline in the X-direction. (Input)
KYORD - Order of the spline in the Y-direction. (Input)
KZORD - Order of the spline in the z-direction. (Input)

XKNOT - Array of length NXCOEF + KXORD containing the knot sequence in the X-direction. (Input)
XKNOT must be nondecreasing.

YKNOT - Array of length NYCOEF + KYORD containing the knot sequence in the Y-direction. (Input) YKNOT must be nondecreasing.

ZKNOT - Array of length NZCOEF + KZORD containing the knot sequence in the z-direction. (Input)
ZKNOT must be nondecreasing.
NXCOEF - Number of B-spline coefficients in the X -direction. (Input)
NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
NZCOEF - Number of B-spline coefficients in the z-direction. (Input)
BSCOEF - Array of length NXCOEF * NYCOEF * NZCOEF containing the tensor-product B-spline coefficients. (Input)
BSCOEF is treated internally as a matrix of size NXCOEF by NYCOEF by NZCOEF.

## FORTRAN 90 Interface

Generic: BS3IG (A, B, C, D, E, F, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Specific: The specific interface names are S_BS3IG and D_BS3IG.

## FORTRAN 77 Interface

Single: BS3IG (A, B, C, D, E, F, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, BSCOEF)

Double: The double precision function name is DBS3IG.

## Description

The routine BS3IG computes the integral of a tensor-product three-dimensional spline, given its B-spline representation. Specifically, given the knot sequence $\mathbf{t}_{x}=$ XKNOT, $\mathbf{t}_{y}=$ YKNOT, $\mathbf{t}_{z}=Z \mathrm{ZKNOT}$, the order $k_{x}=$ KXORD, $k_{y}=$ KYORD, $k_{z}=$ KZORD, the coefficients $\beta=$ BSCOEF, the number of coefficients $n_{x}=$ NXCOEF, $n_{y}=$ NYCOEF, $n_{z}=$ NZCOEF, and a three-dimensional rectangle $[a, b]$ by $[c, d]$ by $[e, f]$, BS3IG returns the value

$$
\int_{a}^{b} \int_{c}^{d} \int_{e}^{f} \sum_{i=1}^{n_{x}} \sum_{j=1}^{n_{y}} \sum_{m=1}^{n_{z}} \beta_{i j m} B_{i j m} d z d y d x
$$

where

$$
B_{i j m}(x, y, z)=B_{i, k_{x}, \mathbf{t}_{x}}(x) B_{j, k_{y}, \mathbf{t}_{y}}(y) B_{m, k_{z} \mathbf{t}_{z}}(z)
$$

This routine uses the identity (22) on page 151 of de Boor (1978). It assumes (for all knot sequences) that the first and last $k$ knots are stacked, that is, $\mathbf{t}_{1}=\ldots=\mathbf{t}_{k}$ and $\mathbf{t}_{n+1}=\ldots=\mathbf{t}_{n+k}$, where $k$ is the order of the spline in the $x, y$, or $z$ direction.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B23IG/DB23IG. The reference is:
```
CALL B23IG(A, B, C, D, E, F, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF,
NYCOEF, NZCOEF, BSCOEF, WK)
```

The additional argument is:

```
WK - Work array of length 4 * (MAX (KXORD, KYORD, KZORD) + 1) + NYCOEF
    + NZCOEF.
```

2. Informational errors

Type Code
31 The lower limit of the X -integration is less than XKNOT(KXORD).
32 The upper limit of the x -integration is greater than XKNOT(NXCOEF + 1).
33 The lower limit of the Y-integration is less than YKNOT(KYORD).
34 The upper limit of the $Y$-integration is greater than YKNOT(NYCOEF + 1).
35 The lower limit of the z - integration is less than ZKNOT(KZORD).
36 The upper limit of the $z$-integration is greater than ZKNOT(NZCOEF + 1).
413 Multiplicity of the knots cannot exceed the order of the spline.
$4 \quad 14$ The knots must be nondecreasing.

## Example

We integrate the three-dimensional tensor-product quartic $\left(k_{x}=5\right)$ by linear $\left(k_{y}=2\right)$ by quadratic $\left(k_{z}=3\right)$ spline which interpolates $x^{3}+x y z$ at the points

$$
\{(i / 10, j / 5, m / 7): i=-10, \ldots, 10, j=0, \ldots, 5, \text { and } m=0, \ldots, 7\}
$$

over the rectangle $[0,1] \times[.5,1] \times[0, .5]$. The exact answer is $11 / 128$.

```
USE BS3IG INT
USE BS3IN -
USE BSNAK_INT
USE UMACH_
IMPLICIT NONE
INTEGER KXORD, KYORD, KZORD, LDF, MDF, NXDATA, NXKNOT, &
    NYDATA, NYKNOT, NZDATA, NZKNOT
PARAMETER (KXORD=5, KYORD=2, KZORD=3, NXDATA=21, NYDATA=6,&
```

```
                NZDATA=8, LDF=NXDATA, MDF=NYDATA, &
                NXKNOT=NXDATA+KXORD, NYKNOT=NYDATA+KYORD,&
                NZKNOT=NZDATA+KZORD)
!
    INTEGER I, J, K, NOUT, NXCOEF, NYCOEF, NZCOEF
    REAL A, B, BSCOEF (NXDATA,NYDATA,NZDATA), C , D, E,&
        F, FDATA(LDF,MDF,NZDATA), FF, FIG, FLOAT, G, H, RI,&
            RJ, VAL, X, XDATA(NXDATA), XKNOT (NXKNOT), Y, &
            YDATA (NYDATA), YKNOT (NYKNOT), Z, ZDATA(NZDATA),&
            ZKNOT (NZKNOT)
    INTRINSIC FLOAT
    F(X,Y,Z) = X*X*X + X*Y*Z
                Define function
                            Set up interpolation points
    DO 10 I=1, NXDATA
        XDATA(I) = FLOAT (I-11)/10.0
    CONTINUE
    CALL BSNAK (NXDATA, XDATA, KXORD, XKNOT)
                            Set up interpolation points
    DO 20 I=1, NYDATA
        YDATA(I) = FLOAT(I-1)/FLOAT(NYDATA-1)
    CONTINUE
                                    Generate knot sequence
    CALL BSNAK (NYDATA, YDATA, KYORD, YKNOT)
                            Set up interpolation points
    DO 30 I=1, NZDATA
        ZDATA(I) = FLOAT(I-1)/FLOAT(NZDATA-1)
    CONTINUE
                            Generate knot sequence
    CALL BSNAK (NZDATA, ZDATA, KZORD, ZKNOT)
                            Generate FDATA
    DO 50 K=1, NZDATA
        DO 40 I=1, NYDATA
            DO 40 J=1, NXDATA
                FDATA(J,I,K) = F(XDATA (J),YDATA (I), ZDATA (K))
    CONTINUE
    CONTINUE
    Get output unit number
    CALL UMACH (2, NOUT) Interpolate
    CALL BS3IN (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD, KZORD, XKNOT, &
        YKNOT, ZKNOT, BSCOEF)
    NXCOEF = NXDATA
    NYCOEF = NYDATA
    NZCOEF = NZDATA
    A = 0.0
    B}=1.
    C = 0.5
    D = 1.0
    E = 0.0
    FF}=0.
    Integrate
    VAL = BS3IG(A, B, C ,D,E,FF,KXORD,KYORD,KZORD, XKNOT, YKNOT, ZKNOT, &
```

```
!
                                    Calculate integral directly
    G = .5*(B**4-A**4)
    H}=(B-A)*(B+A
    RI= G* (D-C )
    RJ =.5*H* (D-C )* (D+C )
    FIG =. 5* (RI* (FF-E) +.5*RJ* (FF-E)* (FF+E))
! Print results
    WRITE (NOUT,99999) VAL, FIG, FIG - VAL
99999 FORMAT (' Computed Integral = ', F10.5, /, ' Exact Integral '&
    ,'= ', F10.5,/,' Error '&
    ', '= ', F10.6, /)
END
```


## Output

```
Computed Integral = 0.08594
Exact Integral = 0.08594
Error = 0.000000
```


## BSCPP

Converts a spline in B-spline representation to piecewise polynomial representation.

## Required Arguments

KORDER - Order of the spline. (Input)
XKNOT - Array of length KORDER + NCOEF containing the knot sequence. (Input) XKNOT must be nondecreasing.

NCOEF - Number of B-spline coefficients. (Input)
BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Input)
NPPCF - Number of piecewise polynomial pieces. (Output) NPPCF is always less than or equal to NCOEF - KORDER +1.

BREAK - Array of length (NPPCF + 1) containing the breakpoints of the piecewise polynomial representation. (Output) BREAK must be dimensioned at least NCOEF - KORDER +2 .

PPCOEF - Array of length KORDER * NPPCF containing the local coefficients of the polynomial pieces. (Output) PPCOEF is treated internally as a matrix of size KORDER by NPPCF.

## FORTRAN 90 Interface

Generic: CALL BSCPP (KORDER, XKNOT, NCOEF, BSCOEF, NPPCF, BREAK, PPCOEF)

Specific: The specific interface names are S_BSCPP and D_BSCPP.

## FORTRAN 77 Interface

Single: CALL BSCPP (KORDER, XKNOT, NCOEF, BSCOEF, NPPCF, BREAK, PPCOEF)

Double: The double precision name is DBSCPP.

## Description

The routine BSCPP is based on the routine BSPLPP by de Boor (1978, page 140). This routine is used to convert a spline in B-spline representation to a piecewise polynomial (pp) representation which can then be evaluated more efficiently. There is some overhead in converting from the B-spline representation to the pp representation, but the conversion to pp form is recommended when 3 or more function values are needed per polynomial piece.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2CPP/DB2CPP. The reference is:

CALL B2CPP (KORDER, XKNOT, NCOEF, BSCOEFF, NPPCF, BREAK, PPCOEF, WK)
The additional argument is
$\boldsymbol{W} \boldsymbol{K}$ - Work array of length (KORDER + 3) * KORDER.
2. Informational errors

Type Code
44 Multiplicity of the knots cannot exceed the order of the spline. $4 \quad 5$ The knots must be nondecreasing.

## Example

For an example of the use of BSCPP, see PPDER.

## PPVAL

This function evaluates a piecewise polynomial.

## Function Return Value

PPVAL — Value of the piecewise polynomial at X . (Output)

## Required Arguments

$\boldsymbol{X}$ - Point at which the polynomial is to be evaluated. (Input)

BREAK - Array of length NINTV + 1 containing the breakpoints of the piecewise polynomial representation. (Input) BREAK must be strictly increasing.

PPCOEF - Array of size KORDER * NINTV containing the local coefficients of the piecewise polynomial pieces. (Input)
PPCOEF is treated internally as a matrix of size KORDER by NINTV.

## Optional Arguments

KORDER - Order of the polynomial. (Input)
Default: KORDER = size (PPCOEF,1).
NINTV — Number of polynomial pieces. (Input)
Default: NINTV = size (PPCOEF,2).

## FORTRAN 90 Interface

Generic: PPVAL (X, BREAK, PPCOEF [, ...])

Specific: The specific interface names are __PPVAL $^{\text {and }} \mathrm{D}_{-}$PPVAL.

## FORTRAN 77 Interface

Single: PPVAL (X, KORDER, NINTV, BREAK, PPCOEF)
Double: The double precision function name is DPPVAL.

## Description

The routine PPVAL evaluates a piecewise polynomial at a given point. This routine is a special case of the routine PPDER, which evaluates the derivative of a piecewise polynomial. (The value of a piecewise polynomial is its zero-th derivative.)

The routine PPDER is based on the routine PPVALU in de Boor (1978, page 89).

## Example

In this example, a spline interpolant to a function $f$ is computed using the IMSL routine BSINT. This routine represents the interpolant as a linear combination of B-splines. This representation is then converted to piecewise polynomial representation by calling the IMSL routine BSCPP. The piecewise polynomial is evaluated using PPVAL. These values are compared to the corresponding values of $f$.

```
USE PPVAL INT
USE BSNAK }\mp@subsup{}{}{-}\mathrm{ INT
USE BSCPP INT
USE BSINT INT
USE UMACH__INT
```

```
    IMPLICIT NONE
    INTEGER KORDER, NCOEF, NDATA, NKNOT
    PARAMETER (KORDER=4, NCOEF=20, NDATA=20, NKNOT=NDATA+KORDER)
!
    INTEGER I, NOUT, NPPCF
    REAL BREAK(NCOEF), BSCOEF (NCOEF), EXP, F, FDATA (NDATA),&
        FLOAT, PPCOEF(KORDER,NCOEF), S, X, XDATA(NDATA), &
        XKNOT (NKNOT)
    INTRINSIC EXP, FLOAT
!
    F(X) = X* EXP(X)
    DO 30 I=1, NDATA
        XDATA(I) = FLOAT (I-1)/FLOAT (NDATA-1)
        FDATA(I) = F(XDATA(I))
    3 0 ~ C O N T I N U E
! Generate knot sequence
    CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
                            Compute the B-spline interpolant
    CALL BSINT (NCOEF, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    Convert to piecewise polynomial
    CALL BSCPP (KORDER, XKNOT, NCOEF, BSCOEF, NPPCF, BREAK, PPCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
    DO 40 I=1, NDATA
        X = FLOAT (I-1)/FLOAT (NDATA-1)
            Compute value of the piecewise
            polynomial
        S = PPVAL (X,BREAK, PPCOEF)
        WRITE (NOUT,'(2F12.3, E14.3)') X, S, F(X) - S
    4 0 ~ C O N T I N U E
99999 FORMAT (11X, 'X', 8X, 'S(X)', 7X, 'Error')
    END
```


## Output

| $X$ | $S(X)$ | Error |
| ---: | ---: | ---: |
| 0.000 | 0.000 | $0.000 \mathrm{E}+00$ |
| 0.053 | 0.055 | $-0.745 \mathrm{E}-08$ |
| 0.105 | 0.117 | $0.000 \mathrm{E}+00$ |
| 0.158 | 0.185 | $0.000 \mathrm{E}+00$ |
| 0.211 | 0.260 | $-0.298 \mathrm{E}-07$ |
| 0.263 | 0.342 | $0.298 \mathrm{E}-07$ |
| 0.316 | 0.433 | $0.000 \mathrm{E}+00$ |
| 0.368 | 0.533 | $0.000 \mathrm{E}+00$ |
| 0.421 | 0.642 | $0.000 \mathrm{E}+00$ |
| 0.474 | 0.761 | $0.596 \mathrm{E}-07$ |
| 0.526 | 0.891 | $0.000 \mathrm{E}+00$ |


| 0.579 | 1.033 | $0.000 \mathrm{E}+00$ |
| :--- | ---: | ---: |
| 0.632 | 1.188 | $0.000 \mathrm{E}+00$ |
| 0.684 | 1.356 | $0.000 \mathrm{E}+00$ |
| 0.737 | 1.540 | $-0.119 \mathrm{E}-06$ |
| 0.789 | 1.739 | $0.000 \mathrm{E}+00$ |
| 0.842 | 1.955 | $0.000 \mathrm{E}+00$ |
| 0.895 | 2.189 | $0.238 \mathrm{E}-06$ |
| 0.947 | 2.443 | $0.238 \mathrm{E}-06$ |
| 1.000 | 2.718 | $0.238 \mathrm{E}-06$ |

## PPDER

This function evaluates the derivative of a piecewise polynomial.

## Function Return Value

PPDER - Value of the IDERIV-th derivative of the piecewise polynomial at X . (Output)

## Required Arguments

$\boldsymbol{X}$ - Point at which the polynomial is to be evaluated. (Input)
BREAK - Array of length NINTV +1 containing the breakpoints of the piecewise polynomial representation. (Input) BREAK must be strictly increasing.

PPCOEF - Array of size KORDER * NINTV containing the local coefficients of the piecewise polynomial pieces. (Input) PPCOEF is treated internally as a matrix of size KORDER by NINTV.

## Optional Arguments

IDERIV - Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the value of the polynomial.
Default: $\operatorname{IDERIV}=1$.

KORDER - Order of the polynomial. (Input)
Default: KORDER $=\operatorname{size}(\operatorname{PPCOEF}, 1)$.

NINTV - Number of polynomial pieces. (Input)
Default: NINTV = size (PPCOEF,2).

## FORTRAN 90 Interface

Generic: PPDER (X, BREAK, PPCOEF [, ...])
Specific: The specific interface names are S_PPDER and D_PPDER.

## FORTRAN 77 Interface

Single: PPDER (IDERIV, X, KORDER, NINTV, BREAK, PPCOEF)
Double: The double precision function name is DPPDER.

## Description

The routine PPDER evaluates the derivative of a piecewise polynomial function $f$ at a given point. This routine is based on the subroutine PPVALU by de Boor (1978, page 89). In particular, if the breakpoint sequence is stored in $\xi$ (a vector of length $N=N I N T V+1$ ), and if the coefficients of the piecewise polynomial representation are stored in $\mathbf{c}$, then the value of the $j$-th derivative of $f$ at $x$ $\operatorname{in}\left[\xi_{i}, \xi_{i+1}\right)$ is

$$
f^{(j)}(x)=\sum_{m=j}^{k-1} c_{m+1, i} \frac{\left(x-\xi_{i}\right)^{m-j}}{(m-j)!}
$$

when $j=0$ to $k-1$ and zero otherwise. Notice that this representation forces the function to be right continuous. If $x$ is less than $\xi_{1}$, then $i$ is set to 1 in the above formula; if $x$ is greater than or equal to $\xi_{N}$, then $i$ is set to $N-1$. This has the effect of extending the piecewise polynomial representation to the real axis by extrapolation of the first and last pieces.

## Example

In this example, a spline interpolant to a function $f$ is computed using the IMSL routine BSINT. This routine represents the interpolant as a linear combination of B-splines. This representation is then converted to piecewise polynomial representation by calling the IMSL routine BSCPP. The piecewise polynomial's zero-th and first derivative are evaluated using PPDER. These values are compared to the corresponding values of $f$.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER KORDER, NCOEF, NDATA, NKNOT
PARAMETER (KORDER=4, NCOEF=20, NDATA=20, NKNOT=NDATA+KORDER)
!
INTEGER I, NOUT, NPPCF
REAL BREAK (NCOEF), BSCOEF (NCOEF), DF, DS, EXP, F, \& FDATA (NDATA), FLOAT, PPCOEF (KORDER,NCOEF), $S, \&$ X, XDATA(NDATA), XKNOT (NKNOT)
INTRINSIC EXP, FLOAT
!
$\mathrm{F}(\mathrm{X})=\mathrm{X} * \operatorname{EXP}(\mathrm{X})$
$\mathrm{DF}(\mathrm{X})=(\mathrm{X}+1). * \operatorname{EXP}(\mathrm{X})$
!
DO $10 \mathrm{I}=1$, NDATA XDATA $(I)=\operatorname{FLOAT}(I-1) /$ FLOAT (NDATA-1) FDATA (I) $=\mathrm{F}($ XDATA (I))
10 CONTINUE
!
CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)

```
! Compute the B-spline interpolant
    CALL BSINT (NCOEF, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    Convert to piecewise polynomial
    CALL BSCPP (KORDER, XKNOT, NCOEF, BSCOEF, NPPCF, BREAK, PPCOEF)
        Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
! Print the interpolant on a uniform
    DO 20 I=1, NDATA
        X = FLOAT (I-1) /FLOAT (NDATA-1)
            Compute value of the piecewise
            polynomial
        S = PPDER(X,BREAK, PPCOEF, IDERIV=0, NINTV=NPPCF)
                        Compute derivative of the piecewise
                    polynomial
        DS = PPDER(X,BREAK,PPCOEF, IDERIV=1, NINTV=NPPCF)
        WRITE (NOUT,'(2F12.3,F12.6,F12.3,F12.6)') X, S, F(X) - S, DS,&
            DF(X), DS
    2 0 ~ C O N T I N U E ~
99999 FORMAT (11X, 'X', 8X, 'S(X)', 7X, 'Error', 7X, 'S''(X)', 7X,&
                'Error')
    END
```


## Output

| X | S (X) | Error | $S^{\prime}(X)$ | Error |
| ---: | ---: | ---: | ---: | ---: |
| 0.000 | 0.000 | 0.000000 | 1.000 | -0.000112 |
| 0.053 | 0.055 | 0.000000 | 1.109 | 0.000030 |
| 0.105 | 0.117 | 0.000000 | 1.228 | -0.000008 |
| 0.158 | 0.185 | 0.000000 | 1.356 | 0.000002 |
| 0.211 | 0.260 | 0.000000 | 1.494 | 0.000000 |
| 0.263 | 0.342 | 0.000000 | 1.643 | 0.000000 |
| 0.316 | 0.433 | 0.000000 | 1.804 | -0.000001 |
| 0.368 | 0.533 | 0.000000 | 1.978 | 0.000002 |
| 0.421 | 0.642 | 0.000000 | 2.165 | 0.000001 |
| 0.474 | 0.761 | 0.000000 | 2.367 | 0.000000 |
| 0.526 | 0.891 | 0.000000 | 2.584 | -0.000001 |
| 0.579 | 1.033 | 0.000000 | 2.817 | 0.000001 |
| 0.632 | 1.188 | 0.000000 | 3.068 | 0.000001 |
| 0.684 | 1.356 | 0.000000 | 3.338 | 0.000001 |
| 0.737 | 1.540 | 0.000000 | 3.629 | 0.000001 |
| 0.789 | 1.739 | 0.000000 | 3.941 | 0.000000 |
| 0.842 | 1.955 | 0.000000 | 4.276 | -0.000006 |
| 0.895 | 2.189 | 0.000000 | 4.636 | 0.000024 |
| 0.947 | 2.443 | 0.000000 | 5.022 | -0.000090 |
| 1.000 | 2.718 | 0.000000 | 5.436 | 0.000341 |

## PP1GD

Evaluates the derivative of a piecewise polynomial on a grid.

## Required Arguments

XVEC - Array of length N containing the points at which the piecewise polynomial is to be evaluated. (Input)
The points in XVEC should be strictly increasing.
BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise polynomial representation. (Input) BREAK must be strictly increasing.

PPCOEF - Matrix of size KORDER by NINTV containing the local coefficients of the polynomial pieces. (Input)

VALUE - Array of length $N$ containing the values of the IDERIV-th derivative of the piecewise polynomial at the points in XVEC. (Output)

## Optional Arguments

IDERIV - Order of the derivative to be evaluated. (Input)
In particular, IDERIV $=0$ returns the values of the piecewise polynomial.
Default: $\operatorname{IDERIV}=1$.
$\mathbf{N}$ - Length of vector XVEC. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{XVEC}, 1)$.
KORDER — Order of the polynomial. (Input)
Default: KORDER $=\operatorname{size}($ PPCOEF, 1$)$.
NINTV - Number of polynomial pieces. (Input)
Default: NINTV = size (PPCOEF,2).

## FORTRAN 90 Interface

Generic: CALL PP1GD (XVEC, BREAK, PPCOEF, VALUE [,...])

Specific: The specific interface names are S_PP1GD and D_PP1GD.

## FORTRAN 77 Interface

Single: CALL PP1GD (IDERIV, N, XVEC, KORDER, NINTV, BREAK, PPCOEF, VALUE)

Double: The double precision name is DPP1GD.

## Description

The routine PP1GD evaluates a piecewise polynomial function $f$ (or its derivative) at a vector of points. That is, given a vector $x$ of length $n$ satisfying $x_{i}<x_{i+1}$ for $i=1, \ldots, n-1$, a derivative
value $j$, and a piecewise polynomial function $f$ that is represented by a breakpoint sequence and coefficient matrix this routine returns the values

$$
f^{(j)}\left(x_{i}\right) \quad i=1, \ldots, n
$$

in the array VALUE. The functionality of this routine is the same as that of PPDER called in a loop, however PP1GD is much more efficient.

## Comments

1. Workspace may be explicitly provided, if desired, by use of P21GD/DP21GD. The reference is:

CALL P21GD (IDERIV, N, XVEC, KORDER, NINTV, BREAK, PPCOEF, VALUE, IWK, WORK1, WORK2)

The additional arguments are as follows:

IWK - Array of length N.
WORK1 - Array of length N.
WORK2 - Array of length N.
2. Informational error

Type Code
44 The points in XVEC must be strictly increasing.

## Example

To illustrate the use of PP1GD, we modify the example program for PPDER. In this example, a piecewise polynomial interpolant to $F$ is computed. The values of this polynomial are then compared with the exact function values. The routine PP1GD is based on the routine PPVALU in de Boor (1978, page 89).

```
USE IMSL_LIBRARIES
```

IMPLICIT NONE
INTEGER KORDER, N, NCOEF, NDATA, NKNOT
PARAMETER (KORDER=4, N=20, NCOEF=20, NDATA=20, \&
NKNOT=NDATA+KORDER)
!
INTEGER I, NINTV, NOUT, NPPCF
REAL BREAK (NCOEF) , BSCOEF (NCOEF), DF, EXP, F, \&
FDATA (NDATA), FLOAT, PPCOEF (KORDER, NCOEF), VALUE1 (N) , \&
VALUE2 (N), X, XDATA (NDATA), XKNOT (NKNOT), XVEC (N)
INTRINSIC EXP, FLOAT
$!$
$\mathrm{F}(\mathrm{X})=\mathrm{X}^{*} \mathrm{EXP}(\mathrm{X})$
$\mathrm{DF}(\mathrm{X})=(\mathrm{X}+1) * .\mathrm{EXP}(\mathrm{X})$
DO $10 \mathrm{I}=1$, NDATA

```
                XDATA(I) = FLOAT(I-1)/FLOAT(NDATA-1)
                FDATA(I) = F(XDATA(I))
    1 0
! Generate knot sequence
    CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
                            Compute the B-spline interpolant
    CALL BSINT (NCOEF, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
    Convert to piecewise polynomial
    CALL BSCPP (KORDER, XKNOT, NCOEF, BSCOEF, NPPCF, BREAK, PPCOEF)
    DO 20 I=1, N
        XVEC(I) = FLOAT(I-1)/FLOAT (N-1)
    20 CONTINUE
    Compute values of the piecewise
    polynomial
    NINTV = NPPCF
    CALL PP1GD (XVEC, BREAK, PPCOEF, VALUE1, IDERIV=0, NINTV=NINTV)
                                    Compute the values of the first
                                    derivative of the piecewise
    polynomial
    CALL PP1GD (XVEC, BREAK, PPCOEF, VALUE2, IDERIV=1, NINTV=NINTV)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
! Print the results on a uniform
    DO 30 I=1, N
        WRITE (NOUT,99999) XVEC(I), VALUE1(I), F(XVEC(I)) - VALUE1(I)&
                , VALUE2(I), DF(XVEC(I)) - VALUE2(I)
    3 0 ~ C O N T I N U E
99998 FORMAT (11X, 'X', 8X, 'S(X)', 7X, 'Error', 7X, 'S''(X)', 7X,&
        'Error')
99999 FORMAT (' ', 2F12.3, F12.6, F12.3, F12.6)
    END
```

        Output
    | $X$ | $S(X)$ | Error | $S^{\prime}(X)$ | Error |
| :---: | :---: | :---: | ---: | ---: |
| 0.000 | 0.000 | 0.000000 | 1.000 | -0.000112 |
| 0.053 | 0.055 | 0.000000 | 1.109 | 0.000030 |
| 0.105 | 0.117 | 0.000000 | 1.228 | -0.000008 |
| 0.158 | 0.185 | 0.000000 | 1.356 | 0.000002 |
| 0.211 | 0.260 | 0.000000 | 1.494 | 0.000000 |
| 0.263 | 0.342 | 0.000000 | 1.643 | 0.000000 |
| 0.316 | 0.433 | 0.000000 | 1.804 | -0.000001 |
| 0.368 | 0.533 | 0.000000 | 1.978 | 0.000002 |
| 0.421 | 0.642 | 0.000000 | 2.165 | 0.000001 |
| 0.474 | 0.761 | 0.000000 | 2.367 | 0.000000 |
| 0.526 | 0.891 | 0.000000 | 2.584 | -0.000001 |
| 0.579 | 1.033 | 0.000000 | 2.817 | 0.000001 |
| 0.632 | 1.188 | 0.000000 | 3.068 | 0.000001 |
| 0.684 | 1.356 | 0.000000 | 3.338 | 0.000001 |
| 0.737 | 1.540 | 0.000000 | 3.629 | 0.000001 |
| 0.789 | 1.739 | 0.000000 | 3.941 | 0.000000 |


| 0.842 | 1.955 | 0.000000 | 4.276 | -0.000006 |
| ---: | ---: | ---: | ---: | ---: |
| 0.895 | 2.189 | 0.000000 | 4.636 | 0.000024 |
| 0.947 | 2.443 | 0.000000 | 5.022 | -0.000090 |
| 1.000 | 2.718 | 0.000000 | 5.436 | 0.000341 |

## PPITG

This function evaluates the integral of a piecewise polynomial.

## Function Return Value

PPITG - Value of the integral from $A$ to $B$ of the piecewise polynomial. (Output)

## Required Arguments

$\boldsymbol{A}-$ Lower limit of integration. (Input)
$\boldsymbol{B}$ - Upper limit of integration. (Input)
BREAK - Array of length NINTV +1 containing the breakpoints for the piecewise polynomial. (Input)
BREAK must be strictly increasing.
PPCOEF - Array of size KORDER * NINTV containing the local coefficients of the piecewise polynomial pieces. (Input) PPCOEF is treated internally as a matrix of size KORDER by NINTV.

## Optional Arguments

KORDER - Order of the polynomial. (Input)
Default: KORDER = size (PPCOEF,1).
NINTV - Number of piecewise polynomial pieces. (Input)
Default: NINTV = size (PPCOEF,2).

## FORTRAN 90 Interface

Generic: PP1TG (A, B, BREAK, PPCOEF [, ...])
Specific: The specific interface names are $S_{-P P 1 T G}$ and D_PP1TG.

## FORTRAN 77 Interface

Single: PP1TG (A, B, KORDER, NINTV, BREAK, PPCOEF)
Double: The double precision function name is DPP1TG.

## Description

The routine PPITG evaluates the integral of a piecewise polynomial over an interval.

## Example

In this example, we compute a quadratic spline interpolant to the function $x^{2}$ using the IMSL routine BSINT. We then evaluate the integral of the spline interpolant over the intervals $[0,1 / 2]$ and $[0,2]$. The interpolant reproduces $x^{2}$, and hence, the values of the integrals are $1 / 24$ and $8 / 3$, respectively.

```
USE IMSL_LIBRARIES
```

IMPLICIT NONE
INTEGER KORDER, NDATA, NKNOT
PARAMETER (KORDER=3, NDATA=10, NKNOT=NDATA+KORDER)
INTEGER I, NOUT, NPPCF
REAL A, B, BREAK (NDATA) , BSCOEF (NDATA), EXACT, F, \&
FDATA (NDATA), FI, FLOAT, PPCOEF (KORDER,NDATA) , \&
VALUE, $\mathrm{X}, \mathrm{XDATA}(\mathrm{NDATA})$, XKNOT (NKNOT)
INTRINSIC FLOAT
$!$
$\mathrm{F}(\mathrm{X})=\mathrm{X} * \mathrm{X}$
$\mathrm{FI}(\mathrm{X})=\mathrm{X} * \mathrm{X} * \mathrm{X} / 3.0$
Set up interpolation points
DO $10 \quad \mathrm{I}=1$, NDATA
XDATA (I) $=$ FLOAT $(I-1) / F L O A T(N D A T A-1)$
FDATA (I) $=F(X D A T A(I))$
10 CONTINUE
CALL BSNAK (NDATA, XDATA, KORDER, XKNOT)
Interpolate
CALL BSINT (NDATA, XDATA, FDATA, KORDER, XKNOT, BSCOEF)
Convert to piecewise polynomial
CALL BSCPP (KORDER, XKNOT, NDATA, BSCOEF, NPPCF, BREAK, PPCOEF)
Compute the integral of $F$ over
[0.0,0.5]
$\mathrm{A}=0.0$
$\mathrm{B}=0.5$
VALUE $=\operatorname{PPITG}(A, B, B R E A K$, PPCOEF, NINTV=NPPCF)
$\mathrm{EXACT}=\mathrm{FI}(\mathrm{B})-\mathrm{FI}(\mathrm{A})$
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) A, B, VALUE, EXACT, EXACT - VALUE
Compute the integral of $F$ over
[0.0,2.0]
$\mathrm{A}=0.0$
$\mathrm{B}=2.0$
VALUE $=\operatorname{PPITG}(A, B, B R E A K$, PPCOEF, NINTV=NPPCF)
$\mathrm{EXACT}=\mathrm{FI}(\mathrm{B})-\mathrm{FI}(\mathrm{A})$
WRITE (NOUT,99999) A, B, VALUE, EXACT, EXACT - VALUE

```
99999 FORMAT (' On the closed interval (', F3.1, ',', F3.1,&
    ') we have :', /, 1X, 'Computed Integral = ', F10.5, /,&
    1X, 'Exact Integral = ', F10.5, /, 1X, 'Error '&
    ,' = ',F10.6, /, /)
!
END
```


## Output

```
On the closed interval (0.0,0.5) we have :
Computed Integral = 0.04167
Exact Integral = 0.04167
Error = 0.000000
On the closed interval (0.0,2.0) we have :
Computed Integral = 2.66667
Exact Integral = 2.66667
Error = 0.000001
```


## QDVAL

This function evaluates a function defined on a set of points using quadratic interpolation.

## Function Return Value

QDVAL - Value of the quadratic interpolant at X. (Output)

## Required Arguments

$\boldsymbol{X}$ - Coordinate of the point at which the function is to be evaluated. (Input)
XDATA - Array of length NDATA containing the location of the data points. (Input) XDATA must be strictly increasing.

FDATA - Array of length NDATA containing the function values. (Input) FDATA(I) is the value of the function at XDATA(I).

## Optional Arguments

NDATA - Number of data points. (Input) NDATA must be at least 3 . Default: NDATA $=\operatorname{size}(\mathrm{XDATA}, 1)$.

CHECK - Logical variable that is .TRUE. if checking of XDATA is required or .FALSE. if checking is not required. (Input)
Default: CHECK = .TRUE.

## FORTRAN 90 Interface

Generic: QDVAL (X, XDATA, FDATA [, ...])
Specific: The specific interface names are S_QDVAL and D_QDVAL.

## FORTRAN 77 Interface

Single: QDVAL (X, NDATA, XDATA, FDATA, CHECK)

Double: The double precision name is DQDVAL.

## Description

The function QDVAL interpolates a table of values, using quadratic polynomials, returning an approximation to the tabulated function. Let $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots, n$ be the tabular data. Given a number $x$ at which an interpolated value is desired, we first find the nearest interior grid point $x_{i}$. A quadratic interpolant $q$ is then formed using the three points $\left(x_{i-1}, f_{i-1}\right),\left(x_{i}, f_{i}\right)$, and $\left(x_{i+1}, f_{i+1}\right)$. The number returned by QDVAL is $q(x)$.

## Comments

Informational error
Type Code
43 The XDATA values must be strictly increasing.

## Example

In this example, the value of $\sin x$ is approximated at $\pi / 4$ by using QDVAL on a table of 33 equally spaced values.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=33)
    INTEGER I, NOUT
    REAL F, FDATA(NDATA), H, PI, QT, SIN, X,&
    XDATA(NDATA)
    INTRINSIC SIN
    F(X) = SIN(X)
    XDATA(1) = 0.0
    FDATA(1) = F(XDATA(1))
    H = 1.0/32.0
    DO 10 I=2, NDATA
    XDATA(I) = XDATA(I-1) + H
    FDATA(I) = F(XDATA(I))
10 CONTINUE
```

$!$

```
! Get value of PI and set X
    PI = CONST('PI')
    X = PI/4.0
! Evaluate at PI/4
    QT = QDVAL(X,XDATA,FDATA)
    CALL UMACH (2, NOUT)
! Print results
    WRITE (NOUT,99999) X, F(X), QT, (F(X)-QT)
!
99999 FORMAT (15X, 'X', 6X, 'F(X)', 6X, 'QDVAL', 5X, 'ERROR', //, 6X,&
        4F10.3, /)
        END
```


## Output

| X | $\mathrm{F}(\mathrm{X})$ | QDVAL | ERROR |
| ---: | :---: | :---: | :---: |
| 0.785 | 0.707 | 0.707 | 0.000 |

## QDDER

This function evaluates the derivative of a function defined on a set of points using quadratic interpolation.

## Function Return Value

QDDER - Value of the IDERIV-th derivative of the quadratic interpolant at X . (Output)

## Required Arguments

IDERIV - Order of the derivative. (Input)
$\boldsymbol{X}$ - Coordinate of the point at which the function is to be evaluated. (Input)
XDATA - Array of length NDATA containing the location of the data points. (Input) XDATA must be strictly increasing.

FDATA - Array of length NDATA containing the function values. (Input) FDATA(I) is the value of the function at XDATA(I).

## Optional Arguments

NDATA - Number of data points. (Input)
NDATA must be at least three.
Default: NDATA $=$ size $(X D A T A, 1)$.

CHECK - Logical variable that is . TRUE. if checking of XDATA is required or . FALSE. if checking is not required. (Input)
Default: $\mathrm{CHECK}=$. TRUE.

## FORTRAN 90 Interface

Generic: QDDER(IDERIV, X, XDATA, FDATA [,...])
Specific: The specific interface names are S_QDVAL and D_QDVAL.

## FORTRAN 77 Interface

Single: QDDER(IDERIV, $\mathrm{X}, \mathrm{NDATA}, \mathrm{XDATA}, ~ F D A T A, ~ C H E C K)$
Double: The double precision function name is DQDVAL.

## Description

The function QDDER interpolates a table of values, using quadratic polynomials, returning an approximation to the derivative of the tabulated function. Let $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots, n$ be the tabular data. Given a number $x$ at which an interpolated value is desired, we first find the nearest interior grid point $x_{i}$. A quadratic interpolant $q$ is then formed using the three points $\left(x_{i_{-1}}, f_{i_{-1}}\right)$ $\left(x_{i}, f_{i}\right)$, and $\left(x_{i+1}, f_{i+1}\right)$. The number returned by QDDER is $q^{(j)}(x)$, where $j=$ IDERIV.

## Comments

1. Informational error
```
Type Code
    4 3 The XDATA values must be strictly increasing.
```

2. Because quadratic interpolation is used, if the order of the derivative is greater than two, then the returned value is zero.

## Example

In this example, the value of $\sin x$ and its derivatives are approximated at $\pi / 4$ by using QDDER on a table of 33 equally spaced values.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=33)
INTEGER I, IDERIV, NOUT
REAL COS, F, F1, F2, FDATA(NDATA), H, PI,&
    QT, SIN, X, XDATA(NDATA)
LOGICAL CHECK
INTRINSIC COS, SIN
```

!

```
! Define function and derivatives
    F(X) = SIN (X)
    F1(X) = COS (X)
    F2(X) = -SIN(X)
!
    XDATA(1) = 0.0
    FDATA(1) = F(XDATA(1))
    H=1.0/32.0
    DO 10 I=2, NDATA
        XDATA(I) = XDATA(I-1) + H
        FDATA(I) = F(XDATA(I))
    10 CONTINUE
! Get value of PI and set X
    PI = CONST('PI')
    X = PI/4.0
    CHECK = .TRUE.
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
    IDERIV = 0
    QT = QDDER(IDERIV,X,XDATA,FDATA, CHECK=CHECK)
    WRITE (NOUT,99999) X, IDERIV, F(X), QT, (F (X)-QT)
    CHECK = .FALSE.
    IDERIV = 1
    QT = QDDER(IDERIV,X,XDATA,FDATA)
    WRITE (NOUT,99999) X, IDERIV, F1(X), QT, (F1(X)-QT)
                            Evaluate second derivative at PI/4
    IDERIV = 2
    QT = QDDER(IDERIV,X,XDATA,FDATA, CHECK=CHECK)
    WRITE (NOUT,99999) X, IDERIV, F2(X), QT, (F2(X)-QT)
!
99998 FORMAT (33X, 'IDER', /, 15X, 'X', 6X, 'IDER', 6X, 'F (X)',&
    5X, 'QDDER', 6X, 'ERROR', //)
99999 FORMAT (7X, F10.3, I8, 3F12.3/)
    END
```


## Output

|  | IDER |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| X | IDER | F (X) | QDDER | ERROR |
| 0.785 | 0 | 0.707 | 0.707 | 0.000 |
| 0.785 | 1 | 0.707 | 0.707 | 0.000 |
| 0.785 | 2 | -0.707 | -0.704 | -0.003 |

## QD2VL

This function evaluates a function defined on a rectangular grid using quadratic interpolation.

## Function Return Value

QD2VL - Value of the function at (X, Y). (Output)

## Required Arguments

$\boldsymbol{X}-x$-coordinate of the point at which the function is to be evaluated. (Input)
$\boldsymbol{Y}-y$-coordinate of the point at which the function is to be evaluated. (Input)
XDATA - Array of length NXDATA containing the location of the data points in the $x$-direction. (Input) XDATA must be increasing.

YDATA - Array of length NYDATA containing the location of the data points in the $y$-direction. (Input) YDATA must be increasing.

FDATA - Array of size NXDATA by NYDATA containing function values. (Input) FDATA ( $I, \mathrm{~J}$ ) is the value of the function at (XDATA (I), YDATA( J$)$ ).

## Optional Arguments

NXDATA - Number of data points in the $x$-direction. (Input) NXDATA must be at least three. Default: NXDATA $=\operatorname{size}($ XDATA, 1$)$.

NYDATA - Number of data points in the $y$-direction. (Input) NYDATA must be at least three. Default: NYDATA $=\operatorname{size}($ YDATA, 1$)$.
$\mathbf{L D F}$ - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
LDF must be at least as large as NXDATA. Default: LDF = size (FDATA,1).

CHECK - Logical variable that is . TRUE. if checking of XDATA and YDATA is required or .FALSE . if checking is not required. (Input)
Default: CHECK = .TRUE.

## FORTRAN 90 Interface

Generic: QD2VL (X, Y, XDATA, YDATA, FDATA [, ...])

Specific: The specific interface names are S_QD2VL and D_QD2VL.

## FORTRAN 77 Interface

Single: QD2VL (X, Y, NXDATA, XDATA, NYDATA, YDATA, FDATA, LDF, CHECK)
Double: The double precision function name is DQD2VL.

## Description

The function QD2VL interpolates a table of values, using quadratic polynomials, returning an approximation to the tabulated function. Let $\left(x_{i}, y_{j}, f_{i j}\right)$ for $i=1, \ldots, n_{x}$ and $j=1, \ldots, n_{y}$ be the tabular data. Given a point $(x, y)$ at which an interpolated value is desired, we first find the nearest interior grid point $\left(x_{i}, y_{j}\right)$. A bivariate quadratic interpolant $q$ is then formed using six points near $(x, y)$. Five of the six points are $\left(x_{i}, y_{j}\right),\left(x_{i \pm 1}, y_{j}\right)$, and $\left(x_{i}, y_{j \pm 1}\right)$. The sixth point is the nearest point to $(x, y)$ of the grid points $\left(x_{i \pm 1}, y_{j \pm 1}\right)$. The value $q(x, y)$ is returned by QD2VL.

## Comments

Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 6 | The XDATA values must be strictly increasing. |
| 4 | 7 | The YDATA values must be strictly increasing. |

## Example

In this example, the value of $\sin (x+y)$ at $x=y=\pi / 4$ is approximated by using QDVAL on a table of size $21 \times 42$ equally spaced values on the unit square.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDF, NXDATA, NYDATA
PARAMETER (NXDATA=21, NYDATA=42, LDF=NXDATA)
INTEGER I, J, NOUT
REAL F, FDATA(LDF,NYDATA), FLOAT, PI, Q, &
    SIN, X, XDATA(NXDATA), Y, YDATA(NYDATA)
INTRINSIC FLOAT, SIN
F(X,Y) = SIN(X+Y)
DO 10 I=1, NXDATA
        XDATA(I) = FLOAT(I-1)/FLOAT (NXDATA-1)
    10 CONTINUE
DO 20 I=1, NYDATA
    YDATA(I) = FLOAT(I-1)/FLOAT(NYDATA-1)
20 CONTINUE
```

```
DO 30 I=1, NXDATA
```

DO 30 I=1, NXDATA
DO 30 J=1, NYDATA
DO 30 J=1, NYDATA
FDATA(I,J) = F(XDATA(I),YDATA(J))

```
    FDATA(I,J) = F(XDATA(I),YDATA(J))
```

!
!
$!$
!

```
    3 0 ~ C O N T I N U E ~
! CAIT UMACH (2, NOUT) Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
    PI = CONST('PI')
    X = PI/4.0
    Y = PI/4.0
! Evaluate quadratic at (X,Y)
    Q = QD2VL (X,Y,XDATA,YDATA,FDATA)
    Print results
    WRITE (NOUT,'(5F12.4)') X, Y, F(X,Y), Q, (Q-F (X,Y))
99999 FORMAT (10X, 'X', 11X, 'Y', 7X, 'F(X,Y)', 7X, 'QD2VL', 9X,&
            'DIF')
    END
```


## Output

| X | Y | $\mathrm{F}(\mathrm{X}, \mathrm{Y})$ | QD2VL | DIF |
| ---: | ---: | ---: | ---: | ---: |
| 0.7854 | 0.7854 | 1.0000 | 1.0000 | 0.0000 |

## QD2DR

This function evaluates the derivative of a function defined on a rectangular grid using quadratic interpolation.

## Function Return Value

QD2DR - Value of the (IXDER, IYDER) derivative of the function at (X, Y). (Output)

## Required Arguments

IXDER - Order of the $x$-derivative. (Input)
IYDER - Order of the $y$-derivative. (Input)
$\boldsymbol{X}$ - X-coordinate of the point at which the function is to be evaluated. (Input)
$\boldsymbol{Y}$ - Y-coordinate of the point at which the function is to be evaluated. (Input)
XDATA - Array of length NXDATA containing the location of the data points in the $x$-direction. (Input)
XDATA must be increasing.
YDATA - Array of length NYDATA containing the location of the data points in the $y$-direction. (Input) YDATA must be increasing.

FDATA - Array of size NXDATA by NYDATA containing function values. (Input) $\operatorname{FDATA}(I, J)$ is the value of the function at (XDATA(I), YDATA(J)).

## Optional Arguments

NXDATA - Number of data points in the $x$-direction. (Input)
NXDATA must be at least three.
Default: NXDATA $=$ size $(X D A T A, 1)$.

NYDATA - Number of data points in the $y$-direction. (Input) NYDATA must be at least three.
Default: NYDATA = size $($ YDATA, 1$)$.
$\boldsymbol{L D F}$ - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
LDF must be at least as large as NXDATA.
Default: LDF = size (FDATA,1).
CHECK - Logical variable that is . TRUE. if checking of XDATA and YDATA is required or .FALSE. if checking is not required. (Input)
Default: CHECK = .TRUE.

## FORTRAN 90 Interface

Generic: QD2DR (IXDER, IYDER, X, Y, XDATA, YDATA, FDATA [ , ...])
Specific: The specific interface names are S_QD2DR and D_QD2DR.

## FORTRAN 77 Interface

Single: $\quad$ QD2DR (IXDER, IYDER, $\mathrm{X}, \mathrm{Y}, \mathrm{NXDATA}, \mathrm{XDATA}, \mathrm{NYDATA}, \mathrm{YDATA}, \mathrm{FDATA}, \mathrm{LDF}$,

Double: The double precision fucntion name is DQD2DR.

## Description

The function QD2DR interpolates a table of values, using quadratic polynomials, returning an approximation to the tabulated function. Let $\left(x_{i}, y_{j}, f_{i j}\right)$ for $i=1, \ldots, n_{x}$ and $j=1, \ldots, n_{y}$ be the tabular data. Given a point $(x, y)$ at which an interpolated value is desired, we first find the nearest interior grid point $\left(x_{i}, y_{j}\right)$. A bivariate quadratic interpolant $q$ is then formed using six points near $(x, y)$. Five of the six points are $\left(x_{i}, y_{j}\right),\left(x_{i \pm 1}, y_{j}\right)$, and $\left(x_{i}, y_{j \pm 1}\right)$. The sixth point is the nearest point to $(x, y)$ of the grid points $\left(x_{i \pm 1}, y_{j \pm 1}\right)$. The value $q^{(p, r)}(x, y)$ is returned by QD 2 DR , where $p=$ IXDER and $r=$ IYDER.

## Comments

1. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 6 | The XDATA values must be strictly increasing. |
| 4 | 7 | The YDATA values must be strictly increasing. |

2. Because quadratic interpolation is used, if the order of any derivative is greater than two, then the returned value is zero.

## Example

In this example, the partial derivatives of $\sin (x+y)$ at $x=y=\pi / 3$ are approximated by using $Q D 2 D R$ on a table of size $21 \times 42$ equally spaced values on the rectangle [0, 2] $\times[0,2]$.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDF, NXDATA, NYDATA
PARAMETER (NXDATA=21, NYDATA=42, LDF=NXDATA)
!
INTEGER I, IXDER, IYDER, J, NOUT
REAL F, FDATA (LDF,NYDATA), FLOAT, FU, FUNC, PI, Q,\& SIN, X, XDATA (NXDATA), Y, YDATA(NYDATA)
INTRINSIC FLOAT, SIN
EXTERNAL FUNC
! Define function
$F(X, Y)=\operatorname{SIN}(X+Y)$
Set up X-grid
DO $10 \mathrm{I}=1$, NXDATA
XDATA $(I)=2.0 *(F L O A T(I-1) / F L O A T(N X D A T A-1))$
10 CONTINUE
Set up Y-grid
DO 20 I=1, NYDATA YDATA (I) $=2.0 *(F L O A T(I-1) / F L O A T(N Y D A T A-1))$
CONTINUE
DO 30 I=1, NXDATA DO $30 \mathrm{~J}=1$, NYDATA FDATA (I, J) = F(XDATA (I), YDATA (J))
30 CONTINUE
!
CALL UMACH (2, NOUT)
WRITE (NOUT,99998)
Get output unit number
Write heading
Check XDATA and YDATA
Get value for $P I$ and set $X$ and $Y$
PI = CONST('PI')
$\mathrm{X}=\mathrm{PI} / 3.0$
$\mathrm{Y}=\mathrm{PI} / 3.0$
Evaluate and print the function and its derivatives at $\mathrm{X}=\mathrm{PI} / 3$ and $\mathrm{Y}=\mathrm{PI} / 3$.

```
    DO 40 IXDER=0, 1
        DO 40 IYDER=0, 1
        Q = QD2DR(IXDER,IYDER,X,Y,XDATA,YDATA,FDATA)
        FU = FUNC(IXDER,IYDER,X,Y)
        WRITE (NOUT,99999) X, Y, IXDER, IYDER, FU, Q, (FU-Q)
    40 CONTINUE
!
99998 FORMAT (32X, '(IDX,IDY)', /, 8X, 'X', 8X, 'Y', 3X, 'IDX', 2X,&
            'IDY', 3X, 'F (X,Y)', 3X, 'QD2DR', 6X, 'ERROR')
99999 FORMAT (2F9.4, 2I5, 3X, F9.4, 2X, 2F11.4)
        END
        REAL FUNCTION FUNC (IX, IY, X, Y)
        INTEGER IX, IY
        REAL X, Y
!
        REAL COS, SIN
        INTRINSIC COS, SIN
!
        IF (IX.EQ.0 .AND. IY.EQ.O) THEN
            FUNC = SIN (X+Y)
        ELSE IF (IX.EQ.O .AND. IY.EQ.1) THEN
            FUNC = COS (X+Y)
        ELSE IF (IX.EQ.1 .AND. IY.EQ.0) THEN
            FUNC = COS (X+Y)
        ELSE IF (IX.EQ.1 .AND. IY.EQ.1) THEN
                                    Define (1,1) derivative
            FUNC = -SIN (X+Y)
        ELSE
            FUNC = 0.0
        END IF
        RETURN
        END
```


## Output

|  |  |  |  | (IDX, ID |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| X | Y | IDX | IDY | F | ( $\mathrm{X}, \mathrm{Y}$ ) | QD2DR | ERROR |
| 1.0472 | 1.0472 | 0 | 0 | 0.8660 |  | 0.8661 | -0.0001 |
| 1.0472 | 1.0472 | 0 | 1 | -0.5000 |  | -0.4993 | -0.0007 |
| 1.0472 | 1.0472 | 1 | 0 | -0.5000 |  | -0.4995 | -0.0005 |
| 1.0472 | 1.0472 | 1 | 1 | -0.8660 |  | -0.8634 | -0.0026 |

## QD3VL

This function evaluates a function defined on a rectangular three-dimensional grid using quadratic interpolation.

## Function Return Value

$\mathbf{Q D} 3 V \mathrm{~L}$ - Value of the function at ( $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ ). (Output)

## Required Arguments

$\boldsymbol{X}-x$-coordinate of the point at which the function is to be evaluated. (Input)
$\boldsymbol{Y}-y$-coordinate of the point at which the function is to be evaluated. (Input)
$\mathbf{Z}-z$-coordinate of the point at which the function is to be evaluated. (Input)

XDATA - Array of length NXDATA containing the location of the data points in the $x$-direction. (Input) XDATA must be increasing.

YDATA - Array of length NYDATA containing the location of the data points in the $y$-direction. (Input) YDATA must be increasing.

ZDATA - Array of length NZDATA containing the location of the data points in the $z$-direction. (Input) ZDATA must be increasing.

FDATA - Array of size NXDATA by NYDATA by NZDATA containing function values. (Input) $\operatorname{FDATA}(I, J, K)$ is the value of the function at (XDATA(I), YDATA(J), ZDATA(K)).

## Optional Arguments

NXDATA - Number of data points in the $x$-direction. (Input) NXDATA must be at least three.
Default: NXDATA $=\operatorname{size}($ XDATA, 1$)$.

NYDATA - Number of data points in the $y$-direction. (Input)
NYDATA must be at least three.
Default: NYDATA = size (YDATA, 1 ).
NZDATA - Number of data points in the $z$-direction. (Input) NZDATA must be at least three.
Default: NZDATA $=\operatorname{size}($ ZDATA, 1$)$.
LDF - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
LDF must be at least as large as nXDATA.
Default: LDF = size (FDATA,1).
MDF - Middle (second) dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
MDF must be at least as large as NYDATA.
Default: MDF $=\operatorname{size}($ FDATA, 2 ).

CHECK - Logical variable that is .TRUE. if checking of XDATA, YDATA, and ZDATA is required or .FALSE. if checking is not required. (Input)
Default: $\mathrm{CHECK}=$. TRUE.

## FORTRAN 90 Interface

Generic: QD3VL (X, Y, Z, XDATA, YDATA, ZDATA, FDATA [, ...])
Specific: The specific interface names are S_QD3VL and D_QD3VL.

## FORTRAN 77 Interface

Single: QD3VL (X, Y, Z, NXDATA, XDATA, NYDATA, YDATA, NZDATA, ZDATA, FDATA, LDF, MDF, CHECK)

Double: The double precision function name is DQD3VL.

## Description

The function QD3VL interpolates a table of values, using quadratic polynomials, returning an approximation to the tabulated function. Let $\left(x_{i}, y_{j}, z_{k}, f_{i j k}\right)$ for $i=1, \ldots, n_{x}, j=1, \ldots, n_{y}$, and $k=1, \ldots, n_{z}$ be the tabular data. Given a point $(x, y, z)$ at which an interpolated value is desired, we first find the nearest interior grid point $\left(x_{i}, y_{j}, z_{k}\right.$ ). A trivariate quadratic interpolant $q$ is then formed. Ten points are needed for this purpose. Seven points have the form

$$
\left(x_{i}, y_{j}, z_{k}\right),\left(x_{i \pm 1}, y_{j}, z_{k}\right),\left(x_{i}, y_{j \pm 1}, z_{k}\right) \text { and }\left(x_{i}, y_{j}, z_{k \pm 1}\right)
$$

The last three points are drawn from the vertices of the octant containing $(x, y, z)$. There are four of these vertices remaining, and we choose to exclude the vertex farthest from the center. This has the slightly deleterious effect of not reproducing the tabular data at the eight exterior corners of the table. The value $q(x, y, z)$ is returned by QD3VL.

## Comments

Informational errors
Type Code
$4 \quad 9$ The XDATA values must be strictly increasing.
410 The YDATA values must be strictly increasing.
411 The ZDATA values must be strictly increasing.

## Example

In this example, the value of $\sin (x+y+z)$ at $x=y=z=\pi / 3$ is approximated by using QD3VL on a grid of size $21 \times 42 \times 18$ equally spaced values on the cube $[0,2]^{3}$.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
```

```
    INTEGER LDF, MDF, NXDATA, NYDATA, NZDATA
    PARAMETER (NXDATA=21, NYDATA=42, NZDATA=18, LDF=NXDATA,&
        MDF=NYDATA)
!
    INTEGER I, J, K, NOUT
    REAL F, FDATA(LDF,MDF,NZDATA), FLOAT, PI, Q, &
        SIN, X, XDATA(NXDATA), Y, YDATA(NYDATA), Z,&
        ZDATA(NZDATA)
    INTRINSIC FLOAT, SIN
    F(X,Y,Z)=SIN(X+Y+Z)
    DO 10 I=1, NXDATA
        XDATA(I) = 2.0*(FLOAT (I-1)/FLOAT (NXDATA-1))
    CONTINUE
    DO 20 J=1, NYDATA
        YDATA(J) = 2.0*(FLOAT (J-1)/FLOAT (NYDATA-1))
    CONTINUE
    DO 30 K=1, NZDATA
        ZDATA(K) = 2.0*(FLOAT (K-1)/FLOAT (NZDATA-1))
    CONTINUE
    DO 40 I=1, NXDATA
        DO 40 J=1, NYDATA
            DO 40 K=1, NZDATA
                FDATA(I,J,K) = F(XDATA(I),YDATA(J),ZDATA (K))
    4 0 ~ C O N T I N U E ~
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
        Get value for PI and set values
        for X, Y, and Z
    PI = CONST('PI')
    X = PI/3.0
    Y = PI/3.0
    Z = PI/3.0
    Evaluate quadratic at (X,Y,Z)
    Q = QD3VL (X,Y,Z, XDATA,YDATA,ZDATA,FDATA)
    Print results
    WRITE (NOUT,'(6F11.4)') X, Y, Z, F(X,Y,Z), Q, (Q-F(X,Y,Z))
    FORMAT (10X, 'X', 10X, 'Y', 10X, 'Z', 5X, 'F(X,Y,Z)', 4X, &
        'QD3VL', 6X, 'ERROR')
    END
```

    Output
    | X | Y | Z | F (X, Y, Z) | QD3VL | ERROR |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 1.0472 | 1.0472 | 1.0472 | 0.0000 | 0.0001 | 0.0001 |

## QD3DR

This function evaluates the derivative of a function defined on a rectangular three-dimensional grid using quadratic interpolation.

## Function Return Value

QD3DR - Value of the appropriate derivative of the function at (X, Y, Z). (Output)

## Required Arguments

IXDER - Order of the $x$-derivative. (Input)
IYDER - Order of the $y$-derivative. (Input)
IZDER - Order of the $z$-derivative. (Input)
$\boldsymbol{X}-x$-coordinate of the point at which the function is to be evaluated. (Input)
$\boldsymbol{Y}-y$-coordinate of the point at which the function is to be evaluated. (Input)
$\mathbf{Z}-z$-coordinate of the point at which the function is to be evaluated. (Input)
XDATA - Array of length NXDATA containing the location of the data points in the $x$-direction. (Input)
XDATA must be increasing.
YDATA - Array of length NYDATA containing the location of the data points in the $y$-direction. (Input) YDATA must be increasing.

ZDATA - Array of length NZDATA containing the location of the data points in the $z$-direction. (Input) ZDATA must be increasing.

FDATA - Array of size NXDATA by NYDATA by NZDATA containing function values. (Input) $\operatorname{FDATA}(I, J, K)$ is the value of the function at (XDATA(I), YDATA(J), ZDATA(K)).

## Optional Arguments

NXDATA - Number of data points in the $x$-direction. (Input)
NXDATA must be at least three.
Default: NXDATA $=\operatorname{size}(\mathrm{XDATA}, 1)$.
NYDATA - Number of data points in the $y$-direction. (Input)
nydAta must be at least three.
Default: NYDATA = size (YDATA, 1 ).

NZDATA - Number of data points in the $z$-direction. (Input) NZDATA must be at least three.
Default: NZDATA $=\operatorname{size}($ ZDATA, 1$)$.

LDF - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
LDF must be at least as large as NXDATA.
Default: LDF $=\operatorname{size}($ FDATA, 1$)$.
MDF - Middle (second) dimension of FDATA exactly as specified in the dimension
statement of the calling program. (Input)
MDF must be at least as large as NYDATA.
Default: MDF $=\operatorname{size}$ (FDATA,2).
CHECK - Logical variable that is .TRUE. if checking of XDATA, YDATA, and ZDATA is required or .FALSE. if checking is not required. (Input)
Default: CHECK = .TRUE.

## FORTRAN 90 Interface

Generic: QD3DR (IXDER, IYDER, IZDER, X, Y, Z, XDATA, YDATA, ZDATA, FDATA [, ...])
Specific: The specific interface names are S_QD3DR and D_QD3DR.

## FORTRAN 77 Interface

Single: QD3DR (IXDER, IYDER, IZDER, X, Y, Z, NXDATA, XDATA, NYDATA, YDATA, NZDATA, ZDATA, FDATA, LDF, MDF, CHECK)

Double: The double precision function name is DQD3DR.

## Description

The function QD3DR interpolates a table of values, using quadratic polynomials, returning an approximation to the partial derivatives of the tabulated function. Let

$$
\left(x_{i}, y_{j}, z_{k}, f_{i j k}\right)
$$

for $i=1, \ldots, n_{x}, j=1, \ldots, n_{y}$, and $k=1, \ldots, n_{z}$ be the tabular data. Given a point $(x, y, z)$ at which an interpolated value is desired, we first find the nearest interior grid point $\left(x_{i}, y_{j}, z_{k}\right)$. A trivariate quadratic interpolant $q$ is then formed. Ten points are needed for this purpose. Seven points have the form

$$
\left(x_{i}, y_{j}, z_{k}\right),\left(x_{i \pm 1}, y_{j}, z_{k}\right),\left(x_{i}, y_{j \pm 1}, z_{k}\right) \text { and }\left(x_{i}, y_{j}, z_{k \pm 1}\right)
$$

The last three points are drawn from the vertices of the octant containing $(x, y, z)$. There are four of these vertices remaining, and we choose to exclude the vertex farthest from the center. This has
the slightly deleterious effect of not reproducing the tabular data at the eight exterior corners of the table. The value $q^{(p, r, t)}(x, y, z)$ is returned by QD3DR, where $p=I X D E R, r=I Y D E R$, and $t=I Z D E R$.

## Comments

1. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 9 | The XDATA values must be strictly increasing. |
| 4 | 10 | The YDATA values must be strictly increasing. |
| 4 | 11 | The ZDATA values must be strictly increasing. |

2. Because quadratic interpolation is used, if the order of any derivative is greater than two, then the returned value is zero.

## Example

In this example, the derivatives of $\sin (x+y+z)$ at $x=y=z=\pi / 5$ are approximated by using QD3DR on a grid of size $21 \times 42 \times 18$ equally spaced values on the cube $[0,2]^{3}$.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDF, MDF, NXDATA, NYDATA, NZDATA
PARAMETER (NXDATA=21, NYDATA=42, NZDATA=18, LDF=NXDATA, \& MDF=NYDATA)
!
INTEGER I, IXDER, IYDER, IZDER, J, K, NOUT
REAL F, FDATA (NXDATA,NYDATA,NZDATA), FLOAT, FU, \& FUNC, PI, Q, SIN, X, XDATA(NXDATA), Y,\& YDATA (NYDATA), Z, ZDATA(NZDATA)
INTRINSIC FLOAT, SIN
EXTERNAL FUNC
!
$\mathrm{F}(\mathrm{X}, \mathrm{Y}, \mathrm{Z})=\operatorname{SIN}(\mathrm{X}+\mathrm{Y}+\mathrm{Z})$
! Set up X-grid
DO 10 I=1, NXDATA
XDATA $(I)=2.0 *(F L O A T(I-1) / F L O A T(N X D A T A-1))$
10 CONTINUE
!
DO 20 J=1, NYDATA YDATA $(\mathrm{J})=2.0 *(F L O A T(J-1) / F L O A T(N Y D A T A-1))$
20 CONTINUE
!
DO $30 \mathrm{~K}=1$, NZDATA ZDATA $(\mathrm{K})=2.0 *(\mathrm{FLOAT}(\mathrm{K}-1) /$ FLOAT $(\mathrm{NZDATA}-1))$
30 CONTINUE
!
DO 40 I=1, NXDATA DO 40 J=1, NYDATA

DO 40 K=1, NZDATA
$\operatorname{FDATA}(I, J, K)=F(X D A T A(I), Y D A T A(J), Z D A T A(K))$
40 CONTINUE
!
Get output unit number

```
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999)
!
    PI = CONST('PI')
    X = PI/5.0
    Y = PI/5.0
    Z = PI/5.0
!
    DO 50 IXDER=0, 1
        DO 50 IYDER=0, 1
            DO 50 IZDER=0, 1
                Q = QD3DR(IXDER,IYDER,IZDER,X,Y,Z,XDATA,YDATA,ZDATA,FDATA)
                FU = FUNC(IXDER,IYDER,IZDER,X,Y,Z)
                    WRITE (NOUT,99998) X, Y, Z, IXDER, IYDER, IZDER, FU, Q,&
                    (FU-Q)
    5 0 ~ C O N T I N U E ~
!
99998 FORMAT (3F7.4, 3I5, 4X, F7.4, 8X, 2F10.4)
99999 FORMAT (39X, '(IDX,IDY,IDZ)', /, 6X, 'X', 6X, 'Y', 6X,&
            'Z', 3X, 'IDX', 2X, 'IDY', 2X, 'IDZ', 2X, 'F ',&
            '(X,Y,Z)', 3X, 'QD3DR', 5X, 'ERROR')
    END
    REAL FUNCTION FUNC (IX, IY, IZ, X, Y, Z)
    INTEGER IX, IY, IZ
    REAL X, Y, Z
    REAL COS, SIN
    INTRINSIC COS, SIN
    IF (IX.EQ.0 .AND. IY.EQ.O .AND. IZ.EQ.O) THEN
                            Define (0,0,0) derivative
        FUNC = SIN(X+Y+Z)
    ELSE IF (IX.EQ.0 .AND. IY.EQ.0 .AND. IZ.EQ.1) THEN
        Define (0,0,1) derivative
        FUNC = COS (X+Y+Z)
    ELSE IF (IX.EQ.0 .AND. IY.EQ.1 .AND. IZ.EQ.0) THEN
        Define (0,1,0,) derivative
        FUNC = COS (X+Y+Z)
    ELSE IF (IX.EQ.O .AND. IY.EQ.1 .AND. IZ.EQ.1) THEN
                        Define (0,1,1) derivative
        FUNC = -SIN(X+Y+Z)
    ELSE IF (IX.EQ.1 .AND. IY.EQ.0 .AND. IZ.EQ.0) THEN
                            Define (1,0,0) derivative
        FUNC = COS (X+Y+Z)
    ELSE IF (IX.EQ.1 .AND. IY.EQ.O .AND. IZ.EQ.1) THEN
                            Define (1,0,1) derivative
        FUNC = -SIN(X+Y+Z)
    ELSE IF (IX.EQ.1 .AND. IY.EQ.1 .AND. IZ.EQ.0) THEN
                            Define (1,1,0) derivative
        FUNC = -SIN(X+Y+Z)
    ELSE IF (IX.EQ.1 .AND. IY.EQ.1 .AND. IZ.EQ.1) THEN
    Define (1,1,1) derivative
```

```
    FUNC = - COS (X+Y+Z)
ELSE
    FUNC = 0.0
END IF
RETURN
END
```


## Output

| (IDX, IDY, IDZ) |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| X | Y | Z | IDX | IDY | IDZ | F | ( $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ ) | QD3DR | ERROR |
| 0.6283 | 0.6283 | 0.6283 | 0 | 0 | 0 | 0.9511 |  | 0.9511 | -0.0001 |
| 0.6283 | 0.6283 | 0.6283 | 0 | 0 | 1 | -0.3090 |  | -0.3080 | -0.0010 |
| 0.6283 | 0.6283 | 0.6283 | 0 | 1 | 0 | -0.3090 |  | -0.3088 | 0.0002 |
| 0.6283 | 0.6283 | 0.6283 | 0 | 1 | 1 | -0.9511 |  | -0.9587 | 0.0077 |
| 0.6283 | 0.6283 | 0.6283 | 1 | 0 | 0 | -0.3090 |  | -0.3078 | -0.0012 |
| 0.6283 | 0.6283 | 0.6283 | 1 | 0 | 1 | -0.9511 |  | -0.9348 | -0.0162 |
| 0.6283 | 0.6283 | 0.6283 | 1 | 1 | 0 | -0.9511 |  | -0.9613 | 0.0103 |
| 0.6283 | 0.6283 | 0.6283 | 1 | 1 | 1 | 0.3090 |  | 0.0000 | 0.3090 |

## SURF

Computes a smooth bivariate interpolant to scattered data that is locally a quintic polynomial in two variables.

## Required Arguments

XYDATA - A 2 by NDATA array containing the coordinates of the interpolation points. (Input)
These points must be distinct. The $x$-coordinate of the I -th data point is stored in $\operatorname{XYDATA}(1, \mathrm{I})$ and the $y$-coordinate of the I-th data point is stored in $\operatorname{XYDATA}(2, \mathrm{I})$.

FDATA - Array of length NDATA containing the interpolation values. (Input) FDATA(I) contains the value at (XYDATA $(1, I), \operatorname{XYDATA}(2, I))$.

XOUT - Array of length nxout containing an increasing sequence of points. (Input) These points are the $x$-coordinates of a grid on which the interpolated surface is to be evaluated.

YOUT - Array of length nYout containing an increasing sequence of points. (Input) These points are the $y$-coordinates of a grid on which the interpolated surface is to be evaluated.

SUR - Matrix of size nxout by nyout. (Output)
This matrix contains the values of the surface on the xout by Yout grid, i.e. SUR(I, J) contains the interpolated value at (XOUT(I), YOUT(J)).

## Optional Arguments

NDATA - Number of data points. (Input)
NDATA must be at least four.
Default: NDATA $=\operatorname{size}($ FDATA, 1$)$.

NXOUT - The number of elements in XOUT. (Input)
Default: NXOUT $=$ size $($ XOUT, 1$)$.
NYOUT - The number of elements in YOUT. (Input)
Default: NYOUT = size (YOUT,1).
LDSUR - Leading dimension of SUR exactly as specified in the dimension statement of the calling program. (Input)
LDSUR must be at least as large as NXOUT.
Default: LDSUR = size $(S U R, 1)$.

## FORTRAN 90 Interface

Generic: CALL SURF (XYDATA, FDATA, XOUT, YOUT, SUR [, ...])

Specific: The specific interface names are $S_{-}$SURF and D_SURF.

## FORTRAN 77 Interface

Single: CALL SURF (NDATA, XYDATA, FDATA, NXOUT, NYOUT, XOUT, YOUT, SUR, LDSUR)

Double: The double precision name is DSURF.

## Description

This routine is designed to compute a $C^{l}$ interpolant to scattered data in the plane. Given the data points

$$
\left\{\left(x_{i}, y_{i}, f_{i}\right)\right\}_{i=1}^{N} \text { in } \mathbf{R}^{3}
$$

SURE returns (in SUR, the user-specified grid) the values of the interpolant $s$. The computation of $s$ is as follows: First the Delaunay triangulation of the points

$$
\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{N}
$$

is computed. On each triangle $T$ in this triangulation, $s$ has the form

$$
s(x, y)=\sum_{m+n \leq 5} c_{m n}^{T} x^{m} y^{n} \quad \forall x, y \in T
$$

Thus, $s$ is a bivariate quintic polynomial on each triangle of the triangulation. In addition, we have

$$
s\left(x_{i}, y_{i}\right)=f_{i} \quad \text { for } i=1, \ldots, N
$$

and $s$ is continuously differentiable across the boundaries of neighboring triangles. These conditions do not exhaust the freedom implied by the above representation. This additional freedom is exploited in an attempt to produce an interpolant that is faithful to the global shape properties implied by the data. For more information on this routine, we refer the reader to the article by Akima (1978). The grid is specified by the two integer variables NXOUT, NYOUT that represent, respectively, the number of grid points in the first (second) variable and by two real vectors that represent, respectively, the first (second) coordinates of the grid.

## Comments

1. Workspace may be explicitly provided, if desired, by use of S2RF/DS2RF. The reference is:
```
CALL S2RF (NDATA, XYDATA, FDATA, NXOUT, NYOUT, XOUT, YOUT, SUR, LDSUR, IWK,
WK)
```

The additional arguments are as follows:
$\boldsymbol{I} \mathbf{W K}$ - Work array of length 31 * NDATA $+2 *($ NXOUT * NYOUT $)$.
WK - Work array of length 6 * NDATA.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 5 | The data point values must be distinct. |
| 4 | 6 | The XOUT values must be strictly increasing. |
| 4 | 7 | The YOUT values must be strictly increasing. |

3. This method of interpolation reproduces linear functions.

## Example

In this example, the interpolant to the linear function $3+7 x+2 y$ is computed from 20 data points equally spaced on the circle of radius 3 . We then print the values on a $3 \times 3$ grid.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDSUR, NDATA, NXOUT, NYOUT
PARAMETER (NDATA=20, NXOUT=3, NYOUT=3, LDSUR=NXOUT)
    INTEGER I, J, NOUT
REAL ABS, COS, F, FDATA(NDATA), FLOAT, PI,&
    SIN, SUR(LDSUR,NYOUT), X, XOUT (NXOUT), &
    XYDATA (2,NDATA), Y, YOUT (NYOUT)
INTRINSIC ABS, COS, FLOAT, SIN
F(X,Y)=3.0 + 7.0*X + 2.0*Y
    Get value for PI
PI = CONST('PI')
```

!

```
    DO 10 I=1, NDATA
        XYDATA(1,I) = 3.0*SIN(2.0*PI*FLOAT (I-1)/FLOAT (NDATA))
        XYDATA (2,I) = 3.0*COS (2.0*PI*FLOAT (I-1)/FLOAT (NDATA))
        FDATA(I) = F(XYDATA (1,I),XYDATA (2,I))
    1 0
!
                Set up XOUT and YOUT data on [0,1] by
                [0,1] grid.
    DO 20 I=1, NXOUT
        XOUT(I) = FLOAT(I-1)/FLOAT (NXOUT-1)
    2 0
    CONTINUE
    DO 30 I=1, NXOUT
        YOUT(I) = FLOAT(I-1)/FLOAT(NYOUT-1)
    3 0 ~ C O N T I N U E
! Interpolate scattered data
    CALL SURF (XYDATA, FDATA, XOUT, YOUT, SUR)
        Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
    DO 40 I=1, NYOUT
        DO 40 J=1, NXOUT
        WRITE (NOUT,99999) XOUT(J), YOUT(I), SUR(J,I),&
                                    F(XOUT(J),YOUT (I)),&
                                    ABS (SUR (J,I) - F (XOUT (J), YOUT (I)))
    4 0 ~ C O N T I N U E ~
99998 FORMAT (' ', 10X, 'X', 11X, 'Y', 9X, 'SURF', 6X, 'F(X,Y)', 7X,&
        'ERROR', /)
99999 FORMAT (1X, 5F12.4)
    END
```


## Output

| X | $Y$ | SURF | $F(X, Y)$ | ERROR |
| ---: | ---: | ---: | ---: | ---: |
| 0.0000 | 0.0000 | 3.0000 | 3.0000 | 0.0000 |
| 0.5000 | 0.0000 | 6.5000 | 6.5000 | 0.0000 |
| 1.0000 | 0.0000 | 10.0000 | 10.0000 | 0.0000 |
| 0.0000 | 0.5000 | 4.0000 | 4.0000 | 0.0000 |
| 0.5000 | 0.5000 | 7.5000 | 7.5000 | 0.0000 |
| 1.0000 | 0.5000 | 11.0000 | 11.0000 | 0.0000 |
| 0.0000 | 1.0000 | 5.0000 | 5.0000 | 0.0000 |
| 0.5000 | 1.0000 | 8.5000 | 8.5000 | 0.0000 |
| 1.0000 | 1.0000 | 12.0000 | 12.0000 | 0.0000 |

## RLINE

Fits a line to a set of data points using least squares.

## Required Arguments

XDATA - Vector of length nobs containing the $x$-values. (Input)

YDATA - Vector of length NOBS containing the $y$-values. (Input)
BO - Estimated intercept of the fitted line. (Output)
B1 - Estimated slope of the fitted line. (Output)

## Optional Arguments

NOBS - Number of observations. (Input)
Default: NOBS $=\operatorname{size}($ XDATA, 1$)$.
STAT - Vector of length 12 containing the statistics described below. (Output)
I ISTAT(I)
1 Mean of XDATA
2 Mean of YDATA
3 Sample variance of XDATA
4 Sample variance of YDATA
5 Correlation
6 Estimated standard error of B0
7 Estimated standard error of B1
8 Degrees of freedom for regression
9 Sum of squares for regression
10 Degrees of freedom for error
11 Sum of squares for error
12 Number of $(x, y)$ points containing NaN (not a number) as either the $x$ or $y$ value

## FORTRAN 90 Interface

Generic: CALL RLINE (XDATA, YDATA, B0, B1 [,...])
Specific: The specific interface names are S_RLINE and D_RLINE.

## FORTRAN 77 Interface

Single: CALL RLINE (NOBS, XDATA, YDATA, B0, B1, STAT)
Double: The double precision name is DRLINE.

## Description

Routine RLINE fits a line to a set of $(x, y)$ data points using the method of least squares. Draper and Smith (1981, pages 1-69) discuss the method. The fitted model is

$$
\hat{y}=\hat{\beta}_{0}+\hat{\beta}_{1} x
$$

where $\hat{\beta}_{0}$ (stored in B 0 ) is the estimated intercept and $\hat{\beta}_{1}$ (stored in B1) is the estimated slope. In addition to the fit, RLINE produces some summary statistics, including the means, sample variances, correlation, and the error (residual) sum of squares. The estimated standard errors of $\hat{\beta}_{0}$ and $\hat{\beta}_{1}$ are computed under the simple linear regression model. The errors in the model are assumed to be uncorrelated and with constant variance.

If the $x$ values are all equal, the model is degenerate. In this case, RLINE sets $\hat{\beta}_{1}$
to zero and $\hat{\beta}_{0}$ to the mean of the $y$ values.

## Comments

Informational error
Type Code
41 Each $(x, y)$ point contains NaN (not a number). There are no valid data.

## Example

This example fits a line to a set of data discussed by Draper and Smith (1981, Table 1.1, pages $9-33$ ). The response $y$ is the amount of steam used per month (in pounds), and the independent variable $x$ is the average atmospheric temperature (in degrees Fahrenheit).

```
USE RLINE_INT
USE UMACH_INT
USE WRRRL_INT
IMPLICIT NONE
INTEGER NOBS
PARAMETER (NOBS=25)
INTEGER NOUT
REAL B0, B1, STAT(12), XDATA(NOBS), YDATA(NOBS)
CHARACTER CLABEL(13)*15, RLABEL(1)*4
DATA XDATA/35.3, 29.7, 30.8, 58.8, 61.4, 71.3, 74.4, 76.7, 70.7.&
    57.5, 46.4, 28.9, 28.1, 39.1, 46.8, 48.5, 59.3, 70.0, 70.0.&
    74.5, 72.1, 58.1, 44.6, 33.4, 28.6/
DATA YDATA/10.98, 11.13, 12.51, 8.4, 9.27, 8.73, 6.36, 8.5,&
    7.82, 9.14, 8.24, 12.19, 11.88, 9.57, 10.94, 9.58, 10.09,&
    8.11, 6.83, 8.88, 7.68, 8.47, 8.86, 10.36, 11.08/
DATA RLABEL/'NONE'/, CLABEL/' ', 'Mean of X', 'Mean of Y',&
    'Variance X', 'Variance Y', 'Corr.', 'Std. Err. BO',&
    'Std. Err. B1', 'DF Reg.', 'SS Reg.', 'DF Error',&
    'SS Error', 'Pts. with NaN'/
CALL RLINE (XDATA, YDATA, B0, B1, STAT=STAT)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) B0, B1
99999 FORMAT (' B0 = ', F7.2, ' B1 = ', F9.5)
CALL WRRRL ('%/STAT', STAT, RLABEL, CLABEL, 1, 12, 1, &
    FMT = '(12W10.4)')
```

!
$!$

## Output

$\mathrm{B} 0=13.62 \mathrm{~B} 1=-0.07983$



Figure 3- 5 Plot of the Data and the Least Squares Line

## RCURV

Fits a polynomial curve using least squares.

## Required Arguments

XDATA - Vector of length NOBS containing the $x$ values. (Input)
YDATA - Vector of length NOBS containing the $y$ values. (Input)
$\boldsymbol{B}-$ Vector of length NDEG +1 containing the coefficients $\hat{\beta}$.
(Output)
The fitted polynomial is

$$
\hat{y}=\hat{\beta}_{0}+\hat{\beta}_{1} x+\hat{\beta}_{2} x^{2}+\cdots+\hat{\beta}_{k} x^{k}
$$

## Optional Arguments

NOBS - Number of observations. (Input)
Default: NOBS $=\operatorname{size}($ XDATA, 1$)$.
NDEG - Degree of polynomial. (Input)
Default: NDEG $=$ size $(B, 1)-1$.
SSPOLY - Vector of length NDEG +1 containing the sequential sums of squares. (Output) $\operatorname{SSPOLY}(1)$ contains the sum of squares due to the mean. For $i=1,2, \ldots, \operatorname{NDEG}$, $\operatorname{SSPOLY}(i+1)$ contains the sum of squares due to $x^{i}$ adjusted for the mean, $x, x^{2}, \ldots$, and $x^{i-1}$.

STAT - Vector of length 10 containing statistics described below. (Output)
i Statistics

1 Mean of $x$
2 Mean of $y$
3 Sample variance of $x$
4 Sample variance of $y$
5 R-squared (in percent)
6 Degrees of freedom for regression
$7 \quad$ Regression sum of squares
8 Degrees of freedom for error
9 Error sum of squares
10 Number of data points $(x, y)$ containing NaN (not a number) as a $x$ or $y$ value

## FORTRAN 90 Interface

Generic: CALL RCURV (XDATA, YDATA, B [, ...])

Specific: The specific interface names are S_RCURV and D_RCURV.

## FORTRAN 77 Interface

Single: CALL RCURV (NOBS, XDATA, YDATA, NDEG, B, SSPOLY, STAT)
Double: The double precision name is DRCURV.

## Description

Routine RCURV computes estimates of the regression coefficients in a polynomial (curvilinear) regression model. In addition to the computation of the fit, RCURV computes some summary statistics. Sequential sums of squares attributable to each power of the independent variable (stored in SSPOLY) are computed. These are useful in assessing the importance of the higher order powers in the fit. Draper and Smith (1981, pages 101-102) and Neter and Wasserman (1974, pages 278-287) discuss the interpretation of the sequential sums of squares. The statistic $R^{2}$ (stored in STAT (5)) is the percentage of the sum of squares of $y$ about its mean explained by the polynomial curve. Specifically,

$$
R^{2}=\frac{\sum_{i=1}^{n}\left(\hat{y}_{i}-\bar{y}\right)^{2}}{\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}} 100 \%
$$

where

$$
\hat{y}_{i}
$$

is the fitted $y$ value at $x_{i}$ and

$$
\bar{y}
$$

(stored in STAT (2) ) is the mean of $y$. This statistic is useful in assessing the overall fit of the curve to the data. $R^{2}$ must be between $0 \%$ and $100 \%$, inclusive. $R^{2}=100 \%$ indicates a perfect fit to the data.

Routine RCURV computes estimates of the regression coefficients in a polynomial model using orthogonal polynomials as the regressor variables. This reparameterization of the polynomial model in terms of orthogonal polynomials has the advantage that the loss of accuracy resulting from forming powers of the $x$-values is avoided. All results are returned to the user for the original model.

The routine RCURV is based on the algorithm of Forsythe (1957). A modification to Forsythe's algorithm suggested by Shampine (1975) is used for computing the polynomial coefficients. A discussion of Forsythe's algorithm and Shampine's modification appears in Kennedy and Gentle (1980, pages 342-347).

## Comments

1. Workspace may be explicitly provided, if desired, by use of R2URV/DR2URV. The reference is:
```
CALL R2URV (NOBS, XDATA, YDATA, NDEG, B, SSPOLY, STAT,WK, IWK)
```

The additional arguments are as follows:

```
WK - Work vector of length 11 * NOBS + 11 * NDEG + \(5+(\) NDEG + 1) *
    (NDEG + 3)
```

IWK - Work vector of length NOBS.
2. Informational errors
Type Code

43 Each $(x, y)$ point contains NaN (not a number). There are no valid data.
$4 \quad 7 \quad$ The $x$ values are constant. At least NDEG +1 distinct $x$ values are needed to fit a NDEG polynomial.
34 The $y$ values are constant. A zero order polynomial is fit. High order coefficients are set to zero.
35 There are too few observations to fit the desired degree polynomial. High order coefficients are set to zero.
36 A perfect fit was obtained with a polynomial of degree less than NDEG. High order coefficients are set to zero.
3. If NDEG is greater than 10 , the accuracy of the results may be questionable.

## Example

A polynomial model is fitted to data discussed by Neter and Wasserman (1974, pages 279-285). The data set contains the response variable $y$ measuring coffee sales (in hundred gallons) and the number of self-service coffee dispensers. Responses for fourteen similar cafeterias are in the data set.

```
USE RCURV_INT
USE WRRRL_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER NDEG, NOBS
PARAMETER (NDEG=2, NOBS=14)
REAL B(NDEG+1), SSPOLY(NDEG+1), STAT(10), XDATA(NOBS),&
    YDATA(NOBS)
CHARACTER CLABEL(11)*15, RLABEL(1)*4
DATA RLABEL/'NONE'/, CLABEL/' ', 'Mean of X', 'Mean of Y',&
            'Variance X', 'Variance Y', 'R-squared',&
            'DF Reg.', 'SS Reg.', 'DF Error', 'SS Error',&
            'Pts. with NaN'/
DATA XDATA/0., 0., 1., 1., 2., 2., 4., 4., 5., 5., 6., 6., 7..&
        7./
DATA YDATA/508.1, 498.4, 568.2, 577.3, 651.7, 657.0, 755.3,&
        758.9, 787.6, 792.1, 841.4, 831.8, 854.7, 871.4/
```

```
CALL RCURV (XDATA, YDATA, B, SSPOLY=SSPOLY, STAT=STAT)
CALL WRRRN ('B', B, 1, NDEG+1, 1)
CALL WRRRN ('SSPOLY', SSPOLY, 1, NDEG+1, 1)
CALL WRRRL ('%/STAT', STAT, RLABEL, CLABEL, 1, 10, 1, &
FMT='(2W10.4)')
END
```


## Output



Figure 3- 6 Plot of Data and Second Degree Polynomial Fit

## FNLSQ

Computes a least-squares approximation with user-supplied basis functions.

## Required Arguments

$\boldsymbol{F}$ - User-supplied function to evaluate basis functions. The form is $\mathrm{F}(\mathrm{K}, \mathrm{X})$, where

K - Number of the basis function. (Input)
K may be equal to $1,2, \ldots$, NBASIS.
X - Argument for evaluation of the $K$-th basis function. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program. The data FDATA is approximated by $A(1) * F(1, X)+A(2) * F(2, X)+\ldots+A($ NBASIS $) * F(N B A S I S, X)$ if INTCEP $=0$ and is approximated by $A(1)+A(2) * F(1, X)+\ldots+A($ NBASIS +1$) * F($ NBASIS, $X)$ if INTCEP $=1$.

XDATA - Array of length NDATA containing the abscissas of the data points. (Input)
FDATA - Array of length NDATA containing the ordinates of the data points. (Input)
A - Array of length INTCEP + NBAS IS containing the coefficients of the approximation. (Output)
If INTCEP $=1, A(1)$ contains the intercept. $A($ INTCEP + I) contains the coefficient of the I-th basis function.
$\boldsymbol{S S E}$ - Sum of squares of the errors. (Output)

## Optional Arguments

INTCEP - Intercept option. (Input)
Default: $\operatorname{INTCEP}=0$.

## INTCEP Action

$0 \quad$ No intercept is automatically included in the model.
1 An intercept is automatically included in the model.
NBASIS - Number of basis functions. (Input)
Default: NBASIS $=\operatorname{size}(\mathrm{A}, 1)$
NDATA - Number of data points. (Input)
Default: NDATA $=\operatorname{size}(X D A T A, 1)$.

IWT - Weighting option. (Input)
Default: $I W T=0$.

IWT Action
$0 \quad$ Weights of one are assumed.
1 Weights are supplied in WEIGHT.
WEIGHT - Array of length NDATA containing the weights. (Input if IWT = 1)
If IWT $=0$, WEIGHT is not referenced and may be dimensioned of length one.

## FORTRAN 90 Interface

Generic: CALL FNLSQ (F, XDATA, FDATA, A, SSE [, ...])
Specific: The specific interface names are S_FNLSQ and D_FNLSQ.

## FORTRAN 77 Interface

Single: CALL FNLSQ (F, INTCEP, NBASIS, NDATA, XDATA, FDATA, IWT, WEIGHT, A, SSE)

Double: The double precision name is DFNLSQ.

## Description

The routine FNLSQ computes a best least-squares approximation to given univariate data of the form

$$
\left\{\left(x_{i}, f_{i}\right)\right\}_{i=1}^{N}
$$

by $M$ basis functions

$$
\left\{F_{j}\right\}_{j=1}^{M}
$$

(where $M=$ NBASIS). In particular, if $\operatorname{INTCEP}=0$, this routine returns the error sum of squares SSE and the coefficients $a$ which minimize

$$
\sum_{i=1}^{N} w_{i}\left(f_{i}-\sum_{j=1}^{M} a_{j} F_{j}\left(x_{i}\right)\right)^{2}
$$

where $w=$ WEIGHT, $N=$ NDATA, $x=$ XDATA, and, $f=$ FDATA.
If $\operatorname{INTCEP}=1$, then an intercept is placed in the model; and the coefficients $a$, returned by FNLSQ, minimize the error sum of squares as indicated below.

$$
\sum_{i=1}^{N} w_{i}\left(f_{i}-a_{1}-\sum_{j=1}^{M} a_{j+1} F_{j}\left(x_{i}\right)\right)^{2}
$$

That is, the first element of the vector $a$ is now the coefficient of the function that is identically 1 and the coefficients of the $F_{j}$ 's are now $a_{j+l}$.

One additional parameter in the calling sequence for FNLSQ is IWT. If IWT is set to 0 , then $w_{i}=1$ is assumed. If IWT is set to 1 , then the user must supply the weights.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2LSQ/DF2LSQ. The reference is:
```
CALL F2LSQ (F, INTCEP, NBASIS, NDATA, XDATA, FDATA, IWT, WEIGHT, A, SSE,
WK)
```

The additional argument is
$\boldsymbol{W} \boldsymbol{K}$ - Work vector of length (INTCEP + NBASIS)**2 $+4 *($ INTCEP + NBASIS $)+$ IWT +1 . On output, the first (INTCEP + NBASIS)**2 elements of WK contain the $R$ matrix from a $Q R$ decomposition of the matrix containing a column of ones (if INTCEP $=1$ ) and the evaluated basis functions in columns INTCEP +1 through INTCEP + NBASIS.
2. Informational errors

Type Code
31 Linear dependence of the basis functions exists. One or more components of $A$ are set to zero.
32 Linear dependence of the constant function and basis functions exists. One or more components of $A$ are set to zero.
41 Negative weight encountered.

## Example

In this example, we fit the following two functions (indexed by $\delta$ )

$$
1+\sin x+7 \sin 3 x+\delta \varepsilon
$$

where $\varepsilon$ is random uniform deviate over the range $[-1,1]$, and $\delta$ is 0 for the first function and 1 for the second. These functions are evaluated at 90 equally spaced points on the interval [0, 6]. We use 4 basis functions, $\sin k x$ for $k=1, \ldots, 4$, with and without the intercept.

```
USE FNLSQ_INT
USE RNSET INT
USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
USE RNUNF_INT
IMPLICIT NONE
INTEGER NBASIS, NDATA
PARAMETER (NBASIS=4, NDATA=90)
```

```
!
    INTEGER I, INTCEP, NOUT
    REAL A(NBASIS+1), F, FDATA(NDATA), FLOAT, G, RNOISE,&
        SIN, SSE, X, XDATA(NDATA)
    INTRINSIC FLOAT, SIN
    EXTERNAL F
    G(X)=1.0 + SIN(X) + 7.0*SIN(3.0*X)
        Set random number seed
        Set up data values
    DO 10 I=1, NDATA
        XDATA(I) = 6.0*(FLOAT(I-1) /FLOAT (NDATA-1))
        FDATA(I) = G(XDATA(I))
    10 CONTINUE
    Compute least squares fit with no
    intercept
    CALL FNLSQ (F, XDATA, FDATA, A, SSE, INTCEP=INTCEP, &
        NBASIS=NBASIS)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99996)
    WRITE (NOUT,99999) SSE, (A(I),I=1,NBASIS)
    INTCEP = 1
        Compute least squares fit with
        intercept
    CALL FNLSQ (F, XDATA, FDATA, A, SSE, INTCEP=INTCEP, &
            NBASIS=NBASIS)
    WRITE (NOUT,99998) SSE, A(1), (A(I),I=2,NBASIS+1)
    DO 20 I=1, NDATA
        RNOISE = RNUNF()
        RNOISE = 2.0*RNOISE - 1.0
        FDATA(I) = FDATA(I) + RNOISE
    20 CONTINUE
    INTCEP = 0
        Compute least squares fit with no
        intercept
    CALL FNLSQ (F, XDATA, FDATA, A, SSE, INTCEP=INTCEP, &
        NBASIS=NBASIS)
    WRITE (NOUT,99997)
    WRITE (NOUT,99999) SSE, (A(I),I=1,NBASIS)
    INTCEP = 1
                                    Compute least squares fit with
                                    intercept
    CALL FNLSQ (F, XDATA, FDATA, A, SSE, INTCEP=INTCEP, &
        NBASIS=NBASIS )
```

```
! Write output
    WRITE (NOUT,99998) SSE, A(1), (A(I),I=2,NBASIS+1)
!
99996 FORMAT (//, ' Without error introduced we have :', /,&
99997 FORMAT (//, 'With error introduced we have :', /, ' SSE '&
, ' Intercept Coefficients ', /)
99998 FORMAT (1X, F8.4, 5X, F9.4, 5X, 4F9.4, /)
99999 FORMAT (1X, F8.4, 14X, 5X, 4F9.4, /)
    END
    REAL FUNCTION F (K, X)
    INTEGER K
    REAL X
!
    REAL SIN
    INTRINSIC SIN
!
    F = SIN (K*X)
    RETURN
    END
```


## Output

Without error introduced we have :
SSE Intercept Coefficients

| 89.8776 |  | 1.0101 | 0.0199 | 7.0291 | 0.0374 |
| ---: | :--- | :--- | :--- | :--- | :--- |
| 0.0000 | 1.0000 | 1.0000 | 0.0000 | 7.0000 | 0.0000 |

With error introduced we have :
SSE Intercept Coefficients

| 112.4662 |  | 0.9963 | -0.0675 | 6.9825 | 0.0133 |
| ---: | ---: | ---: | ---: | ---: | ---: |
| 30.9831 | 0.9522 | 0.9867 | -0.0864 | 6.9548 | -0.0223 |

## BSLSQ

Computes the least-squares spline approximation, and return the B-spline coefficients.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input)
FDATA - Array of length NDATA containing the data point ordinates. (Input)

KORDER - Order of the spline. (Input)
KORDER must be less than or equal to NDATA.
XKNOT - Array of length NCOEF + KORDER containing the knot sequence. (Input) XKNOT must be nondecreasing.

NCOEF - Number of B-spline coefficients. (Input)
NCOEF cannot be greater than NDATA.
BSCOEF - Array of length ncoef containing the B-spline coefficients. (Output)

## Optional Arguments

NDATA - Number of data points. (Input)
Default: NDATA $=\operatorname{size}(X D A T A, 1)$
WEIGHT - Array of length NDATA containing the weights. (Input)
Default: WE IGHT $=1.0$.

## FORTRAN 90 Interface

Generic: CALL BSLSQ (XDATA, FDATA, KORDER, XKNOT, NCOEF, BSCOEF
$[, \ldots])$

Specific: The specific interface names are _ $^{\text {BSLSQ }}$ and D_BSLSQ.

## FORTRAN 77 Interface

Single: CALL BSLSQ (NDATA, XDATA, FDATA, WEIGHT, KORDER, XKNOT, NCOEF, BSCOEF)

Double: The double precision name is DBSLSQ.

## Description

The routine BSLSQ is based on the routine L2APPR by de Boor (1978, page 255). The IMSL routine BSLSQ computes a weighted discrete $L_{2}$ approximation from a spline subspace to a given data set $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots, N$ (where $N=$ NDATA $)$. In other words, it finds B-spline coefficients, $a=\mathrm{BSCOEF}$, such that

$$
\sum_{i=1}^{N}\left|f_{i}-\sum_{j=1}^{m} a_{j} B_{j}\left(x_{i}\right)\right|^{2} w_{i}
$$

is a minimum, where $m=$ NCOEF and $B_{j}$ denotes the $j$-th B-spline for the given order, KORDER, and knot sequence, XKNOT. This linear least squares problem is solved by computing and solving the normal equations. While the normal equations can sometimes cause numerical difficulties, their use here should not cause a problem because the B-spline basis generally leads to well-conditioned banded matrices.

The choice of weights depends on the problem. In some cases, there is a natural choice for the weights based on the relative importance of the data points. To approximate a continuous function (if the location of the data points can be chosen), then the use of Gauss quadrature weights and points is reasonable. This follows because BSLSQ is minimizing an approximation to the integral

$$
\int|F-s|^{2} d x
$$

The Gauss quadrature weights and points can be obtained using the IMSL routine GQRUL (see Chapter 4, Integration and Differentiation).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2LSQ/DB2LSQ. The reference is:

CALL B2LSQ (NDATA, XDATA, FDATA, WEIGHT, KORDER, XKNOT, NCOEF, BSCOEF, WK1, WK2, WK3, WK4, IWK)

The additional arguments are as follows:

$$
\begin{aligned}
& \text { WK1 - Work array of length (3 + NCOEF) * KORDER. } \\
& \text { WK2 - Work array of length NDATA. } \\
& \text { WK3 - Work array of length NDATA. } \\
& \text { WK4 — Work array of length NDATA. } \\
& \boldsymbol{I W K} \text { — Work array of length NDATA. }
\end{aligned}
$$

2. Informational errors

Type Code
45 Multiplicity of the knots cannot exceed the order of the spline.
46 The knots must be nondecreasing.
$4 \quad 7 \quad$ All weights must be greater than zero.
48 The smallest element of the data point array must be greater than or equal to the KORDth knot.
49 The largest element of the data point array must be less than or equal to the (NCOEF +1 )st knot.
3. The B-spline representation can be evaluated using BSVAL, and its derivative can be evaluated using BSDER.

## Example

In this example, we try to recover a quadratic polynomial using a quadratic spline with one interior knot from two different data sets. The first data set is generated by evaluating the quadratic at 50 equally spaced points in the interval $(0,1)$ and then adding uniformly distributed noise to the data. The second data set includes the first data set, and, additionally, the values at 0 and at 1 with no noise added. Since the first and last data points are uncontaminated by noise, we have chosen weights equal to $10^{5}$ for these two points in this second problem. The quadratic, the first approximation, and the second approximation are then evaluated at 11 equally spaced points. This example illustrates the use of the weights to enforce interpolation at certain of the data points.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER KORDER, NCOEF
    PARAMETER (KORDER=3, NCOEF=4)
!
    INTEGER I, NDATA, NOUT
    REAL ABS, BSCOF1(NCOEF), BSCOF2 (NCOEF), F,&
            FDATA1(50), FDATA2(52), FLOAT, RNOISE, S1,&
            S2, WEIGHT (52), X, XDATA1(50), XDATA2(52), &
            XKNOT(KORDER+NCOEF), XT, YT
    INTRINSIC ABS, FLOAT
!
! Define function
    F(X) = 8.0*X* (1.0-X)
    CALL RNSET (12345679)
    NDATA = 50
    DO 10 I=1, NCOEF - KORDER + 2
        XKNOT(I+KORDER-1) = FLOAT(I-1)/FLOAT (NCOEF-KORDER+1)
    CONTINUE
    DO 20 I=1, KORDER - 1
        XKNOT(I) = XKNOT (KORDER)
        XKNOT(I+NCOEF+1) = XKNOT (NCOEF+1)
    20 CONTINUE
                                    Set up data points excluding
                    the endpoints 0 and 1.
                            The function values have noise
                            introduced.
    DO 30 I=1, NDATA
        XDATA1(I) = FLOAT(I)/51.0
        RNOISE = RNUNF()
        RNOISE = RNOISE - 0.5
        FDATA1(I) = F(XDATA1(I)) + RNOISE
    3 0 ~ C O N T I N U E
                                    Compute least squares B-spline
                                    representation.
    CALL BSLSQ (XDATA1, FDATA1, KORDER, XKNOT, NCOEF, BSCOF1)
                                    Now use same XDATA values but with
                                    the endpoints included. These
                                    points will have large weights.
    NDATA = 52
    CALL SCOPY (50, XDATA1, 1, XDATA2 (2:), 1)
    CALL SCOPY (50, FDATA1, 1, FDATA2 (2:), 1)
    WEIGHT(1) = 1.0E5
    XDATA2(1) = 0.0
    FDATA2(1) = F(XDATA2(1))
    WEIGHT (NDATA) = 1.0E5
    XDATA2 (NDATA) = 1.0
    FDATA2 (NDATA) = F (XDATA2 (NDATA))
    Compute least squares B-spline
```

```
! representation.
    CALL BSLSQ (XDATA2, FDATA2, KORDER, XKNOT, NCOEF, BSCOF2, &
            WEIGHT=WEIGHT)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
! Print the two interpolants
    DO 40 I=1, 11
    XT = FLOAT (I-1)/10.0
    YT = F(XT)
! Evaluate splines
    S1 = BSVAL (XT,KORDER,XKNOT,NCOEF,BSCOF1)
    S2 = BSVAL (XT,KORDER,XKNOT,NCOEF,BSCOF2)
    WRITE (NOUT,99999) XT, YT, S1, S2, (S1-YT), (S2-YT)
    40 CONTINUE
!
99998 FORMAT (7X, 'X', 9X, 'F(X)', 6X, 'S1(X)', 5X, 'S2(X)', 7X,&
            'F(X)-S1(X)', 7X, 'F(X)-S2(X)')
99999 FORMAT (' ', 4F10.4, 4X, F10.4, 7X, F10.4)
    END
```


## Output

| $X$ | $F(X)$ | $S 1(X)$ | $S 2(X)$ | $F(X)-S 1(X)$ | $F(X)-S 2(X)$ |
| :--- | :--- | :---: | :---: | :---: | :---: |
| 0.0000 | 0.0000 | 0.0515 | 0.0000 | 0.0515 | 0.0000 |
| 0.1000 | 0.7200 | 0.7594 | 0.7490 | 0.0394 | 0.0290 |
| 0.2000 | 1.2800 | 1.3142 | 1.3277 | 0.0342 | 0.0477 |
| 0.3000 | 1.6800 | 1.7158 | 1.7362 | 0.0358 | 0.0562 |
| 0.4000 | 1.9200 | 1.9641 | 1.9744 | 0.0441 | 0.0544 |
| 0.5000 | 2.0000 | 2.0593 | 2.0423 | 0.0593 | 0.0423 |
| 0.6000 | 1.9200 | 1.9842 | 1.9468 | 0.0642 | 0.0268 |
| 0.7000 | 1.6800 | 1.7220 | 1.6948 | 0.0420 | 0.0148 |
| 0.8000 | 1.2800 | 1.2726 | 1.2863 | -0.0074 | 0.0063 |
| 0.9000 | 0.7200 | 0.6360 | 0.7214 | -0.0840 | 0.0014 |
| 1.0000 | 0.0000 | -0.1878 | 0.0000 | -0.1878 | 0.0000 |

## BSVLS

Computes the variable knot B-spline least squares approximation to given data.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input)
FDATA - Array of length NDATA containing the data point ordinates. (Input)
KORDER - Order of the spline. (Input)
KORDER must be less than or equal to NDATA.

NCOEF - Number of B-spline coefficients. (Input)
nCOEF must be less than or equal to NDATA.
XGUESS - Array of length NCOEF + KORDER containing the initial guess of knots. (Input) XGUESS must be nondecreasing.

XKNOT - Array of length NCOEF + KORDER containing the (nondecreasing) knot sequence. (Output)

BSCOEF - Array of length NCOEF containing the B-spline representation. (Output)
SSQ - The square root of the sum of the squares of the error. (Output)

## Optonal Arguments

NDATA - Number of data points. (Input)
NDATA must be at least 2 .
Default: NDATA $=\operatorname{size}(X D A T A, 1)$
WEIGHT - Array of length NDATA containing the weights. (Input)
Default: WEIGHT $=1.0$.

## FORTRAN 90 Interface

Generic: CALL BSVLS (NDATA, XDATA, FDATA, WEIGHT, KORDER, NCOEF, XGUESS, XKNOT, BSCOEF, SSQ)

Specific: The specific interface names are _ BSVLS $^{\text {and D_BSVLS. }}$

## FORTRAN 77 Interface

Single: CALL BSVLS (XDATA, FDATA, KORDER, NCOEF, XGUESS, XKNOT, BSCOEF, SSQ [,...])

Double: The double precision name is DBSVLS.

## Description

The routine BSVLS attempts to find the best placement of knots that will minimize the leastsquares error to given data by a spline of order $k=$ KORDER with $N=$ NCOEF coefficients. The user provides the order $k$ of the spline and the number of coefficients $N$. For this problem to make sense, it is necessary that $N>k$. We then attempt to find the minimum of the functional

$$
F(a, \mathbf{t})=\sum_{i=1}^{M} w_{i}\left(f_{i}-\sum_{j=1}^{N} a_{j} B_{j, k, \mathrm{t}}\left(x_{j}\right)\right)^{2}
$$

The user must provide the weights $w=$ WEIGHT, the data $x_{i}=$ XDATA and $f_{i}$ FFDATA, and $M=$ nDATA. The minimum is taken over all admissible knot sequences $\mathbf{t}$.

The technique employed in BSVLS uses the fact that for a fixed knot sequence $\mathbf{t}$ the minimization in $a$ is a linear least-squares problem that can be solved by calling the IMSL routine BSLSQ. Thus, we can think of our objective function $F$ as a function of just $\mathbf{t}$ by setting

$$
G(\mathbf{t})=\min _{a} F(a, \mathbf{t})
$$

A Gauss-Seidel (cyclic coordinate) method is then used to reduce the value of the new objective function $G$. In addition to this local method, there is a global heuristic built into the algorithm that will be useful if the data arise from a smooth function. This heuristic is based on the routine NEWNOT of de Boor (1978, pages 184 and 258-261).

The user must input an initial guess, $\mathbf{t}^{g}=$ XGUESS, for the knot sequence. This guess must be a valid knot sequence for the splines of order $k$ with

$$
\mathbf{t}_{1}^{g} \leq \ldots \leq \mathbf{t}_{k}^{g} \leq x_{i} \leq \mathbf{t}_{N+1}^{g} \leq \ldots \leq \mathbf{t}_{N+k}^{g}, \quad i=1, \ldots, M
$$

with $\mathbf{t}^{g}$ nondecreasing, and

$$
\mathbf{t}_{i}^{g}<\mathbf{t}_{i+k}^{g} \quad i=1, \ldots, N
$$

The routine BSVLS returns the B-spline representation of the best fit found by the algorithm as well as the square root of the sum of squares error in SSQ. If this answer is unsatisfactory, you may reinitialize BSVLS with the return from BSVLS to see if an improvement will occur. We have found that this option does not usually (substantially) improve the result. In regard to execution speed, this routine can be several orders of magnitude slower than one call to the least-squares routine BSLSQ.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2VLS / DB2VLS. The reference is:

CALL B2VLS (NDATA, XDATA, FDATA, WEIGHT, KORDER, NCOEF, XGUESS, XKNOT, BSCOEF, SSQ, IWK, WK)

The additional arguments are as follows:
IWK - Work array of length NDATA.
WK - Work array of length NCOEF * $(6+2 *$ KORDER $)+$ KORDER * $(7-\operatorname{KORDER})+$ $3 *$ NDATA +3 .
2. Informational errors

Type Code
312 The knots found to be optimal are stacked more than KORDER. This indicates fewer knots will produce the same error sum of squares. The knots have been separated slightly.

9 The multiplicity of the knots in XGUESS cannot exceed the order of the spline.
$4 \quad 10 \quad$ XGUESS must be nondecreasing.

## Example

In this example, we try to fit the function $|x-.33|$ evaluated at 100 equally spaced points on $[0,1]$. We first use quadratic splines with 2 interior knots initially at .2 and .8 . The eventual error should be zero since the function is a quadratic spline with two knots stacked at .33 . As a second example, we try to fit the same data with cubic splines with three interior knots initially located at $.1, .2$, and, .5 . Again, the theoretical error is zero when the three knots are stacked at .33 .

We include a graph of the initial least-squares fit using the IMSL routine BSLSQ for the above quadratic spline example with knots at .2 and .8 . This graph overlays the graph of the spline computed by BSVLS, which is indistinguishable from the data.

```
    USE BSVLS INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER KORD1, KORD2, NCOEF1, NCOEF2, NDATA
    PARAMETER (KORD1=3, KORD2=4, NCOEF1=5, NCOEF2=7, NDATA=100)
    INTEGER I, NOUT
    REAL ABS, BSCOEF(NCOEF2), F, FDATA(NDATA), FLOAT, SSQ,&
            WEIGHT (NDATA), X, XDATA (NDATA), XGUES1 (NCOEF1+KORD1) , &
            XGUES2 (KORD2+NCOEF2), XKNOT (NCOEF2+KORD2)
    INTRINSIC ABS, FLOAT
    DATA XGUES1/3*0.0, .2, . 8, 3*1.0001/
    DATA XGUES2/4*0.0, .1, .2, .5, 4*1.0001/
    DATA WEIGHT/NDATA*.01/
! Define function
    F(X) = ABS (X-. 33)
    DO 10 I=1, NDATA
        XDATA(I) = FLOAT(I-1)/FLOAT (NDATA)
        FDATA(I) = F(XDATA(I))
    10 CONTINUE
Compute least squares B-spline
                    representation with KORD1, NCOEF1,
                    and XGUES1.
CALL BSVLS (XDATA, FDATA, KORD1, NCOEF1, XGUES1, &
        XKNOT, BSCOEF, SSQ, WEIGHT=WEIGHT)
                            Get output unit number
CALL UMACH (2, NOUT)
                            Print heading
WRITE (NOUT,99998) 'quadratic'
                            Print SSQ and the knots
WRITE (NOUT,99999) SSQ, (XKNOT(I),I=1,KORD1+NCOEF1)
Compute least squares B-spline
                    representation with KORD2, NCOEF2,
                    and XGUES2.
CALL BSVLS (XDATA, FDATA, KORD2, NCOEF2, XGUES2,&
```

!
!

```
            XKNOT, BSCOEF, SSQ, WEIGHT=WEIGHT)
! Print SSQ and the knots
    WRITE (NOUT,99998) 'cubic'
    WRITE (NOUT,99999) SSQ, (XKNOT(I),I=1,KORD2+NCOEF2)
!
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ P i e c e w i s e ~ ' , ~ A , ~ / ) ~
99999 FORMAT (' Square root of the sum of squares : ', F9.4, /,&
    ' Knot sequence : ', /, 1X, 11(F9.4,/,1X))
    END
```


## Output

Piecewise quadratic

Square root of the sum of squares : 0.0008
Knot sequence :
0.0000
0.0000
0.0000
0.3137
0.3464
1.0001
1.0001
1.0001

Piecewise cubic

Square root of the sum of squares : 0.0005
Knot sequence :
0.0000
0.0000
0.0000
0.0000
0.3167
0.3273
0.3464
1.0001
1.0001
1.0001
1.0001


Figure 3- 7 BSVLS vs. BSLSQ

## CONFT

Computes the least-squares constrained spline approximation, returning the B-spline coefficients.

## Required Arguments

XDATA - Array of length NDATA containing the data point abscissas. (Input)
FDATA - Array of size NDATA containing the values to be approximated. (Input) FDATA(I) contains the value at XDATA(I).

XVAL - Array of length NXVAL containing the abscissas at which the fit is to be constrained. (Input)

NHARD - Number of entries of XVAL involved in the 'hard' constraints. (Input)
Note: $(0 \leq$ NHARD $\leq N X V A L)$. Setting NHARD to zero always results in a fit, while setting NHARD to NXVAL forces all constraints to be met. The 'hard' constraints must be satisfied or else the routine signals failure. The 'soft' constraints need not be satisfied, but there will be an attempt to satisfy the 'soft' constraints. The constraints must be ordered in terms of priority with the most important constraints first. Thus, all of the 'hard' constraints must preceed the 'soft' constraints. If infeasibility is detected among the soft constraints, we satisfy (in order) as many of the soft constraints as possible.

IDER - Array of length NXVAL containing the derivative value of the spline that is to be constrained. (Input)

If we want to constrain the integral of the spline over the closed interval $(c, d)$, then we set $\operatorname{IDER}(\mathrm{I})=\operatorname{IDER}(\mathrm{I}+1)=-1$ and $\operatorname{XVAL}(\mathrm{I})=c$ and $\operatorname{XVAL}(I+1)=d$. For consistency, we insist that $\operatorname{ITYPE}(\mathrm{I})=\operatorname{ITYPE}(\mathrm{I}+1)$.GE. 0 and $c . \operatorname{LE} . d$. Note that every entry in IDER must be at least -1 .

ITYPE - Array of length NXVAL indicating the types of general constraints. (Input)

| ITYPE(I) | I-th Constraint |
| ---: | :--- |
| 1 | $\mathrm{BL}(\mathrm{I})=f^{\left(d_{i}\right)}\left(x_{i}\right)$ |
| 2 | $f^{\left(d_{i}\right)}\left(x_{i}\right) \leq \mathrm{BU}(\mathrm{I})$ |
| 3 | $f^{\left(d_{i}\right)}\left(x_{i}\right) \geq \mathrm{BL}(\mathrm{I})$ |
| 4 | $\mathrm{BL}(\mathrm{I})=\leq f^{\left(d_{i}\right)}\left(x_{i}\right) \leq \mathrm{BU}(\mathrm{I})$ |
| $\left(d_{i}=-1\right) 1$ | $\mathrm{BL}(\mathrm{I})=\int_{c}^{d} f(t) d t$ |
| $\left(d_{i}=-1\right) 2$ | $\int_{c}^{d} f(t) d t \leq \mathrm{BU}(\mathrm{I})$ |
| $\left(d_{i}=-1\right) 3$ | $\int_{c}^{d} f(t) d t \geq \mathrm{BL}(\mathrm{I})$ |
| $\left(d_{i}=-1\right) 4$ | $\mathrm{BL}(\mathrm{I}) \leq \int_{c}^{d} f(t) d t \leq \mathrm{BU}(\mathrm{I})$ |
| 10 | periodic end conditions |
| 99 | disregard this constraint |

In order to set two point constraints, we must have $\operatorname{ITYPE}(I)=\operatorname{ITYPE}(I+1)$ and $\operatorname{ITYPE}(\mathrm{I})$ must be negative.

| ITYPE (I) | I - th Contraint |
| ---: | :--- |
| -1 | $\mathrm{BL}(\mathrm{I})=f^{\left(d_{i}\right)}\left(x_{i}\right)-f^{\left(d_{i+1}\right)}\left(x_{i+1}\right)$ |
| -2 | $f^{\left(d_{i}\right)}\left(x_{i}\right)-f^{\left(d_{i+1}\right)}\left(x_{i+1}\right) \leq \mathrm{BU}(\mathrm{I})$ |
| -3 | $f^{\left(d_{i}\right)}\left(x_{i}\right)-f^{\left(d_{i+1}\right)}\left(x_{i+1}\right) \geq \mathrm{BL}(\mathrm{I})$ |
| -4 | $\mathrm{BL}(\mathrm{I}) \leq f^{\left(d_{i}\right)}\left(x_{i}\right)-f^{\left(d_{i+1}\right)}\left(x_{i+1}\right) \leq \mathrm{BU}(\mathrm{I})$ |

$\boldsymbol{B L}$ - Array of length NXVAL containing the lower limit of the general constraints, if there is no lower limit on the I-th constraint, then $\mathrm{BL}(\mathrm{I})$ is not referenced. (Input)
$\boldsymbol{B U}$ - Array of length NXVAL containing the upper limit of the general constraints, if there is no upper limit on the I-th constraint, then $\operatorname{BU}(I)$ is not referenced; if there is no range constraint, BL and BU can share the same storage locations. (Input)
If the I-th constraint is an equality constraint, $B U(I)$ is not referenced.
KORDER - Order of the spline. (Input)

XKNOT - Array of length NCOEF + KORDER containing the knot sequence. (Input) The entries of XKNOT must be nondecreasing.

BSCOEF - Array of length NCOEF containing the B-spline coefficients. (Output)

## Optional Arguments

NDATA - Number of data points. (Input) Default: NDATA $=$ size $(X D A T A, 1)$.

WEIGHT - Array of length NDATA containing the weights. (Input)
Default: $\mathrm{WEIGHT}=1.0$.
NXVAL - Number of points in the vector XVAL. (Input) Default: NXVAL $=\operatorname{size}($ XVAL, 1$)$.

NCOEF - Number of B-spline coefficients. (Input)
Default: $\mathrm{NCOEF}=\operatorname{size}(\mathrm{BSCOEF}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CONFT (XDATA, FDATA, XVAL, NHARD, IDER, ITYPE, BL, BU, KORDER, XKNOT, BSCOEF [,...])

Specific: The specific interface names are S_CONFT and D_CONFT.

## FORTRAN 77 Interface

Single: CALL CONFT (NDATA, XDATA, FDATA, WEIGHT, NXVAL, XVAL, NHARD, IDER, ITYPE, BL, BU, KORDER, XKNOT, NCOEF, BSCOEF)

Double: $\quad$ The double precision name is DCONFT.

## Description

The routine CONFT produces a constrained, weighted least-squares fit to data from a spline subspace. Constraints involving one point, two points, or integrals over an interval are allowed. The types of constraints supported by the routine are of four types.

$$
\begin{aligned}
E_{p}[f] & & =f^{\left(j_{p}\right)}\left(y_{p}\right) \\
\text { or } & & =f^{\left(j_{p}\right)}\left(y_{p}\right)-f^{\left(j_{p+1}\right)}\left(y_{p+1}\right) \\
\text { or } & & =\int_{y_{p}}^{y_{p+1}} f(t) d t \\
\text { or } & & =\text { periodic end conditions }
\end{aligned}
$$

An interval, $I_{p}$, (which may be a point, a finite interval, or semi-infinite interval) is associated with each of these constraints.

The input for this routine consists of several items, first, the data set $\left(x_{i}, f_{i}\right)$ for $i=1, \ldots, N$ (where $N=$ NDATA ), that is the data which is to be fit. Second, we have the weights to be used in the least squares fit ( $w=$ WEIGHT $)$. The vector XVAL of length NXVAL contains the abscissas of the points involved in specifying the constraints. The algorithm tries to satisfy all the constraints, but if the constraints are inconsistent then it will drop constraints, in the reverse order specified, until either a consistent set of constraints is found or the "hard" constraints are determined to be inconsistent (the "hard" constraints are those involving XVAL(1), .., XVAL(NHARD)). Thus, the algorithm satisfies as many constraints as possible in the order specified by the user. In the case when constraints are dropped, the user will receive a message explaining how many constraints had to be dropped to obtain the fit. The next several arguments are related to the type of constraint and the constraint interval. The last four arguments determine the spline solution. The user chooses the spline subspace (KORDER, XKNOT, and NCOEF), and the routine returns the B-spline coefficients in BSCOEF.

Let $n_{f}$ denote the number of feasible constraints as described above. Then, the routine solves the problem.

$$
\begin{aligned}
& \sum_{i=1}^{N}\left|f_{i}-\sum_{j=1}^{m} a_{j} B_{j}\left(x_{i}\right)\right|^{2} w_{i} \\
& E_{p}\left[\sum_{j=1}^{m} a_{j} B_{j}\right] \in I_{p} \quad p=1, \ldots, n_{f}
\end{aligned}
$$

This linearly constrained least-squares problem is treated as a quadratic program and is solved by invoking the IMSL routine QPROG (see Chapter 8, Optimization).
The choice of weights depends on the data uncertainty in the problem. In some cases, there is a natural choice for the weights based on the estimates of errors in the data points.
Determining feasibility of linear constraints is a numerically sensitive task. If you encounter difficulties, a quick fix would be to widen the constraint intervals $I_{p}$.

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2NFT/DC2NFT. The reference is:

CALL C2NFT (NDATA, XDATA, FDATA, WEIGHT, NXVAL, XVAL, NHARD, IDER, ITYPE, BL, BU, KORDER, XKNOT, NCOEF, BSCOEF, H, G, A, RHS, WK, IPERM, IWK)

The additional arguments are as follows:
$\boldsymbol{H}$ - Work array of size ncoef by ncoef. Upon output, h contains the Hessian matrix of the objective function used in the call to $Q P R O G$ (see Chapter 8 , Optimization).
$\boldsymbol{G}$ - Work array of size NCOEF. Upon output, G contains the coefficients of the linear term used in the call to QPROG.
$\boldsymbol{A}-$ Work array of size ( 2 * NXVAL + KORDER) by (NCOEF + 1). Upon output, A contains the constraint matrix used in the call QPROG. The last column of A is used to keep record of the original order of the constraints.
$\boldsymbol{R H S}$ - Work array of size 2 * NXVAL + KORDER. Upon output, RHS contains the right hand side of the constraint matrix $A$ used in the call to $Q P R O G$.
$\boldsymbol{W} \boldsymbol{K}$ - Work array of size $($ KORDER +1$)$ * $(2$ * KORDER +1$)+(3$ * NCOEF * NCOEF + 13 * NCOEF $) / 2+(2 *$ NXVAL + KORDER +30$) \star(2 \star$ NXVAL + KORDER $)+$ NDATA + 1.

IPERM - Work array of size NXVAL. Upon output, IPERM contains the permutaion of the original constraints used to generate the matrix $A$.
$\boldsymbol{I} \boldsymbol{W} \boldsymbol{K}$ - Work array of size NDATA +30 * $(2 *$ NXVAL + KORDER $)+4$ * NCOEF.
2. Informational errors

Type Code

| 3 | 11 | Soft constraints had to be removed in order to get a fit. |
| :--- | :--- | :--- |
| 4 | 12 | Multiplicity of the knots cannot exceed the order of the spline. |
| 4 | 13 | The knots must be nondecreasing. |
| 4 | 14 | The smallest element of the data point array must be greater than or <br> equal to the KORD-th knot. |
| 4 | 15 | The largest element of the data point array must be less than or equal <br> to the (NCOEF + 1)st knot. |
| 4 | 16 | All weights must be greater than zero. |
| 4 | 17 | The hard constraints could not be met. |
| 4 | 18 | The abscissas of the constrained points must lie within knot interval. <br> 4 |
| 4 | 20 | The upperbound must be greater than or equal to the lowerbound for <br> a range constaint. <br> The upper limit of integration must be greater than the lower limit of <br> integration for constraints involving the integral of the <br> approximation. |

## Example 1

This is a simple application of CONFT. We generate data from the function

$$
\frac{x}{2}+\sin \left(\frac{x}{2}\right)
$$

contaminated with random noise and fit it with cubic splines. The function is increasing so we would hope that our least-squares fit would also be increasing. This is not the case for the unconstrained least squares fit generated by BSLSQ. We then force the derivative to be greater than 0 at NXVAL $=15$ equally spaced points and call CONFT. The resulting curve is monotone. We print the error for the two fits averaged over 100 equally spaced points.

```
USE IMSL_LIBRARIES
```

```
    IMPLICIT NONE
    INTEGER KORDER, NCOEF, NDATA, NXVAL
    PARAMETER (KORDER=4, NCOEF=8, NDATA=15, NXVAL=15)
!
    INTEGER I, IDER(NXVAL), ITYPE(NXVAL), NHARD, NOUT
    REAL ABS, BL (NXVAL), BSCLSQ (NDATA), BSCNFT (NDATA), &
        BU(NXVAL), ERRLSQ, ERRNFT, F1, FDATA(NDATA), FLOAT,&
        GRDSIZ, SIN, WEIGHT(NDATA), X, XDATA(NDATA), &
        XKNOT (KORDER+NDATA), XVAL (NXVAL)
    INTRINSIC ABS, FLOAT, SIN
    F1 (X) = . 5*X + SIN(.5*X)
                Initialize random number generator
                and get output unit number.
    CALL RNSET (234579)
    CALL UMACH (2, NOUT)
                Use default weights of one.
                                    Compute original XDATA and FDATA
                                    with random noise.
    GRDSIZ = 10.0
    DO 10 I=1, NDATA
        XDATA(I) = GRDSIZ*((FLOAT (I-1)/FLOAT (NDATA-1)))
        FDATA(I) = RNUNF()
        FDATA(I) = F1(XDATA(I)) + (FDATA(I) -.5)
    CONTINUE
                                    Compute knots
    DO 20 I=1, NCOEF - KORDER + 2
        XKNOT(I+KORDER-1) = GRDSIZ*((FLOAT (I-1)/FLOAT (NCOEF-KORDER+1)) &
                            )
    20 CONTINUE
    DO 30 I=1, KORDER - 1
        XKNOT(I) = XKNOT (KORDER)
        XKNOT(I+NCOEF+1) = XKNOT (NCOEF+1)
    CONTINUE
            Compute BSLSQ fit.
    CALL BSLSQ (XDATA, FDATA, KORDER, XKNOT, NCOEF, BSCLSQ)
                                    Construct the constraints for
            CONFT.
    DO 40 I=1, NXVAL
        XVAL(I) = GRDSIZ*FLOAT(I-1)/FLOAT (NXVAL-1)
        ITYPE(I) = 3
        IDER(I) = 1
        BL(I) = 0.0
    4 0 ~ C O N T I N U E
CALL CONFT (XDATA, FDATA, XVAL, NHARD, IDER, ITYPE, BL, BU, KORDER, &
                XKNOT, BSCNFT, NCOEF=NCOEF)
                    Compute the average error
                    of }100\mathrm{ points in the interval.
ERRLSQ = 0.0
ERRNFT = 0.0
DO 50 I=1, 100
```

```
                X = GRDSIZ*FLOAT (I-1)/99.0
                ERRNFT = ERRNFT + ABS(F1 (X) -BSVAL(X,KORDER,XKNOT,NCOEF,BSCNFT) &
                )
                ERRLSQ = ERRLSQ + ABS(F1 (X) -BSVAL(X,KORDER,XKNOT,NCOEF,BSCLSQ)&
    5 0 ~ C O N T I N U E
!
WRITE (NOUT,99998) ERRLSQ/100.0
WRITE (NOUT,99999) ERRNFT/100.0
!
99998 FORMAT (' Average error with BSLSQ fit: ', F8.5)
99999 FORMAT (' Average error with CONFT fit: ', F8.5)
END
```


## Output

Average error with BSLSQ fit: 0.20250
Average error with CONFT fit: 0.14334


Figure 3- 8 CONFT vs. BSLSQ Forcing Monotonicity

## Additional Examples

## Example 2

We now try to recover the function

$$
\frac{1}{1+x^{4}}
$$

from noisy data. We first try the unconstrained least-squares fit using BSLSQ. Finding that fit somewhat unsatisfactory, we apply several constraints using CONFT. First, notice that the unconstrained fit oscillates through the true function at both ends of the interval. This is common for flat data. To remove this oscillation, we constrain the cubic spline to have zero second derivative at the first and last four knots. This forces the cubic spline to reduce to a linear polynomial on the first and last three knot intervals. In addition, we constrain the fit (which we will call $s$ ) as follows:

$$
\begin{aligned}
s(-7) & \geq 0 \\
\int_{-7}^{7} s(x) d x & \leq 2.3 \\
s(-7) & =s(7)
\end{aligned}
$$

Notice that the last constraint was generated using the periodic option (requiring only the zeroeth derivative to be periodic). We print the error for the two fits averaged over 100 equally spaced points.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER KORDER, NCOEF, NDATA, NXVAL
    PARAMETER (KORDER=4, NCOEF=13, NDATA=51, NXVAL=12)
    INTEGER I, IDER(NXVAL), ITYPE(NXVAL), NHARPT, NOUT
    REAL ABS, BL(NXVAL), BSCLSQ(NDATA), BSCNFT (NDATA),&
        BU(NXVAL), ERRLSQ, ERRNFT, F1, FDATA(NDATA), FLOAT,&
        GRDSIZ, WEIGHT (NDATA), X, XDATA(NDATA),&
        XKNOT (KORDER+NDATA), XVAL (NXVAL)
    INTRINSIC ABS, FLOAT
    F1(X) = 1.0/(1.0+X**4)
! Initialize random number generator
                and get output unit number.
    CALL UMACH (2, NOUT)
    CALL RNSET (234579)
        Use deafult weights of one.
        Compute original XDATA and FDATA
        with random noise.
    GRDSIZ = 14.0
    DO 10 I=1, NDATA
        XDATA(I) = GRDSIZ*((FLOAT(I-1)/FLOAT (NDATA-1))) - GRDSIZ/2.0
        FDATA(I) = RNUNF()
        FDATA(I) = F1(XDATA(I)) + 0.125*(FDATA(I)-.5)
    10 CONTINUE
    DO 20 I=1, NCOEF - KORDER + 2
    XKNOT(I+KORDER-1) = GRDSIZ*((FLOAT(I-1)/FLOAT(NCOEF-KORDER+1))&
        ) - GRDSIZ/2.0
    2 0 ~ C O N T I N U E ~
```

!

```
    DO 30 I=1, KORDER - 1
        XKNOT(I) = XKNOT (KORDER)
        XKNOT(I+NCOEF+1) = XKNOT (NCOEF+1)
    3 0
    CONTINUE
    CALL BSLSQ (XDATA, FDATA, KORDER, XKNOT, NCOEF, BSCLSQ)
!
    DO 40 I=1, 4
        XVAL(I) = XKNOT (KORDER+I-1)
        XVAL(I+4) = XKNOT (NCOEF-3+I)
        ITYPE(I) = 1
        ITYPE(I+4) = 1
        IDER(I) = 2
        IDER(I+4) = 2
        BL(I) = 0.0
        BL(I+4) = 0.0
    40 CONTINUE
!
    XVAL(9) = -7.0
    ITYPE(9) = 3
    IDER(9) = 0
    BL(9) = 0.0
!
    XVAL(10) = -7.0
    ITYPE(10) = 2
    IDER(10) = -1
    BU(10) = 2.3
    XVAL(11) = 7.0
    ITYPE(11) = 2
    IDER(11) = -1
    BU(11) = 2.3
    XVAL(12) = -7.0
    ITYPE (12) = 10
    IDER(12) = 0
    CALL CONFT (XDATA, FDATA, XVAL, NHARPT, IDER, ITYPE, BL, BU,&
        KORDER, XKNOT, BSCNFT, NCOEF=NCOEF)
                                    Compute the average error
                                    of }100\mathrm{ points in the interval.
    ERRLSQ = 0.0
    ERRNFT = 0.0
    DO 50 I=1, 100
        X = GRDSIZ*FLOAT(I-1)/99.0 - GRDSIZ/2.0
        ERRNFT = ERRNFT + ABS(F1 (X)-BSVAL (X,KORDER,XKNOT,NCOEF,BSCNFT)&
        )
        ERRLSQ = ERRLSQ + ABS(F1 (X) -BSVAL(X,KORDER,XKNOT,NCOEF,BSCLSQ)&
            )
    5 0 ~ C O N T I N U E
!
    WRITE (NOUT,99998) ERRLSQ/100.0
    WRITE (NOUT,99999) ERRNFT/100.0
!
```

```
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ A v e r a g e ~ e r r o r ~ w i t h ~ B S L S Q ~ f i t : ~ ' , ~ F 8 . 5 )
99999 FORMAT (' Average error with CONFT fit: ', F8.5)
    END
```


## Output

```
Average error with BSLSQ fit: 0.01783
Average error with CONFT fit: 0.01339
```



Figure 3-9 CONFT vs. BSLSQ Approximating $1 /\left(1+x^{4}\right)$

## BSLS2

Computes a two-dimensional tensor-product spline approximant using least squares, returning the tensor-product B-spline coefficients.

## Required Arguments

XDATA - Array of length NXDATA containing the data points in the X -direction. (Input) XDATA must be nondecreasing.

YDATA - Array of length NYDATA containing the data points in the Y-direction. (Input) YDATA must be nondecreasing.

FDATA - Array of size NXDATA by NYDATA containing the values on the $\mathrm{X}-\mathrm{Y}$ grid to be interpolated. (Input)
$\operatorname{FDATA}(I, J)$ contains the value at (XDATA(I), YDATA(I)).
KXORD - Order of the spline in the X -direction. (Input)

KYORD - Order of the spline in the Y-direction. (Input)
XKNOT - Array of length KXORD + NXCOEF containing the knots in the X-direction. (Input) XKNOT must be nondecreasing.

YKNOT - Array of length KYORD + NYCOEF containing the knots in the Y-direction. (Input) YKNOT must be nondecreasing.

BSCOEF - Array of length NXCOEF * NYCOEF that contains the tensor product B-spline coefficients. (Output) BSCOEF is treated internally as an array of size nXCOEF by NYCOEF.

## Optional Arguments

NXDATA - Number of data points in the x -direction. (Input) Default: NXDATA $=\operatorname{size}(\mathrm{XDATA}, 1)$.

NYDATA - Number of data points in the Y-direction. (Input) Default: NYDATA $=\operatorname{size}($ YDATA, 1$)$.
$\boldsymbol{L D F}$ - Leading dimension of FDATA exactly as specified in the dimension statement of calling program. (Input) Default: LDF $=\operatorname{size}($ FDATA, $)$.

NXCOEF - Number of B-spline coefficients in the X -direction. (Input)
Default: NXCOEF $=\operatorname{size}(X K N O T, 1)-$ KXORD.
NYCOEF - Number of B-spline coefficients in the Y-direction. (Input)
Default: NYCOEF $=\operatorname{size}($ YKNOT, 1$)-$ KYORD.
XWEIGH - Array of length NXDATA containing the positive weights of XDATA. (Input) Default: XWEIGH $=1.0$.

YWEIGH - Array of length NYDATA containing the positive weights of YDATA. (Input) Default: YWEIGH $=1.0$.

## FORTRAN 90 Interface

Generic: CALL BSLS2 (XDATA, YDATA, FDATA, KXORD, KYORD, XKNOT, YKNOT, BSCOEF [, ...])

Specific: The specific interface names are S_BSLS2 and D_BSLS2.

## FORTRAN 77 Interface

Single: CALL BSLS2 (NXDATA, XDATA, NYDATA, YDATA, FDATA, LDF, KXORD, KYORD, XKNOT, YKNOT, NXCOEF, NYCOEF, XWEIGH, YWEIGH, BSCOEF)

Double: The double precision name is DBSLS2.

## Description

The routine BSLS2 computes the coefficients of a tensor-product spline least-squares approximation to weighted tensor-product data. The input for this subroutine consists of data vectors to specify the tensor-product grid for the data, two vectors with the weights, the values of the surface on the grid, and the specification for the tensor-product spline. The grid is specified by the two vectors $x=$ XDATA and $y=$ YDATA of length $n=$ NXDATA and $m=$ NYDATA, respectively. A two-dimensional array $f=$ FDATA contains the data values that are to be fit. The two vectors $w_{x}=$ XWEIGH and $w_{y}$ YYWEIGH contain the weights for the weighted least-squares problem. The information for the approximating tensor-product spline must also be provided. This information is contained in $k_{x}=$ KXORD, $\mathbf{t}_{x}=$ XKNOT, and $N=$ NXCOEF for the spline in the first variable, and in $k_{y}=$ KYORD , $\mathbf{t}_{y}=$ YKNOT and $M=$ NYCOEF for the spline in the second variable. The coefficients of the resulting tensor-product spline are returned in $c=$ BSCOEF, which is an $N * M$ array. The procedure computes coefficients by solving the normal equations in tensor-product form as discussed
in de Boor (1978, Chapter 17). The interested reader might also want to study the paper by E. Grosse (1980).
The final result produces coefficients $c$ minimizing

$$
\sum_{i=1}^{n} \sum_{j=1}^{m} w_{x}(i) w_{y}(j)\left[\sum_{k=1}^{N} \sum_{l=1}^{M} c_{k l} B_{k l}\left(x_{i}, y_{j}\right)-f_{i j}\right]^{2}
$$

where the function $B_{k l}$ is the tensor-product of two B-splines of order $k_{x}$ and $k_{y}$. Specifically, we have

$$
B_{k l}(x, y)=B_{k, k_{x}, \mathbf{t}_{x}}(x) B_{l, k_{y}, \mathbf{t}_{y}}(y)
$$

The spline

$$
\sum_{k=1}^{N} \sum_{l=1}^{M} c_{k l} B_{k l}
$$

can be evaluated using BS2VL and its partial derivatives can be evaluated using BS2DR.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2LS2/DB2LS2. The reference is:
```
CALL B2LS2 (NXDATA, XDATA, NYDATA, YDATA, FDATA, LDF, KXORD, KYORD, XKNOT,
``` YKNOT, NXCOEF, NYCOEF, XWEIGH, YWEIGH, BSCOEF, WK)

The additional argument is:
```

WK - Work array of length (NXCOEF + 1) * NYDATA + KXORD * NXCOEF + KYORD *
NYCOEF + 3 * MAX(KXORD, KYORD).

```
2. Informational errors
Type Code

314 There may be less than one digit of accuracy in the least squares fit. Try using higher precision if possible.
45 Multiplicity of the knots cannot exceed the order of the spline.
\(4 \quad 6\) The knots must be nondecreasing.
\(4 \quad 7\) All weights must be greater than zero.
49 The data point abscissae must be nondecreasing.
410 The smallest element of the data point array must be greater than or equal to the K_ORDth knot.
411 The largest element of the data point array must be less than or equal to the (n_COEF + 1)st knot.

\section*{Example}

The data for this example arise from the function \(e^{x} \sin (x+y)+\varepsilon\) on the rectangle \([0,3] \times[0,5]\). Here, \(\varepsilon\) is a uniform random variable with range \([-1,1]\). We sample this function on a \(100 \times 50\) grid and then try to recover it by using cubic splines in the \(x\) variable and quadratic splines in the \(y\) variable. We print out the values of the function \(e^{x} \sin (x+y)\) on a \(3 \times 5\) grid and compare these values with the values of the tensor-product spline that was computed using the IMSL routine BSLS2.

USE IMSL_LIBRARIES
```

IMPLICIT NONE
INTEGER KXORD, KYORD, LDF, NXCOEF, NXDATA, NXVEC, NYCOEF,\&
NYDATA, NYVEC
PARAMETER (KXORD=4, KYORD=3, NXCOEF=15, NXDATA=100, NXVEC=4,\&
NYCOEF=7, NYDATA=50, NYVEC=6, LDF=NXDATA)
INTEGER I, J, NOUT
REAL BSCOEF(NXCOEF,NYCOEF), EXP, F, FDATA (NXDATA,NYDATA),\&
FLOAT, RNOISE, SIN, VALUE (NXVEC,NYVEC), X, \&
XDATA(NXDATA), XKNOT (NXCOEF+KXORD), XVEC (NXVEC) , \&
XWEIGH (NXDATA), Y, YDATA(NYDATA) ,\&
YKNOT (NYCOEF+KYORD), YVEC(NYVEC), YWEIGH (NYDATA)
INTRINSIC EXP, FLOAT, SIN
! Define function
F(X,Y) = EXP(X)*SIN (X+Y)
CALL RNSET (1234579)
DO 10 I=1, NXCOEF - KXORD + 2

```
!
```

        XKNOT (I+KXORD-1) = 3.0*(FLOAT (I-1)/FLOAT (NXCOEF-KXORD+1))
    10 CONTINUE
    XKNOT (NXCOEF+1) = XKNOT (NXCOEF+1) + 0.001
    !
Stack knots.
DO 20 I=1, KXORD - 1
XKNOT(I) = XKNOT (KXORD)
XKNOT(I+NXCOEF+1) = XKNOT (NXCOEF+1)
20 CONTINUE
DO 30 I=1, NYCOEF - KYORD + 2
YKNOT(I+KYORD-1) = 5.0*(FLOAT(I-1)/FLOAT (NYCOEF-KYORD+1))
3 0 ~ C O N T I N U E ~
YKNOT (NYCOEF+1) = YKNOT (NYCOEF+1) + 0.001
!
DO 40 I=1, KYORD - 1
YKNOT(I) = YKNOT (KYORD)
YKNOT(I+NYCOEF+1) = YKNOT(NYCOEF+1)
4 0 ~ C O N T I N U E ~
DO 50 I=1, NXDATA
XDATA(I) = 3.0*(FLOAT (I-1)/FLOAT (NXDATA-1))
5 0 ~ C O N T I N U E
DO 60 I=1, NYDATA
YDATA(I) = 5.0*(FLOAT (I-1)/FLOAT (NYDATA-1))
CONTINUE
Evaluate function on grid and
introduce random noise in [1,-1].
DO 70 I=1, NYDATA
DO 70 J=1, NXDATA
RNOISE = RNUNF()
RNOISE = 2.0*RNOISE - 1.0
FDATA(J,I) = F(XDATA(J),YDATA(I)) + RNOISE
7 0 ~ C O N T I N U E
Use default weights equal to 1.
Compute least squares approximation.
CALL BSLS2 (XDATA, YDATA, FDATA, KXORD, KYORD, \&
XKNOT, YKNOT, BSCOEF)
Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999)
! Print interpolated values
DO 80 I=1, NXVEC
XVEC(I) = FLOAT(I-1)
80 CONTINUE
DO 90 I=1, NYVEC
YVEC(I) = FLOAT(I-1)
9 0 ~ C O N T I N U E ~
!
Evaluate spline
CALL BS2GD (0, 0, XVEC, YVEC, KXORD, KYORD, XKNOT,\&
YKNOT, BSCOEF, VALUE)
DO 110 I=1, NXVEC

```
```

            DO 100 J=1, NYVEC
                WRITE (NOUT,'(5F15.4)') XVEC(I), YVEC(J),&
                        F(XVEC(I), YVEC(J)), VALUE(I,J),&
                        (F (XVEC (I) , YVEC (J)) -VALUE (I,J))
    100 CONTINUE
    110 CONTINUE
    99999 FORMAT (13X, 'X', 14X, 'Y', 10X, 'F(X,Y)', 9X, 'S(X,Y)', 10X,\&
'Error')
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline X & Y & F (X, Y) & S (X, Y) & Error \\
\hline 0.0000 & 0.0000 & 0.0000 & 0.2782 & -0.2782 \\
\hline 0.0000 & 1.0000 & 0.8415 & 0.7762 & 0.0653 \\
\hline 0.0000 & 2.0000 & 0.9093 & 0.8203 & 0.0890 \\
\hline 0.0000 & 3.0000 & 0.1411 & 0.1391 & 0.0020 \\
\hline 0.0000 & 4.0000 & -0.7568 & -0.5705 & -0.1863 \\
\hline 0.0000 & 5.0000 & -0.9589 & -1.0290 & 0.0701 \\
\hline 1.0000 & 0.0000 & 2.2874 & 2.2678 & 0.0196 \\
\hline 1.0000 & 1.0000 & 2.4717 & 2.4490 & 0.0227 \\
\hline 1.0000 & 2.0000 & 0.3836 & 0.4947 & -0.1111 \\
\hline 1.0000 & 3.0000 & -2.0572 & -2.0378 & -0.0195 \\
\hline 1.0000 & 4.0000 & -2.6066 & -2.6218 & 0.0151 \\
\hline 1.0000 & 5.0000 & -0.7595 & -0.7274 & -0.0321 \\
\hline 2.0000 & 0.0000 & 6.7188 & 6.6923 & 0.0265 \\
\hline 2.0000 & 1.0000 & 1.0427 & 0.8492 & 0.1935 \\
\hline 2.0000 & 2.0000 & -5.5921 & -5.5885 & -0.0035 \\
\hline 2.0000 & 3.0000 & -7.0855 & -7.0955 & 0.0099 \\
\hline 2.0000 & 4.0000 & -2.0646 & -2.1588 & 0.0942 \\
\hline 2.0000 & 5.0000 & 4.8545 & 4.7339 & 0.1206 \\
\hline 3.0000 & 0.0000 & 2.8345 & 2.5971 & 0.2373 \\
\hline 3.0000 & 1.0000 & -15.2008 & -15.1079 & -0.0929 \\
\hline 3.0000 & 2.0000 & -19.2605 & -19.1698 & -0.0907 \\
\hline 3.0000 & 3.0000 & -5.6122 & -5.5820 & -0.0302 \\
\hline 3.0000 & 4.0000 & 13.1959 & 12.6659 & 0.5300 \\
\hline 3.0000 & 5.0000 & 19.8718 & 20.5170 & -0.6452 \\
\hline
\end{tabular}

\section*{BSLS3}

Computes a three-dimensional tensor-product spline approximant using least squares, returning the tensor-product B-spline coefficients.

\section*{Required Arguments}

XDATA - Array of length NXDATA containing the data points in the \(x\)-direction. (Input) XDATA must be nondecreasing.

YDATA - Array of length NYDATA containing the data points in the \(y\)-direction. (Input) YDATA must be nondecreasing.

ZDATA - Array of length NZDATA containing the data points in the \(z\)-direction. (Input) ZDATA must be nondecreasing.

FDATA - Array of size NXDATA by NYDATA by NZDATA containing the values to be interpolated. (Input)
FDATA(I, J, K) contains the value at (XDATA(I), YDATA(J), ZDATA(K)).
KXORD - Order of the spline in the \(x\)-direction. (Input)
KYORD - Order of the spline in the \(y\)-direction. (Input)
KZORD - Order of the spline in the \(z\)-direction. (Input)

XKNOT - Array of length KXORD + NXCOEF containing the knots in the \(x\)-direction. (Input) XKNOT must be nondecreasing.

YKNOT - Array of length KYORD + NYCOEF containing the knots in the \(y\)-direction. (Input) YKNOT must be nondecreasing.

ZKNOT - Array of length KZORD + NZCOEF containing the knots in the \(z\)-direction. (Input) ZKNOT must be nondecreasing.

BSCOEF - Array of length NXCOEF*NYCOEF*NZCOEF that contains the tensor product B-spline coefficients. (Output)

\section*{Optional Arguments}

NXDATA - Number of data points in the \(x\)-direction. (Input) NXDATA must be greater than or equal to NXCOEF.
Default: NXDATA \(=\operatorname{size}(\) XDATA, 1\()\).
NYDATA - Number of data points in the \(y\)-direction. (Input) NYDATA must be greater than or equal to NYCOEF. Default: NYDATA \(=\operatorname{size}(\) YDATA, 1\()\).

NZDATA - Number of data points in the \(z\)-direction. (Input) NZDATA must be greater than or equal to NZCOEF. Default: NZDATA = size (ZDATA,1).

LDFDAT - Leading dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFDAT \(=\operatorname{size}(\) FDATA, 1\()\).
MDFDAT - Second dimension of FDATA exactly as specified in the dimension statement of the calling program. (Input) Default: MDFDAT \(=\operatorname{size}(\) FDATA, 2\()\).

NXCOEF - Number of B-spline coefficients in the \(x\)-direction. (Input) Default: NXCOEF \(=\operatorname{size}(\) XKNOT, 1\()-\) KXORD.

NYCOEF - Number of B-spline coefficients in the \(y\)-direction. (Input)
Default: NYCOEF \(=\operatorname{size}(\) YKNOT, 1\()-\) KYORD.
NZCOEF - Number of B-spline coefficients in the \(z\)-direction. (Input)
Default: NZCOEF \(=\operatorname{size}(\) ZKNOT, 1\()-\) KZORD.
XWEIGH - Array of length NXDATA containing the positive weights of XDATA. (Input) Default: XWEIGH \(=1.0\).

YWEIGH - Array of length NYDATA containing the positive weights of YDATA. (Input) Default: YWEIGH \(=1.0\).

ZWEIGH - Array of length NZDATA containing the positive weights of zDATA. (Input) Default: ZWEIGH \(=1.0\).

\section*{FORTRAN 90 Interface}

Generic: CALL BSLS3 (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, BSCOEF [, ...])

Specific: The specific interface names are S_BSLS3 and D_BSLS 3 .

\section*{FORTRAN 77 Interface}

Single:
CALL BSLS3 (NXDATA, XDATA, NYDATA, YDATA, NZDATA, ZDATA, FDATA, LDFDAT, MDFDAT, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, XWEIGH, YWEIGH, ZWEIGH, BSCOEF)

Double: The double precision name is DBSLS3.

\section*{Description}

The routine BSLS3 computes the coefficients of a tensor-product spline least-squares approximation to weighted tensor-product data. The input for this subroutine consists of data vectors to specify the tensor-product grid for the data, three vectors with the weights, the values of the surface on the grid, and the specification for the tensor-product spline. The grid is specified by the three vectors \(x=\) XDATA, \(y=\) YDATA, and \(z=\) ZDATA of length \(k=\) NXDATA, \(l=\) NYDATA , and \(m=\) NYDATA, respectively. A three-dimensional array \(f=\) FDATA contains the data values which are to be fit. The three vectors \(w_{x}=\) XWEIGH, \(w_{y}\) = YWEIGH, and \(w_{z}=\) ZWEIGH contain the weights for the weighted least-squares problem. The information for the approximating tensor-product spline must also be provided. This information is contained in \(k_{x}=\mathrm{KXORD}, \mathbf{t}_{x}=\) XKNOT, and \(K=\) NXCOEF for the spline in the first variable, in \(k_{y}=\) KYORD, \(\mathbf{t}_{y}=\) YKNOT and \(L=\) NYCOEF for the spline in the second variable, and in \(k_{z}=\mathrm{KZORD}, \mathbf{t}_{z}=\mathrm{ZKNOT}\) and \(M=\mathrm{NZCOEF}\) for the spline in the third variable.

The coefficients of the resulting tensor product spline are returned in \(c=\) BSCOEF, which is an \(K \times L \times M\) array. The procedure computes coefficients by solving the normal equations in tensorproduct form as discussed in de Boor (1978, Chapter 17). The interested reader might also want to study the paper by E. Grosse (1980).

The final result produces coefficients \(c\) minimizing
\[
\sum_{i=l}^{k} \sum_{j=1}^{l} \sum_{p=1}^{m} w_{x}(i) w_{y}(j) w_{z}(p)\left[\sum_{s=1}^{K} \sum_{t=1}^{L} \sum_{u=1}^{M} c_{s t u} B_{s t u}\left(x_{i}, y_{j}, z_{p}\right)-f_{i j p}\right]^{2}
\]
where the function \(B_{s t u}\) is the tensor-product of three B -splines of order \(k_{x}, k_{y}\), and \(k_{z}\). Specifically, we have
\[
B_{s t u}(x, y, z)=B_{s, k_{x}, \mathbf{t}_{x}}(x) B_{t, k_{y}, t_{y}}(y) B_{u, k_{z}, \mathbf{t}_{z}}(z)
\]

The spline
\[
\sum_{s=1}^{K} \sum_{t=1}^{L} \sum_{u=1}^{M} c_{s t u} B_{s t u}
\]
can be evaluated at one point using BS3VL and its partial derivatives can be evaluated using \(B S 3 D R\). If the values on a grid are desired then we recommend BS3GD.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of B2LS3 / DB2LS3. The reference is:

CALL B2LS3 (NXDATA, XDATA, NYDATA, NZDATA, ZDATA, YDATA, FDATA, LDFDAT, KXORD, KYORD, KZORD, XKNOT, YKNOT, ZKNOT, NXCOEF, NYCOEF, NZCOEF, XWEIGH, YWEIGH, ZWEIGH, BSCOEF, WK)

The additional argument is:
WK - Work array of length NYCOEF * (NZDATA + KYORD + NZCOEF) + NZDATA *
\((1+\) NYDATA \()+\) NXCOEF * \((\) KXORD + NYDATA * NZDATA \()+\) KZORD * NZCOEF + 3 * MAXO(KXORD, KYORD, KZORD).
2. Informational errors
\begin{tabular}{ccl} 
Type & Code \\
3 & 13 & \begin{tabular}{l} 
There may be less than one digit of accuracy in the least squares fit. \\
Try using higher precision if possible.
\end{tabular} \\
4 & 7 & Multiplicity of knots cannot exceed the order of the spline. \\
4 & 8 & The knots must be nondecreasing. \\
4 & 9 & All weights must be greater than zero. \\
4 & 10 & \begin{tabular}{l} 
The data point abscissae must be nondecreasing. \\
4
\end{tabular} \\
11 & \begin{tabular}{l} 
The smallest element of the data point array must be greater than or \\
equal to the K_ORDth knot. \\
The largest element of the data point array must be less than or equal \\
to the (N_COEF + 1) st knot.
\end{tabular}
\end{tabular}

\section*{Example}

The data for this example arise from the function \(e^{(y-z)} \sin (x+y)+\varepsilon\) on the rectangle \([0,3] \times[0,2] \times[0,1]\). Here, \(\varepsilon\) is a uniform random variable with range \([-.5, .5]\). We sample this function on a \(4 \times 3 \times 2\) grid and then try to recover it by using tensor-product cubic splines in all variables. We print out the values of the function \(e^{(y-z)} \sin (x+y)\) on a \(4 \times 3 \times 2\) grid and compare these values with the values of the tensor-product spline that was computed using the IMSL routine BSLS3.
```

    USE BSLS3 INT
    USE RNSET - INT
    USE RNUNF_INT
    USE UMACH - INT
    USE BS3GD_INT
    IMPLICIT NONE
    INTEGER KXORD, KYORD, KZORD, LDFDAT, MDFDAT, NXCOEF, NXDATA,&
            NXVAL, NYCOEF, NYDATA, NYVAL, NZCOEF, NZDATA, NZVAL
    PARAMETER (KXORD=4, KYORD=4, KZORD=4, NXCOEF=8, NXDATA=15,&
            NXVAL=4, NYCOEF=8, NYDATA=15, NYVAL=3, NZCOEF=8,&
            NZDATA=15, NZVAL=2, LDFDAT=NXDATA, MDFDAT=NYDATA)
    INTEGER I, J, K, NOUT
    REAL BSCOEF(NXCOEF,NYCOEF,NZCOEF), EXP, F,&
            FDATA(NXDATA,NYDATA,NZDATA), FLOAT, RNOISE,&
            SIN, SPXYZ(NXVAL,NYVAL,NZVAL), X, XDATA(NXDATA),&
            XKNOT (NXCOEF+KXORD), XVAL(NXVAL), XWEIGH (NXDATA), Y,&
            YDATA(NYDATA), YKNOT (NYCOEF+KYORD), YVAL(NYVAL),&
            YWEIGH (NYDATA), Z, ZDATA(NZDATA),&
            ZKNOT(NZCOEF+KZORD), ZVAL(NZVAL), ZWEIGH(NZDATA)
    INTRINSIC EXP, FLOAT, SIN
    ! Define a function
F(X,Y,Z) = EXP(Y-Z)*SIN (X+Y)
CALL RNSET (1234579)
CALL UMACH (2, NOUT)
Set up knot sequences
X-knots
DO 10 I=1, NXCOEF - KXORD + 2
XKNOT(I+KXORD-1) = 3.0*(FLOAT(I-1)/FLOAT (NXCOEF-KXORD+1))
10 CONTINUE
DO 20 I=1, KXORD - 1
XKNOT(I) = XKNOT(KXORD)
XKNOT(I+NXCOEF+1) = XKNOT (NXCOEF+1)
2 0 ~ C O N T I N U E ~
Y-knots
DO 30 I=1, NYCOEF - KYORD + 2
YKNOT(I+KYORD-1) = 2.0*(FLOAT(I-1)/FLOAT (NYCOEF-KYORD+1))
30 CONTINUE
DO 40 I=1, KYORD - 1
YKNOT(I) = YKNOT(KYORD)
YKNOT(I+NYCOEF+1) = YKNOT(NYCOEF+1)
40 CONTINUE
z-knots

```
\(!\)
!
```

            DO 50 I=1, NZCOEF - KZORD + 2
        ZKNOT (I+KZORD-1) = 1.0*(FLOAT (I-1) /FLOAT (NZCOEF-KZORD+1))
    5 0 ~ C O N T I N U E
    DO 60 I=1, KZORD - 1
        ZKNOT(I) = ZKNOT(KZORD)
        ZKNOT (I+NZCOEF+1) = ZKNOT (NZCOEF+1)
    6 0 ~ C O N T I N U E ~
    DO 70 I=1, NXDATA
        XDATA(I) = 3.0*(FLOAT (I-1)/FLOAT (NXDATA-1))
    7 0 ~ C O N T I N U E
    !
DO 80 I=1, NYDATA
YDATA(I) = 2.0*(FLOAT (I-1)/FLOAT (NYDATA-1))
8 0 ~ C O N T I N U E ~
DO 90 I=1, NZDATA
ZDATA(I) = 1.0*(FLOAT (I-1)/FLOAT (NZDATA-1))
90 CONTINUE
Evaluate the function on the grid
and add noise.
DO 100 I=1, NXDATA
DO 100 J=1, NYDATA
DO 100 K=1, NZDATA
RNOISE = RNUNF()
RNOISE = RNOISE - 0.5
FDATA(I,J,K) = F(XDATA(I),YDATA(J),ZDATA(K)) + RNOISE
100 CONTINUE
!
!
Use default weights equal to 1.0
Compute least-squares
CALL BSLS3 (XDATA, YDATA, ZDATA, FDATA, KXORD, KYORD, KZORD, XKNOT, \&
YKNOT, ZKNOT, BSCOEF)
DO 110 I=1, NXVAL
XVAL(I) = FLOAT(I-1)
110 CONTINUE
DO 120 I=1, NYVAL
YVAL(I) = FLOAT(I-1)
120 CONTINUE
DO 130 I=1, NZVAL
ZVAL(I) = FLOAT(I-1)
130 CONTINUE
! Evaluate on the grid.
CALL BS3GD (0, 0, 0, XVAL, YVAL, ZVAL, KXORD, KYORD, KZORD, XKNOT, \&
YKNOT, ZKNOT, BSCOEF, SPXYZ)
Print results.
WRITE (NOUT,99998)
DO 140 I=1, NXVAL
DO 140 J=1, NYVAL
DO 140 K=1, NZVAL
WRITE (NOUT,99999) XVAL(I), YVAL(J), ZVAL(K),\&
F(XVAL (I) , YVAL (J), ZVAL (K)),\&
SPXYZ(I,J,K), F(XVAL(I),YVAL(J),ZVAL (K)\&
) - SPXYZ(I,J,K)

```
```

    140 CONTINUE
    99998 FORMAT (8X, 'X', 9X, 'Y', 9X, 'Z', 6X, 'F(X,Y,Z)', 3X,\&
'S(X,Y,Z)', 4X, 'Error')
99999 FORMAT (' ', 3F10.3, 3F11.4)
END

```

\section*{Output}
\begin{tabular}{ccccrr}
\(X\) & \(Y\) & \(Z\) & \(F(X, Y, Z)\) & \(S(X, Y, Z)\) & Error \\
0.000 & 0.000 & 0.000 & 0.0000 & 0.1987 & -0.1987 \\
0.000 & 0.000 & 1.000 & 0.0000 & 0.1447 & -0.1447 \\
0.000 & 1.000 & 0.000 & 2.2874 & 2.2854 & 0.0019 \\
0.000 & 1.000 & 1.000 & 0.8415 & 1.0557 & -0.2142 \\
0.000 & 2.000 & 0.000 & 6.7188 & 6.4704 & 0.2484 \\
0.000 & 2.000 & 1.000 & 2.4717 & 2.2054 & 0.2664 \\
1.000 & 0.000 & 0.000 & 0.8415 & 0.8779 & -0.0365 \\
1.000 & 0.000 & 1.000 & 0.3096 & 0.2571 & 0.0524 \\
1.000 & 1.000 & 0.000 & 2.4717 & 2.4015 & 0.0703 \\
1.000 & 1.000 & 1.000 & 0.9093 & 0.8995 & 0.0098 \\
1.000 & 2.000 & 0.000 & 1.0427 & 1.1330 & -0.0902 \\
1.000 & 2.000 & 1.000 & 0.3836 & 0.4951 & -0.1115 \\
2.000 & 0.000 & 0.000 & 0.9093 & 0.8269 & 0.0824 \\
2.000 & 0.000 & 1.000 & 0.3345 & 0.3258 & 0.0087 \\
2.000 & 1.000 & 0.000 & 0.3836 & 0.3564 & 0.0272 \\
2.000 & 1.000 & 1.000 & 0.1411 & 0.1905 & -0.0494 \\
2.000 & 2.000 & 0.000 & -5.5921 & -5.5362 & -0.0559 \\
2.000 & 2.000 & 1.000 & -2.0572 & -1.9659 & -0.0913 \\
3.000 & 0.000 & 0.000 & 0.1411 & 0.4841 & -0.3430 \\
3.000 & 0.000 & 1.000 & 0.0519 & -0.4257 & 0.4776 \\
3.000 & 1.000 & 0.000 & -2.0572 & -1.9710 & -0.0862 \\
3.000 & 1.000 & 1.000 & -0.7568 & -0.8479 & 0.0911 \\
3.000 & 2.000 & 0.000 & -7.0855 & -7.0957 & 0.0101 \\
3.000 & 2.000 & 1.000 & -2.6066 & -2.1650 & -0.4416
\end{tabular}

\section*{CSSED}

Smooths one-dimensional data by error detection.

\section*{Required Arguments}

XDATA - Array of length NDATA containing the abscissas of the data points. (Input)

FDATA - Array of length NDATA containing the ordinates (function values) of the data points. (Input)

DIS - Proportion of the distance the ordinate in error is moved to its interpolating curve. (Input)
It must be in the range 0.0 to 1.0. A suggested value for DIS is one.
\(\mathbf{S C}\) - Stopping criterion. (Input)
SC should be greater than or equal to zero. A suggested value for SC is zero.

MAXIT - Maximum number of iterations allowed. (Input)
SDATA - Array of length NDATA containing the smoothed data. (Output)

\section*{Optional Arguments}

NDATA - Number of data points. (Input)
Default: NDATA \(=\) size \((X D A T A, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL CSSED (XDATA, FDATA, DIS, SC, MAXIT, SDATA [, ...])

Specific: The specific interface names are S_CSSED and D_CSSED.

\section*{FORTRAN 77 Interface}

Single: CALL CSSED (NDATA, XDATA, FDATA, DIS, SC, MAXIT, SDATA)
Double: The double precision name is DCSSED.

\section*{Description}

The routine CSSED is designed to smooth a data set that is mildly contaminated with isolated errors. In general, the routine will not work well if more than \(25 \%\) of the data points are in error. The routine CSSED is based on an algorithm of Guerra and Tapia (1974).

Setting NDATA \(=n\), FDATA \(=f, \operatorname{SDATA}=s\) and \(\mathrm{XDATA}=x\), the algorithm proceeds as follows. Although the user need not input an ordered XDATA sequence, we will assume that \(x\) is increasing for simplicity. The algorithm first sorts the XDATA values into an increasing sequence and then continues. A cubic spline interpolant is computed for each of the 6 -point data sets (initially setting \(s=f\) )
\[
\left(x_{j}, s_{j}\right) \quad j=i-3, \ldots, i+3 j \neq i,
\]
where \(i=4, \ldots, n-3\) using CSAKM. For each \(i\) the interpolant, which we will call \(S_{i}\), is compared with the current value of \(s_{i}\), and a 'point energy' is computed as
\[
p e_{i}=S_{i}\left(x_{i}\right)-s_{i}
\]

Setting \(s c=S C\), the algorithm terminates either if MAXIT iterations have taken place or if
\[
\left|p e_{i}\right| \leq s c\left(x_{i+3}-x_{i-3}\right) / 6 \quad i=4, \ldots, n-3
\]

If the above inequality is violated for any \(i\), then we update the \(i\)-th element of \(s\) by setting \(s_{i}=s_{i}+d\left(p e_{i}\right)\), where \(d=\) DIS. Note that neither the first three nor the last three data points are changed. Thus, if these points are inaccurate, care must be taken to interpret the results.

The choice of the parameters \(d, s c\) and MAXIT are crucial to the successful usage of this subroutine. If the user has specific information about the extent of the contamination, then he
should choose the parameters as follows: \(d=1, s c=0\) and MAXIT to be the number of data points in error. On the other hand, if no such specific information is available, then choose \(d=.5\), MAXIT \(\leq 2 n\), and
\[
s c=.5 \frac{\max s-\min s}{\left(x_{n}-x_{1}\right)}
\]

In any case, we would encourage the user to experiment with these values.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of C2SED/DC2SED. The reference is:

CALL C2SED (NDATA, XDATA, FDATA, DIS, SC, MAXIT, DATA, WK, IWK)
The additional arguments are as follows:
\(\boldsymbol{W K}\) — Work array of length 4 * NDATA +30.
\(\boldsymbol{I W K}\) - Work array of length 2 * NDATA.
2. Informational error

Type Code
31 The maximum number of iterations allowed has been reached.
3. The arrays FDATA and SDATA may the the same.

\section*{Example}

We take 91 uniform samples from the function \(5+\left(5+t^{2} \sin t\right) / t\) on the interval \([1,10]\). Then, we contaminate 10 of the samples and try to recover the original function values.
```

USE CSSED_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=91)
INTEGER I, MAXIT, NOUT, ISB(10)
REAL DIS, F, FDATA(91), SC, SDATA(91), SIN, X, XDATA(91),\&
RNOISE(10)
INTRINSIC SIN
DATA ISB/6, 17, 26, 34, 42, 49, 56, 62, 75, 83/
DATA RNOISE/2.5, -3.0, -2.0, 2.5, 3.0, -2.0, -2.5, 2.0, -2.0, 3.0/
F(X) = (X*X*SIN (X) +5.0)/X + 5.0
EX. \#1; No specific information
available
DIS = 0.5

```
!
!
```

    SC = 0.56
    MAXIT = 182
    !
XDATA(1) = 1.0
FDATA(1) = F(XDATA(1))
DO 10 I=2, NDATA
XDATA(I) = XDATA(I-1) + .1
FDATA(I) = F(XDATA(I))
1 0 ~ C O N T I N U E
!
DO 20 I=1, 10
FDATA(ISB(I)) = FDATA(ISB(I)) + RNOISE(I)
20 CONTINUE
CALL CSSED (XDATA, FDATA, DIS, SC, MAXIT, SDATA)
Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99997)
DO 30 I=1, 10
WRITE (NOUT,99999) F(XDATA(ISB(I))), FDATA(ISB(I)),\&
SDATA(ISB(I))
3 0 ~ C O N T I N U E ~
DIS = 1.0
SC = 0.0
MAXIT = 10
A warning message is produced
because the maximum number of
iterations is reached.
Smooth data
CALL CSSED (XDATA, FDATA, DIS, SC, MAXIT, SDATA)
Write heading
Write data
DO 40 I=1, 10
WRITE (NOUT,99999) F(XDATA(ISB(I))), FDATA(ISB(I)),\&
SDATA(ISB(I))
4 0 ~ C O N T I N U E ~
!
9 9 9 9 7 ~ F O R M A T ~ ( ' ~ C a s e ~ A ~ - ~ N o ~ s p e c i f i c ~ i n f o r m a t i o n ~ a v a i l a b l e ' , ~ / , \& ~
F(X) F(X)+NOISE SDATA(X)', /)
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ C a s e ~ B ~ - ~ S p e c i f i c ~ i n f o r m a t i o n ~ a v a i l a b l e ' , ~ / , \& ~
F(X) F(X)+NOISE SDATA(X)', /)
99999 FORMAT (' ', F7.3, 8X, F7.3, 11X, F7.3)
END

```

\section*{Output}
```

Case A - No specific information available
F(X) F(X)+NOISE SDATA(X)

```


\section*{CSSMH}

Computes a smooth cubic spline approximation to noisy data.

\section*{Required Arguments}

XDATA - Array of length NDATA containing the data point abscissas. (Input) XDATA must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
SMPAR - A nonnegative number which controls the smoothing. (Input) The spline function \(S\) returned is such that the sum from \(I=1\) to NDATA of \(((S(\operatorname{XDATA}(\mathrm{I})) \operatorname{FDATA}(\mathrm{I})) /\) weight(I))\() * 2\) is less than or equal to SMPAR. It is recommended that SMPAR lie in the confidence interval of this sum, i.e., NDATA \(-\operatorname{SQRT}(2\) * NDATA).LE. SMPAR.LE. NDATA \(+\operatorname{SQRT}(2\) * NDATA).

BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces. (Output)

\section*{Optional Arguments}

NDATA - Number of data points. (Input)
NDATA must be at least 2 .
Default: NDATA \(=\operatorname{size}(X D A T A, 1)\).
WEIGHT - Array of length NDATA containing estimates of the standard deviations of FDATA. (Input)
All elements of WEIGHT must be positive.
Default: WE IGHT \(=1.0\).

\section*{FORTRAN 90 Interface}

Generic: CALL CSSMH (XDATA, FDATA, SMPAR, BREAK, CSCOEF [, ...])

Specific: The specific interface names are S_CSSMH and D_CSSMH.

\section*{FORTRAN 77 Interface}

Single: CALL CSSMH (NDATA, XDATA, FDATA, WEIGHT, SMPAR, BREAK, CSCOEF)

Double: The double precision name is DCSSMH.

\section*{Description}

The routine CSSMH is designed to produce a \(C^{2}\) cubic spline approximation to a data set in which the function values are noisy. This spline is called a smoothing spline. It is a natural cubic spline with knots at all the data abscissas \(x=\) XDATA, but it does not interpolate the data \(\left(x_{i}, f_{i}\right)\). The smoothing spline \(S\) is the unique \(C^{2}\) function which minimizes
\[
\int_{a}^{b} S^{\prime \prime}(x)^{2} d x
\]
subject to the constraint
\[
\sum_{i=1}^{N}\left|\frac{S\left(x_{i}\right)-f_{i}}{w_{i}}\right|^{2} \leq \sigma
\]
where \(w=\) WEIGHT, \(\sigma=\) SMPAR is the smoothing parameter, and \(N=\) NDATA.
Recommended values for \(\sigma\) depend on the weights \(w\). If an estimate for the standard deviation of the error in the value \(f_{i}\) is available, then \(w_{i}\) should be set to this value and the smoothing parameter \(\sigma\) should be chosen in the confidence interval corresponding to the left side of the above inequality. That is,
\[
N-\sqrt{2 N} \leq \sigma \leq N+\sqrt{2 N}
\]

The routine CSSMH is based on an algorithm of Reinsch (1967). This algorithm is also discussed in de Boor (1978, pages 235-243).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of C2SMH/DC2SMH. The reference is:

CALL C2SMH (NDATA, XDATA, FDATA, WEIGHT, SMPAR, BREAK, CSCOEF, WK, IWK)
The additional arguments are as follows:
\(\boldsymbol{W K}-\) Work array of length \(8 *\) NDATA +5.
\(I W K\) - Work array of length ndATA.
2. Informational errors
Type Code

31 The maximum number of iterations has been reached. The best approximation is returned.
43 All weights must be greater than zero.
3. The cubic spline can be evaluated using CSVAL; its derivative can be evaluated using CSDER.

\section*{Example}

In this example, function values are contaminated by adding a small "random" amount to the correct values. The routine CSSMH is used to approximate the original, uncontaminated data.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=300)
!
INTEGER I, NOUT
REAL BREAK (NDATA), \(\operatorname{CSCOEF}(4, N D A T A), ~ E R R O R, F, \&\)
FDATA (NDATA), FLOAT, FVAL, SDEV, SMPAR, SQRT, \&
SVAL, WEIGHT (NDATA), X, XDATA (NDATA), XT, RN
INTRINSIC FLOAT, SQRT
!
\(F(X)=1.0 /(.1+(3.0 *(X-1.0)) * * 4)\)
! Set up a grid
DO 10 I=1, NDATA XDATA \((I)=3.0 *(F L O A T(I-1) / F L O A T(N D A T A-1))\) FDATA (I) \(=F(X D A T A(I))\)
10 CONTINUE
!
CALL RNSET (1234579)

DO 20 I=1, NDATA
RN = RNUNF ()
FDATA \((I)=\) FDATA \((I)+2.0 * R N-1.0\)
20 CONTINUE
```

! Set the WEIGHT vector
SDEV = 1.0/SQRT (3.0)
CALL SSET (NDATA, SDEV, WEIGHT, 1)
SMPAR = NDATA
! Smooth the data
CALL CSSMH (XDATA, FDATA, SMPAR, BREAK, CSCOEF, WEIGHT=WEIGHT)
Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999)
DO 30 I=1, 10
XT = 90.0*(FLOAT(I-1)/FLOAT (NDATA-1))
Evaluate the spline
SVAL = CSVAL (XT, BREAK, CSCOEF)
FVAL = F(XT)
ERROR = SVAL - FVAL
WRITE (NOUT,'(4F15.4)') XT, FVAL, SVAL, ERROR
3 0 ~ C O N T I N U E
!
99999 FORMAT (12X, 'X', 9X, 'Function', 7X, 'Smoothed', 10X,\&
'Error')
END

```

\section*{Output}
\begin{tabular}{cccr} 
X & Function & Smoothed & Error \\
0.0000 & 0.0123 & 0.1118 & 0.0995 \\
0.3010 & 0.0514 & 0.0646 & 0.0131 \\
0.6020 & 0.4690 & 0.2972 & -0.1718 \\
0.9030 & 9.3312 & 8.7022 & -0.6289 \\
1.2040 & 4.1611 & 4.7887 & 0.6276 \\
1.5050 & 0.1863 & 0.2718 & 0.0856 \\
1.8060 & 0.0292 & 0.1408 & 0.1116 \\
2.1070 & 0.0082 & 0.0826 & 0.0743 \\
2.4080 & 0.0031 & 0.0076 & 0.0045 \\
2.7090 & 0.0014 & -0.1789 & -0.1803
\end{tabular}

\section*{CSSCV}

Computes a smooth cubic spline approximation to noisy data using cross-validation to estimate the smoothing parameter.

\section*{Required Arguments}

XDATA - Array of length NDATA containing the data point abscissas. (Input) XDATA must be distinct.

FDATA - Array of length NDATA containing the data point ordinates. (Input)
IEQUAL - A flag alerting the subroutine that the data is equally spaced. (Input)

BREAK - Array of length NDATA containing the breakpoints for the piecewise cubic representation. (Output)

CSCOEF - Matrix of size 4 by NDATA containing the local coefficients of the cubic pieces. (Output)

\section*{Optional Arguments}

NDATA - Number of data points. (Input)
NDATA must be at least 3 .
Default: NDATA \(=\operatorname{size}(\mathrm{XDATA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL CSSCV (XDATA, FDATA, IEQUAL, BREAK, CSCOEF [,...])
Specific: The specific interface names are S_CSSCV and D_CSSCV.

\section*{FORTRAN 77 Interface}

Single: CALL CSSCV (NDATA, XDATA, FDATA, IEQUAL, BREAK, CSCOEF)
Double: The double precision name is DCSSCV.

\section*{Description}

The routine \(\operatorname{CSSCV}\) is designed to produce a \(C^{2}\) cubic spline approximation to a data set in which the function values are noisy. This spline is called a smoothing spline. It is a natural cubic spline with knots at all the data abscissas \(x=\) XDATA, but it does not interpolate the data \(\left(x_{i}, f_{i}\right)\). The smoothing spline \(S_{S}\) is the unique \(C^{2}\) function that minimizes
\[
\int_{a}^{b} S_{\sigma}^{\prime \prime}(x)^{2} d x
\]
subject to the constraint
\[
\sum_{i=1}^{N}\left|S_{\sigma}\left(x_{i}\right)-f_{i}\right|^{2} \leq \sigma
\]
where \(\sigma\) is the smoothing parameter and \(N=\) NDATA. The reader should consult Reinsch (1967) for more information concerning smoothing splines. The IMSL subroutine CSSMH solves the above problem when the user provides the smoothing parameter \(\sigma\). This routine attempts to find the 'optimal' smoothing parameter using the statistical technique known as cross-validation. This means that (in a very rough sense) one chooses the value of \(\sigma\) so that the smoothing spline \(\left(S_{s}\right)\) best approximates the value of the data at \(x_{i}\), if it is computed using all the data except the \(i\)-th; this is true for all \(i=1, \ldots, N\). For more information on this topic, we refer the reader to Craven and Wahba (1979).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(\mathrm{C} 2 \mathrm{SCV} / \mathrm{DC} 2 \mathrm{SCV}\). The reference is:

CALL C2SCV (NDATA, XDATA, FDATA, IEQUAL, BREAK, CSCOEF, WK, SDWK, IPVT)
The additional arguments are as follows:
\(W K-\) Work array of length 7 * (NDATA +2 ).
\(\boldsymbol{S D W K}\) - Work array of length 2 * NDATA.
IPVT - Work array of length NDATA.
2. Informational error

Type Code
42 Points in the data point abscissas array, XDATA, must be distinct.

\section*{Example}

In this example, function values are computed and are contaminated by adding a small "random" amount. The routine CSSCV is used to try to reproduce the original, uncontaminated data.
```

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NDATA
PARAMETER (NDATA=300)
INTEGER I, IEQUAL, NOUT
REAL BREAK(NDATA), CSCOEF (4,NDATA), ERROR, F,\&
FDATA(NDATA), FLOAT, FVAL, SVAL, X,\&
XDATA(NDATA), XT, RN
INTRINSIC FLOAT
F(X) = 1.0/(.1+(3.0*(X-1.0))**4)
CALL UMACH (2, NOUT)
DO 10 I=1, NDATA
XDATA(I) = 3.0*(FLOAT (I-1)/FLOAT (NDATA-1))
FDATA(I) = F(XDATA(I))
1 0 ~ C O N T I N U E
Introduce noise on [-.5,.5]
Contaminate the data
CALL RNSET (1234579)
DO 20 I=1, NDATA
RN = RNUNF ()
FDATA(I) = FDATA(I) + 2.0*RN - 1.0
20 CONTINUE

```
!
\(!\)
\(!\)
!
```

    IEQUAL = 1
    ! Smooth data
CALL CSSCV (XDATA, FDATA, IEQUAL, BREAK, CSCOEF)
!
WRITE (NOUT,99999)
DO 30 I=1, 10
XT = 90.0*(FLOAT(I-1)/FLOAT (NDATA-1))
SVAL = CSVAL(XT,BREAK,CSCOEF)
FVAL = F(XT)
ERROR = SVAL - FVAL
WRITE (NOUT,'(4F15.4)') XT, FVAL, SVAL, ERROR
3 0 ~ C O N T I N U E ~
99999 FORMAT (12X, 'X', 9X, 'Function', 7X, 'Smoothed', 10X,\&
'Error')
END

```

\section*{Output}
\begin{tabular}{crrr} 
X & Function & Smoothed & Error \\
0.0000 & 0.0123 & 0.2528 & 0.2405 \\
0.3010 & 0.0514 & 0.1054 & 0.0540 \\
0.6020 & 0.4690 & 0.3117 & -0.1572 \\
0.9030 & 9.3312 & 8.9461 & -0.3850 \\
1.2040 & 4.1611 & 4.6847 & 0.5235 \\
1.5050 & 0.1863 & 0.3819 & 0.1956 \\
1.8060 & 0.0292 & 0.1168 & 0.0877 \\
2.1070 & 0.0082 & 0.0658 & 0.0575 \\
2.4080 & 0.0031 & 0.0395 & 0.0364 \\
2.7090 & 0.0014 & -0.2155 & -0.2169
\end{tabular}

\section*{RATCH}

Computes a rational weighted Chebyshev approximation to a continuous function on an interval.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be approximated. The form is \(\mathrm{F}(\mathrm{X})\), where

X - Independent variable. (Input)
F - The function value. (Output)

F must be declared EXTERNAL in the calling program.

PHI - User-supplied FUNCTION to supply the variable transformation which must be continuous and monotonic. The form is PHI(X), where

X - Independent variable. (Input)
PHI - The function value. (Output)

PHI must be declared EXTERNAL in the calling program.

WEIGHT - User-supplied FUNCTION to scale the maximum error. It must be continuous and nonvanishing on the closed interval (A, B). The form is WEIGHT(X), where

X - Independent variable. (Input)
WEIGHT - The function value. (Output)
WEIGHT must be declared EXTERNAL in the calling program.
\(\boldsymbol{A}\) - Lower end of the interval on which the approximation is desired. (Input)
\(\boldsymbol{B}\) - Upper end of the interval on which the approximation is desired. (Input)
\(\boldsymbol{P}-\) Vector of length \(\mathrm{N}+1\) containing the coefficients of the numerator polynomial. (Output)
\(\boldsymbol{Q}\) - Vector of length \(\mathrm{M}+1\) containing the coefficients of the denominator polynomial. (Output)
\(\boldsymbol{E R R O R}\) - Min-max error of approximation. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - The degree of the numerator. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{P}, 1)-1\).
\(\boldsymbol{M}\) - The degree of the denominator. (Input)
Default: \(\mathrm{M}=\operatorname{size}(\mathrm{e}, 1)-1\).

\section*{FORTRAN 90 Interface}

Generic: CALL RATCH (F, PHI, WEIGHT, A, B, P, Q, ERROR [, ...])
Specific: The specific interface names are S_RATCH and D_RATCH.

\section*{FORTRAN 77 Interface}

Single: CALL RATCH (F, PHI, WEIGHT, A, B, N, M, P, Q, ERROR)
Double: The double precision name is DRATCH.

\section*{Description}

The routine RATCH is designed to compute the best weighted \(L_{\infty}\) (Chebyshev) approximant to a given function. Specifically, given a weight function \(w=W E I G H T\), a monotone function \(\phi=P H I\), and a function \(f\) to be approximated on the interval \([a, b]\), the subroutine RATCH returns the coefficients (in \(P\) and \(Q\) ) for a rational approximation to \(f\) on \([a, b]\). The user must supply the degree of the numerator \(N\) and the degree of the denominator \(M\) of the rational function
\[
R_{M}^{N}
\]

The goal is to produce coefficients which minimize the expression
\[
\left\|\frac{f-R_{M}^{N}}{w}\right\|:=\max _{x \in[a, b]}^{\left|f(x)-\frac{\sum_{i=1}^{N+1} P_{i} \phi^{i-1}(x)}{\sum_{i=1}^{M+1} Q_{i} \phi^{i-1}(x)}\right|} \frac{w(x)}{|c|}
\]

Notice that setting \(\phi(x)=x\) yields ordinary rational approximation. A typical use of the function \(\phi\) occurs when one wants to approximate an even function on a symmetric interval, say \([-a, a]\) using ordinary rational functions. In this case, it is known that the answer must be an even function. Hence, one can set \(\phi(x)=x^{2}\), only approximate on \([0, a]\), and decrease by one half the degrees in the numerator and denominator.

The algorithm implemented in this subroutine is designed for fast execution. It assumes that the best approximant has precisely \(N+M+2\) equi-oscillations. That is, that there exist \(N+M+2\) points \(\mathbf{t}_{1}<\ldots<\mathbf{t}_{N+M+2}\) satisfying
\[
e\left(\mathbf{t}_{i}\right)=-e\left(\mathbf{t}_{i+1}\right)= \pm\left\|\frac{f-R_{M}^{N}}{w}\right\|
\]

Such points are called alternants. Unfortunately, there are many instances in which the best rational approximant to the given function has either fewer alternants or more alternants. In this case, it is not expected that this subroutine will perform well. For more information on rational Chebyshev approximation, the reader can consult Cheney (1966). The subroutine is based on work of Cody, Fraser, and Hart (1968).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of R2TCH/DR2TCH. The reference is:
```

CALL R2TCH (F, PHI,WEIGHT, A, B, N, M, P, Q, ERROR, ITMAX, IWK, WK)

```

The additional arguments are as follows:
ITMAX - Maximum number of iterations. (Input) The default value is 20 .

IWK - Workspace vector of length \((\mathrm{N}+\mathrm{M}+2\) ). (Workspace)
WK - Workspace vector of length \((\mathrm{N}+\mathrm{M}+8) *(\mathrm{~N}+\mathrm{M}+2)\). (Workspace)
2. Informational errors

\section*{Type Code}

31 The maximum number of iterations has been reached. The routine R2TCH may be called directly to set a larger value for ItmAx.
32 The error was reduced as far as numerically possible. A good approximation is returned in \(P\) and \(Q\), but this does not necessarily give the Chebyshev approximation.
43 The linear system that defines \(P\) and \(Q\) was found to be algorithmically singular. This indicates the possibility of a degenerate approximation.
\(4 \quad 4\) A sequence of critical points that was not monotonic generated. This indicates the possibility of a degenerate approximation.
45 The value of the error curve at some critical point is too large. This indicates the possibility of poles in the rational function.
46 The weight function cannot be zero on the closed interval (A, B).

\section*{Example}

In this example, we compute the best rational approximation to the gamma function, \(\Gamma\), on the interval \([2,3]\) with weight function \(w=1\) and \(N=M=2\). We display the maximum error and the coefficients. This problem is taken from the paper of Cody, Fraser, and Hart (1968). We compute in double precision due to the conditioning of this problem.
```

USE RATCH_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER M, N
PARAMETER (M=2, N=2)
INTEGER NOUT
DOUBLE PRECISION A, B, ERROR, F, P(N+1), PHI, Q(M+1), WEIGHT
EXTERNAL F, PHI, WEIGHT
A = 2.0D0
B = 3.0D0
! Compute double precision rational
approximation
CALL RATCH (F, PHI, WEIGHT, A, B, P, Q, ERROR)
Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,'(1X,A)') 'In double precision we have:'
WRITE (NOUT,99999) 'P = ', P
WRITE (NOUT,99999) 'Q = ', Q
WRITE (NOUT,99999) 'ERROR = ', ERROR
99999 FORMAT (' ', A, 5X, 3F20.12, /)
END
! ---------------------------------------------------------------------------------
DOUBLE PRECISION FUNCTION F (X)
DOUBLE PRECISION X

```
!
```

!
DOUBLE PRECISION DGAMMA
EXTERNAL DGAMMA
!
F = DGAMMA(X)
RETURN
END
! -----------------------------------------------------------------------------
DOUBLE PRECISION FUNCTION PHI (X)
DOUBLE PRECISION X
!
PHI = X
RETURN
END
! -----------------------------------------------------------------------------
!
DOUBLE PRECISION FUNCTION WEIGHT (X)
DOUBLE PRECISION X
!
DOUBLE PRECISION DGAMMA
EXTERNAL DGAMMA
!
WEIGHT = DGAMMA(X)
RETURN
END

```

\section*{Output}
\begin{tabular}{llll} 
In double precision we have: & \\
P & 1.265583562487 & -0.650585004466 & 0.197868699191 \\
\(Q\) & 1.000000000000 & -0.064342721236 & -0.028851461855 \\
ERROR \(=\) & -0.000026934190 & &
\end{tabular}

\section*{Chapter 4: Integration and Differentiation}

\section*{Routines}
4.1. Univariate Quadrature
Adaptive general-purpose endpoint singularities. ..... QDAGS ..... 862
Adaptive general purpose QDAG ..... 865
Adaptive general-purpose points of singularity ..... QDAGP ..... 869
Adaptive general-purpose infinite interval ..... QDAGI ..... 872
Adaptive weighted oscillatory (trigonometric) QDAWO ..... 875
Adaptive weighted Fourier (trigonometric) QDAWF ..... 879
Adaptive weighted algebraic endpoint singularities ..... QDAWS ..... 883
Adaptive weighted Cauchy principal value QDAWC ..... 886
Nonadaptive general purpose QDNG ..... 889
4.2. Multidimensional Quadrature
Two-dimensional quadrature (iterated integral) TWODQ ..... 891
Adaptive N -dimensional quadrature over a hyper-rectangle .QAND ..... 896
Integrates a function over a hyperrectangle using a quasi-Monte Carlo method QMC ..... 899
4.3. Gauss Rules and Three-term Recurrences
Gauss quadrature rule for classical weights GQRUL ..... 901
Gauss quadrature rule from recurrence coefficients GQRCF ..... 905
Recurrence coefficients for classical weights RECCF ..... 908
Recurrence coefficients from quadrature rule RECQR ..... 911
Fejer quadrature rule .FQRUL ..... 914
4.4. Differentiation
Approximation to first, second, or third derivative. DERIV ..... 918

\section*{Usage Notes}

\section*{Univariate Quadrature}

The first nine routines described in this chapter are designed to compute approximations to integrals of the form
\[
\int_{a}^{b} f(x) w(x) d x
\]

The weight function \(w\) is used to incorporate known singularities (either algebraic or logarithmic), to incorporate oscillations, or to indicate that a Cauchy principal value is desired. For general purpose integration, we recommend the use of QDAGS (even if no endpoint singularities are present). If more efficiency is desired, then the use of QDAG (or QDAG*) should be considered. These routines are organized as follows:
- \(\mathrm{w}=1\)
- QDAGS
- QDAG
- QDAGP
- QDAGI
- QDNG
- \(w(x)=\sin \omega x\) or \(w(x)=\cos \omega x\)
- QDAWO (for a finite interval)
- QDAWF (for an infinite interval)
- \(\quad w(x)=(x-a)^{\alpha}(b-x)^{\beta} \ln (x-a) \ln (b-x)\), where the \(\ln\) factors are optional
- QDAWS
- \(\quad w(x)=1 /(x-c) \quad\) Cauchy principal value
- QDAWC

The calling sequences for these routines are very similar. The function to be integrated is always F ; the lower and upper limits are, respectively, A and B. The requested absolute error \(\varepsilon\) is ERRABS, while the requested relative error \(\rho\) is ERRREL. These quadrature routines return two numbers of interest, namely, RESULT and ERREST, which are the approximate integral \(R\) and the error estimate \(E\), respectively. These numbers are related as follows:
\[
\left|\int_{a}^{b} f(x) w(x) d x-R\right| \leq E \leq \max \left\{\varepsilon, \rho\left|\int_{a}^{b} f(x) w(x) d x\right|\right\}
\]

One situation that occasionally arises in univariate quadrature concerns the approximation of integrals when only tabular data are given. The routines described above do not directly address this question. However, the standard method for handling this problem is first to interpolate the data and then to integrate the interpolant. This can be accomplished by using the IMSL spline
interpolation routines described in Chapter 3, "Interpolation and Apprximation", with one of the integration routines CSINT, BSINT, or PPITG.

\section*{Multivariate Quadrature}

Two routines are described in this chapter that are of use in approximating certain multivariate integrals. In particular, the routine TWODQ returns an approximation to an iterated two-dimensional integral of the form
\[
\int_{a}^{b} \int_{g(x)}^{h(x)} f(x, y) d y d x
\]

The second routine, QAND, returns an approximation to the integral of a function of \(n\) variables over a hyper-rectangle
\[
\int_{a_{1}}^{b_{1}} \ldots \int_{a_{n}}^{b_{n}} f\left(x_{1}, \ldots, x_{n}\right) d x_{n} \ldots d x_{1}
\]

If one has two- or three-dimensional tensor-product tabular data, use the IMSL spline interpolation routines BS2IN or BS3IN, followed by the IMSL spline integration routines BS2IG and BS3IG that are described in Chapter 3, Interpolation and Approximation.

\section*{Gauss rules and three-term recurrences}

The routines described in this section deal with the constellation of problems encountered in Gauss quadrature. These problems arise when quadrature formulas, which integrate polynomials of the highest degree possible, are computed. Once a member of a family of seven weight functions is specified, the routine GQRUL produces the points \(\left\{x_{i}\right\}\) and weights \(\left\{w_{i}\right\}\) for \(i=1, \ldots, N\) that satisfy
\[
\int_{a}^{b} f(x) w(x) d x=\sum_{i=1}^{N} f\left(x_{i}\right) w_{i}
\]
for all functions \(f\) that are polynomials of degree less than \(2 N\). The weight functions \(w\) may be selected from the following table:
\begin{tabular}{|c|c|c|}
\hline\(w(x)\) & Interval & Name \\
\hline 1 & \((-1,1)\) & Legendre \\
\(1 / \sqrt{1-x^{2}}\) & \((-1,1)\) & Chebyshev 1st kind \\
\(\sqrt{1-x^{2}}\) & \((-1,1)\) & Chebyshev 2nd kind \\
\(e^{-x^{2}}\) & \((-\infty, \infty)\) & Hermite \\
\((1+x)^{\alpha}(1-x)^{\beta}\) & \((-1,1)\) & Jacobi \\
\(e^{-x} x^{\alpha}\) & \((0, \infty)\) & Generalized Laguerre \\
\(1 / \cosh (x)\) & \((-\infty \infty)\) & Hyperbolic cosine \\
\hline
\end{tabular}

Where permissible, GQRUL will also compute Gauss-Radau and Gauss-Lobatto quadrature rules. The routine RECCF produces the three-term recurrence relation for the monic orthogonal polynomials with respect to the above weight functions.

Another routine, \(G Q R C F\), produces the Gauss, Gauss-Radau, or Gauss-Lobatto quadrature rule from the three-term recurrence relation. This means Gauss rules for general weight functions may be obtained if the three-term recursion for the orthogonal polynomials is known. The routine \(R E C Q R\) is an inverse to \(G Q R C F\) in the sense that it produces the recurrence coefficients given the Gauss quadrature formula.

The last routine described in this section, \(F Q R U L\), generates the Fejér quadrature rules for the following family of weights:
\[
\begin{aligned}
& w(x)=1 \\
& w(x)=1 /(x-\alpha) \\
& w(x)=(b-x)^{\alpha}(x-a)^{\beta} \\
& w(x)=(b-x)^{\alpha}(x-a)^{\beta} \ln (x-a) \\
& w(x)=(b-x)^{\alpha}(x-a)^{\beta} \ln (b-x)
\end{aligned}
\]

\section*{Numerical differentiation}

We provide one routine, DERIV, for numerical differentiation. This routine provides an estimate for the first, second, or third derivative of a user-supplied function.

\section*{QDAGS}

Integrates a function (which may have endpoint singularities).

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
A- Lower limit of integration. (Input)
\(\boldsymbol{B}\) - Upper limit of integration. (Input)
RESULT — Estimate of the integral from A to B of F . (Output)

\section*{Optional Required Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . e-3\) for single precision and 1.d-8 for double precision.
\(\boldsymbol{E R R R E L}\) - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . e-3\) for single precision and 1.d-8 for double precision.
ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAGS (F, A, B, RESULT [, ...])
Specific: The specific interface names are S_QDAGS and D_QDAGS.

\section*{FORTRAN 77 Interface}

Single: CALL QDAGS (F, A, B, ERRABS, ERRREL, RESULT, ERREST)
Double: The double precision name is DQDAGS.

\section*{Description}

The routine \(Q D A G S\) is a general-purpose integrator that uses a globally adaptive scheme to reduce the absolute error. It subdivides the interval \([A, B]\) and uses a 21-point Gauss-Kronrod rule to estimate the integral over each subinterval. The error for each subinterval is estimated by comparison with the 10 -point Gauss quadrature rule. This routine is designed to handle functions with endpoint singularities. However, the performance on functions, which are well-behaved at the endpoints, is quite good also. In addition to the general strategy described in QDAG, this routine uses an extrapolation procedure known as the \(\varepsilon\)-algorithm. The routine QDAGS is an implementation of the routine QAGS, which is fully documented by Piessens et al. (1983). Should QDAGS fail to produce acceptable results, then either IMSL routines QDAG or QDAG* may be appropriate. These routines are documented in this chapter.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AGS / DQ2AGS. The reference is
```

CALL Q2AGS (F, A, B, ERRABS, ERRREL, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD)

```

The additional arguments are as follows:
MAXSUB - Number of subintervals allowed. (Input)
A value of 500 is used by QDAGS.
NEVAL - Number of evaluations of F. (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. (Output)
Let \(k\) be
NSUBIN \(\quad\) if \(\operatorname{NSUBIN} \leq(\) MAXSUB \(/ 2+2)\);
MAXSUB +1 - NSUBIN otherwise.
The first \(k\) locations contain pointers to the error estimates over the subintervals such that ELIST (IORD(1)), \(\ldots, \operatorname{ELIST}(\operatorname{IORD}(k))\) form a decreasing sequence.
2. Informational errors
Type Code
\(4 \quad 1 \quad\) The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
34 Roundoff error in the extrapolation table, preventing the requested tolerance from being achieved, has been detected.
45 Integral is probably divergent or slowly convergent.
3. If EXACT is the exact value, QDAGS attempts to find RESULT such that \(\mid\) EXACT \(-\operatorname{RESULT} \mid \leq \max (E R R A B S\), ERRREL * \(|E X A C T|)\). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{1} \ln (x) x^{-1 / 2} d x=-4
\]
is estimated. The values of the actual and estimated error are machine dependent.
```

USE QDAGS_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
REAL A, ABS, B, ERRABS, ERREST, ERROR, ERRREL, EXACT, F, \&
RESULT
INTRINSIC ABS
EXTERNAL F
CALL UMACH (2, NOUT)
A = 0.0
B = 1.0

```
```

! Set error tolerances
ERRABS = 0.0
CALL QDAGS (F, A, B, RESULT, ERRABS=ERRABS, ERREST=ERREST)
EXACT = -4.0
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
!
REAL FUNCTION F (X)
REAL X
REAL ALOG, SQRT
INTRINSIC ALOG, SQRT
F = ALOG(X)/SQRT(X)
RETURN
END

```

\section*{Output}
```

Computed = -4.000 Exact = -4.000

```
Error estimate \(=1.519 \mathrm{E}-04 \quad\) Error \(=2.098 \mathrm{E}-05\)

\section*{QDAG}

Integrates a function using a globally adaptive scheme based on Gauss-Kronrod rules.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.

A- Lower limit of integration. (Input)
B- Upper limit of integration. (Input)
RESULT - Estimate of the integral from A to B of F . (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS = 1.e-3 for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . e-3\) for single precision and 1.d-8 for double precision.

IRULE - Choice of quadrature rule. (Input)
Default: IRULE \(=2\).
The Gauss-Kronrod rule is used with the following points:
IRULE Points
\(1 \quad 7-15\)
\(2 \quad 10-21\)
\(3 \quad 15-31\)
\(4 \quad 20-41\)
\(5 \quad 25-51\)
\(6 \quad 30-61\)

IRULE \(=2\) is recommended for most functions. If the function has a peak singularity, use IRULE \(=1\). If the function is oscillatory, use \(\operatorname{IRULE}=6\).

ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAG (F, A, B, RESULT [, ...])

Specific: The specific interface names are S_QDAG and D_QDAG.

\section*{FORTRAN 77 Interface}

Single: CALL QDAG (F, A, B, ERRABS, ERRREL, IRULE, RESULT, ERREST)
Double: The double precision name is DQDAG.

\section*{Description}

The routine \(Q D A G\) is a general-purpose integrator that uses a globally adaptive scheme in order to reduce the absolute error. It subdivides the interval \([A, B]\) and uses a \((2 k+1)\)-point Gauss-Kronrod rule to estimate the integral over each subinterval. The error for each subinterval is estimated by comparison with the \(k\)-point Gauss quadrature rule. The subinterval with the largest estimated error is then bisected and the same procedure is applied to both halves. The bisection process is continued until either the error criterion is satisfied, roundoff error is detected, the subintervals become too small, or the maximum number of subintervals allowed is reached. The routine QDAG is based on the subroutine QAG by Piessens et al. (1983).

Should QDAG fail to produce acceptable results, then one of the IMSL routines QDAG* may be appropriate. These routines are documented in this chapter.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AG/DQ2AG. The reference is:

CALL Q2AG (F, A, B, ERRABS, ERRREL, IRULE, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD)

The additional arguments are as follows:
MAXSUB - Number of subintervals allowed. (Input)
A value of 500 is used by QDAG.
NEVAL - Number of evaluations of F . (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. (Output) Let K be NSUBIN if nSUBIN.LE.(MAXSUB/2 +2 ), MAXSUB \(+1-\) NSUBIN otherwise. The first K locations contain pointers to the error estimates over the corresponding subintervals, such that \(\operatorname{ELIST}(\operatorname{IORD}(1)), \ldots, \operatorname{ELIST}(\operatorname{IORD}(\mathrm{K}))\) form a decreasing sequence.
2. Informational errors

Type Code
41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
3. If EXACT is the exact value, QDAG attempts to find RESULT such that ABS(EXACT - RESULT). LE.MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{2} x e^{x} d x=e^{2}+1
\]
is estimated. Since the integrand is not oscillatory, IRULE \(=1\) is used. The values of the actual and estimated error are machine dependent.
```

USE QDAG_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER IRULE, NOUT
REAL A, ABS, B, ERRABS, ERREST, ERROR, EXACT, EXP, \&
F, RESULT
INTRINSIC ABS, EXP
EXTERNAL F
! Get output unit number
CALL UMACH (2, NOUT)
A = 0.0
B = 2.0
ERRABS = 0.0
Set error tolerances
Parameter for non-oscillatory
function
IRULE = 1
CALL QDAG (F, A, B, RESULT, ERRABS=ERRABS, IRULE=IRULE, ERREST=ERREST)
EXACT = 1.0 + EXP(2.0)
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
REAL FUNCTION F (X)
REAL X
REAL EXP
INTRINSIC EXP
F = X*EXP(X)
RETURN
END

```
!

\section*{Output}
\begin{tabular}{ll} 
Computed \(=8.389\) & Exact \(=8.389\) \\
Error estimate \(=5.000 \mathrm{E}-05\) & Error \(=9.537 \mathrm{E}-07\)
\end{tabular}

\section*{QDAGP}

Integrates a function with singularity points given.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
A- Lower limit of integration. (Input)
\(\boldsymbol{B}\) - Upper limit of integration. (Input)
POINTS - Array of length NPTS containing breakpoints in the range of integration. (Input) Usually these are points where the integrand has singularities.

RESULT - Estimate of the integral from A to B of F . (Output)

\section*{Optional Arguments}

NPTS - Number of break points given. (Input)
Default: NPTS \(=\operatorname{size}(\) POINTS, 1\()\).
ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . \mathrm{e}-3\) for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL = 1.e-3 for single precision and 1.d-8 for double precision.
ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAGP (F, A, B, POINTS, RESULT [, ...])
Specific: The specific interface names are S_QDAGP and D_QDAGP.

\section*{FORTRAN 77 Interface}

Single: CALL QDAGP (F, A, B, NPTS, POINTS, ERRABS, ERRREL, RESULT, ERREST)

Double: The double precision name is DQDAGP.

\section*{Description}

The routine QDAGP uses a globally adaptive scheme in order to reduce the absolute error. It initially subdivides the interval \([A, B]\) into NPTS +1 user-supplied subintervals and uses a 21-point Gauss-Kronrod rule to estimate the integral over each subinterval. The error for each subinterval is estimated by comparison with the 10 -point Gauss quadrature rule. This routine is designed to handle endpoint as well as interior singularities. In addition to the general strategy described in the IMSL routine \(Q D A G\), this routine employs an extrapolation procedure known as the \(\varepsilon\)-algorithm. The routine QDAGP is an implementation of the subroutine QAGP, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AGP/DQ2AGP. The reference is:

CALL Q2AGP (F, A, B, NPTS, POINTS, ERRABS, ERRREL, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD, LEVEL, WK, IWK)

The additional arguments are as follows:
MAXSUB - Number of subintervals allowed. (Input)
A value of 450 is used by QDAGP.
NEVAL - Number of evaluations of \(F\). (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. (Output)
Let K be NSUBIN if NSUBIN.LE.(MAXSUB/2 +2 ), MAXSUB \(+1-\) NSUBIN otherwise. The first \(k\) locations contain pointers to the error estimates over the subintervals, such that \(\operatorname{ELIST}(\operatorname{IORD}(1)), \ldots, \operatorname{ELIST}(\operatorname{IORD}(\mathrm{K}))\) form a decreasing sequence.

LEVEL - Array of length MAXSUB, containing the subdivision levels of the subinterval. (Output)
That is, if (AA, BB) is a subinterval of (P1, P2) where P1 as well as P2 is a
user-provided break point or integration limit, then ( \(\mathrm{AA}, \mathrm{BB}\) ) has level L if \(\mathrm{ABS}(\mathrm{BB}-\mathrm{AA})=\mathrm{ABS}(\mathrm{P} 2-\mathrm{P} 1)\) * \(2 * *(-\mathrm{L})\).
\(W K-\) Work array of length NPTS +2.
\(I W K\) - Work array of length NPTS +2.
2. Informational errors

Type Code
41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
34 Roundoff error in the extrapolation table, preventing the requested tolerance from being achieved, has been detected.
45 Integral is probably divergent or slowly convergent.
3. If EXACT is the exact value, QDAGP attempts to find RESULT such that

ABS(EXACT - RESULT). LE.MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{3} x^{3} \ln \left|\left(x^{2}-1\right)\left(x^{2}-2\right)\right| d x=61 \ln 2+\frac{77}{4} \ln 7-27
\]
is estimated. The values of the actual and estimated error are machine dependent. Note that this subroutine never evaluates the user-supplied function at the user-supplied breakpoints.
```

USE QDAGP_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT, NPTS
REAL A, ABS, ALOG, B, ERRABS, ERREST, ERROR, ERRREL, \&
EXACT, F, POINTS(2), RESULT, SQRT
INTRINSIC ABS, ALOG, SQRT
EXTERNAL F
CALL UMACH (2, NOUT)
A = 0.0
B = 3.0
ERRABS = 0.0
ERRREL = 0.01
NPTS = 2
POINTS(1) = 1.0

```
```

POINTS(2) = SQRT(2.0)
CALL QDAGP (F, A, B, POINTS, RESULT, ERRABS=ERRABS, ERRREL=ERRREL, \&
ERREST=ERREST)
! Print results
EXACT = 61.0*ALOG(2.0) + 77.0/4.0*ALOG(7.0) - 27.0
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
!
END
!
REAL FUNCTION F (X)
REAL X
REAL ABS, ALOG
INTRINSIC ABS, ALOG
F = X** 3*ALOG (ABS ((X* X-1.0)* (X*X-2.0)))
RETURN
END

```

\section*{Output}

Computed \(=52.741 \quad\) Exact \(=52.741\)
Error estimate \(=5.062 \mathrm{E}-01 \quad\) Error \(=6.104 \mathrm{E}-04\)

\section*{QDAGI}

Integrates a function over an infinite or semi-infinite interval.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(F(X)\), where

X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
BOUND - Finite bound of the integration range. (Input) Ignored if INTERV \(=2\).

INTERV - Flag indicating integration interval. (Input)
INTERV Interval
\(-1 \quad(-\infty\), BOUND \()\)

1 (BOUND, \(+\infty\) )

2
\((-\infty,+\infty)\)

RESULT - Estimate of the integral from A to B of F . (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: \(E R R A B S=1 . e-3\) for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . e-3\) for single precision and 1.d-8 for double precision.
ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAGI (F, BOUND, INTERV, RESULT [, ...])

Specific: The specific interface names are S_QDAGI and D_QDAGI.

\section*{FORTRAN 77 Interface}

Single: CALL QDAGI (F, BOUND, INTERV, ERRABS, ERRREL, RESULT, ERREST)

Double: The double precision name is DQDAGI.

\section*{Description}

The routine \(Q D A G I\) uses a globally adaptive scheme in an attempt to reduce the absolute error. It initially transforms an infinite or semi-infinite interval into the finite interval [ 0,1\(]\). Then, QDAGI uses a 21-point Gauss-Kronrod rule to estimate the integral and the error. It bisects any interval with an unacceptable error estimate and continues this process until termination. This routine is designed to handle endpoint singularities. In addition to the general strategy described in QDAG, this subroutine employs an extrapolation procedure known as the \(\varepsilon\)-algorithm. The routine QDAGI is an implementation of the subroutine QAGI, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AGI / DQ2AGI. The reference is

CALL Q2AGI (F, BOUND, INTERV, ERRABS, ERRREL, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD)

The additional arguments are as follows:

MAXSUB - Number of subintervals allowed. (Input)
A value of 500 is used by QDAGI.

NEVAL - Number of evaluations of F. (Output)
NSUBIN - Number of subintervals generated. (Output)

ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

\section*{IORD - Array of length MAXSUB. (Output)}

Let \(k\) be NSUBIN if NSUBIN .LE. (MAXSUB/2 + 2), MAXSUB + 1 - NSUBIN otherwise. The first k locations contain pointers to the error estimates over the subintervals, such that ELIST (IORD (1)) , ..., ELIST (IORD (K) ) form a decreasing sequence.
2. Informational errors

Type Code
\(4 \quad 1 \quad\) The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
34 Roundoff error in the extrapolation table, preventing the requested tolerance from being achieved, has been detected.
45 Integral is divergent or slowly convergent.
3. If EXACT is the exact value, QDAGI attempts to find RESULT such that ABS(EXACT - RESULT). LE. MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{\infty} \frac{\ln (x)}{1+(10 x)^{2}} d x=\frac{-\pi \ln (10)}{20}
\]
is estimated. The values of the actual and estimated error are machine dependent. Note that we have requested an absolute error of 0 and a relative error of .001 . The effect of these requests, as documented in Comment 3 above, is to ignore the absolute error requirement.
```

    USE QDAGI_INT
    USE UMACH INT
    USE CONST-
    IMPLICIT NONE
    INTEGER INTERV, NOUT
    REAL ABS, ALOG, BOUND, ERRABS, ERREST, ERROR, &
        ERRREL, EXACT, F, PI, RESULT
    INTRINSIC ABS, ALOG
    EXTERNAL F
    CALL UMACH (2, NOUT)
    BOUND = 0.0
    INTERV = 1
    ERRABS = 0.0
    CALL QDAGI (F, BOUND, INTERV, RESULT, ERRABS=ERRABS, &
        ERREST=ERREST)
            Print results
    PI = CONST('PI')
    EXACT = -PI*ALOG(10.)/20.
    ERROR = ABS (RESULT-EXACT)
    WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
    9 9 9 9 9 ~ F O R M A T ~ ( ' ~ C o m p u t e d ~ = ' , ~ F 8 . 3 , ~ 1 3 X , ~ ' ~ E x a c t ~ = ' , ~ F 8 . 3 / / ' ~ E r r o r ~ ' , ~ \& ~
'estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
!
REAL FUNCTION F (X)
REAL X
REAL ALOG
INTRINSIC ALOG
F = ALOG (X)/(1.+(10.*X)**2)
RETURN
END

```

\section*{Output}
\begin{tabular}{ll} 
Computed \(=-0.362\) & Exact \(=-0.362\) \\
Error estimate \(=2.652 \mathrm{E}-06\) & Error \(=5.960 \mathrm{E}-08\)
\end{tabular}

\section*{QDAWO}

Integrates a function containing a sine or a cosine.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied function to be integrated. The form is \(\mathrm{F}(\mathrm{x})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
\(\boldsymbol{A}\) - Lower limit of integration. (Input)
B - Upper limit of integration. (Input)
IWEIGH - Type of weight function used. (Input)

\section*{IWEIGH Weight}
\(1 \cos (O M E G A * x)\)
2 SIN(OMEGA * X)
OMEGA - Parameter in the weight function. (Input)
RESULT - Estimate of the integral from A to B of F * WEIGHT. (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . e-3\) for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . \mathrm{e}-3\) for single precision and 1.d-8 for double precision.
ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAWO (F, A, B, IWEIGH, OMEGA, RESULT [, ...])
Specific: The specific interface names are S_QDAWO and D_QDAWO.

\section*{FORTRAN 77 Interface}

Single: CALL QDAWO (F, A, B, IWEIGH, OMEGA, ERRABS, ERRREL, RESULT,
Double: The double precision name is DQDAWO.

\section*{Description}

The routine QDAWO uses a globally adaptive scheme in an attempt to reduce the absolute error. This routine computes integrals whose integrands have the special form \(w(x) f(x)\), where \(w(x)\) is either \(\cos \omega x\) or \(\sin \omega x\). Depending on the length of the subinterval in relation to the size of \(\omega\), either a modified Clenshaw-Curtis procedure or a Gauss-Kronrod \(7 / 15\) rule is employed to approximate the integral on a subinterval. In addition to the general strategy described for the IMSL routine QDAG, this subroutine uses an extrapolation procedure known as the \(\varepsilon\)-algorithm.

The routine QDAWO is an implementation of the subroutine QAWO, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AWO/DQ2AWO. The reference is:
```

CALL Q2AWO (F, A, B, IWEIGH, OMEGA, ERRABS, ERRREL, RESULT, ERREST, MAXSUB,
MAXCBY, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD, NNLOG, WK)

```

The additional arguments are as follows:
MAXSUB - Maximum number of subintervals allowed. (Input)
A value of 390 is used by QDAWO.
MAXCBY - Upper bound on the number of Chebyshev moments which can be stored. That is, for the intervals of lengths ABS \((B-A) * 2 * *(-L), L=0\), \(1, \ldots\), MAXCBY -2 , MAXCBY. GE.1. The routine QDAWO uses 21. (Input)

NEVAL - Number of evaluations of F . (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. Let K be NSUBIN if NSUBIN. LE . (MAXSUB/2 + 2), MAXSUB + 1 - NSUBIN otherwise. The first K locations contain pointers to the error estimates over the subintervals, such that ELIST(IORD(1)), ..., ELIST(IORD(K)) form a decreasing sequence. (Output)

NNLOG - Array of length MAXSUB containing the subdivision levels of the subintervals, i.e. \(\operatorname{NNLOG}(I)=L\) means that the subinterval numbered \(I\) is of length ABS (B - A) * (1- L). (Output)

WK - Array of length 25 * MAXCBY. (Workspace)
2. Informational errors
Type Code

41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
34 Roundoff error in the extrapolation table, preventing the requested tolerances from being achieved, has been detected.
45 Integral is probably divergent or slowly convergent.
3. If EXACT is the exact value, QDAWO attempts to find RESULT such that ABS(EXACT - RESULT) . LE. MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{1} \ln (x) \sin (10 \pi x) d x
\]
is estimated. The values of the actual and estimated error are machine dependent. Notice that the \(\log\) function is coded to protect for the singularity at zero.
```

USE QDAWO_INT
USE UMACH INT
USE CONST_INT
IMPLICIT NONE
INTEGER IWEIGH, NOUT
REAL A, ABS, B, ERRABS, ERREST, ERROR, \&
EXACT, F, OMEGA, PI, RESULT
INTRINSIC ABS
EXTERNAL F
CALL UMACH (2, NOUT)
A = 0.0
B = 1.0
IWEIGH = 2
PI = CONST('PI')
OMEGA = 10.*PI
ERRABS = 0.0
CALL QDAWO (F, A, B, IWEIGH, OMEGA, RESULT, ERRABS=ERRABS, \&
ERREST=ERREST)
EXACT = -0.1281316
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&

```
```

        ' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
    END
    !
REAL FUNCTION F (X)
REAL X
REAL ALOG
INTRINSIC ALOG
IF (X .EQ. O.) THEN
F=0.0
ELSE
F = ALOG (X)
END IF
RETURN
END

```

\section*{Output}

Computed \(=-0.128 \quad\) Exact \(=-0.128\)
Error estimate \(=7.504 \mathrm{E}-05 \quad\) Error \(=5.260 \mathrm{E}-06\)

\section*{QDAWF}

Computes a Fourier integral.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUnCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
x - Independent variable. (Input)
F - The function value. (Output)
F must be declared external in the calling program.
A-Lower limit of integration. (Input)
IWEIGH - Type of weight function used. (Input)
IWEIGH Weight

1
COS(OMEGA * X)

2
SIN(OMEGA * X)
OMEGA - Parameter in the weight function. (Input)
RESULT - Estimate of the integral from A to infinity of F * WEIGHT. (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . e-3\) for single precision and 1.d-8 for double precision.

ERREST - Estimate of the absolute value of the error. (Output)
Default: ERREST \(=1 . \mathrm{e}-3\) for single precision and 1.d-8 for double precision.

\section*{FORTRAN 90 Interface}

Generic: CALL QDAWF (F, A, IWEIGH, OMEGA, RESULT [,...])

Specific: The specific interface names are S_QDAWF and D_QDAWF.

\section*{FORTRAN 77 Interface}

Single: CALL QDAWF (F, A, IWEIGH, OMEGA, ERRABS, RESULT, ERREST)
Double: The double precision name is DQDAWF.

\section*{Description}

The routine \(Q D A W F\) uses a globally adaptive scheme in an attempt to reduce the absolute error. This routine computes integrals whose integrands have the special form \(w(x) f(x)\), where \(w(x)\) is either \(\cos \omega x\) or \(\sin \omega x\). The integration interval is always semi-infinite of the form \([A, \infty]\). These Fourier integrals are approximated by repeated calls to the IMSL routine QDAWO followed by extrapolation. The routine QDAWF is an implementation of the subroutine QAWF, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AWF/DQ2AWF. The reference is:

CALL Q2AWF (F, A, IWEIGH, OMEGA, ERRABS, RESULT, ERREST, MAXCYL, MAXSUB, MAXCBY, NEVAL, NCYCLE, RSLIST, ERLIST, IERLST, NSUBIN, WK, IWK)

The additional arguments are as follows:
MAXSUB - Maximum number of subintervals allowed. (Input)
A value of 365 is used by QDAWF.
MAXCYL - Maximum number of cycles allowed. (Input) MAXCYL must be at least 3 . QDAWF uses 50 .

MAXCBY - Maximum number of Chebyshev moments allowed. (Input) QDAWF uses 21.

NEVAL - Number of evaluations of F . (Output)
NCYCLE - Number of cycles used. (Output)

RSLIST - Array of length MAXCYL containing the contributions to the integral over the interval \(\left(\mathrm{A}+(k-1) * \mathrm{C}, \mathrm{A}+k^{*} \mathrm{C}\right)\), for \(k=1, \ldots, \mathrm{NCYCLE}\). (Output) \(C=(2 * \operatorname{INT}(A B S(O M E G A))+1) * \operatorname{PI} / A B S(O M E G A)\).

ERLIST - Array of length MAXCYL containing the error estimates for the intervals defined in RSLIST. (Output)

IERLST - Array of length MAXCYL containing error flags for the intervals defined in RSLIST. (Output)

\section*{IERLST(K) Meaning}

1 The maximum number of subdivisions (MAXSUB) has been achieved on the \(k\)-th cycle.

2
Roundoff error prevents the desired accuracy from being achieved on the \(k\)-th cycle.

3
Extremely bad integrand behavior occurs at some points of the \(k\)-th cycle.

4
Integration procedure does not converge (to the desired accuracy) due to roundoff in the extrapolation procedure on the \(k\)-th cycle. It is assumed that the result on this interval is the best that can be obtained.

5
Integral over the \(k\)-th cycle is divergent or slowly convergent.
NSUBIN - Number of subintervals generated. (Output)
WK - Work array of length 4 * MAXSUB +25 * MAXCBY.
IWK - Work array of length 2 * MAXSUB.
2. Informational errors

Type Code
31 Bad integrand behavior occurred in one or more cycles.
42 Maximum number of cycles allowed has been reached.
33 Extrapolation table constructed for convergence acceleration of the series formed by the integral contributions of the cycles does not converge to the requested accuracy.
3. If EXACT is the exact value, QDAWF attempts to find RESULT such that ABS(EXACT - RESULT) .LE. ERRABS.

\section*{Example}

The value of
\[
\int_{0}^{\infty} x^{-1 / 2} \cos (\pi x / 2) d x=1
\]
is estimated. The values of the actual and estimated error are machine dependent. Notice that \(F\) is coded to protect for the singularity at zero.
```

USE QDAWF_INT
USE UMACH INT
USE CONST_INT
IMPLICIT NONE
INTEGER IWEIGH, NOUT
REAL A, ABS, ERRABS, ERREST, ERROR, EXACT, F, \&
OMEGA, PI, RESULT
INTRINSIC ABS
EXTERNAL F
CALL UMACH (2, NOUT)
A = 0.0
IWEIGH = 1
PI = CONST('PI')
OMEGA = PI/2.0
CALL QDAWF (F, A, IWEIGH, OMEGA, RESULT, ERREST=ERREST)
EXACT = 1.0
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
REAL FUNCTION F (X)
REAL X
REAL SQRT
INTRINSIC SQRT
IF (X .GT. O.O) THEN
F = 1.0/SQRT (X)
ELSE
F=0.0
END IF
RETURN
END

```
!

\section*{Output}
\begin{tabular}{ll} 
Computed \(=1.000\) & Exact \(=1.000\) \\
Error estimate \(=6.267 \mathrm{E}-04\) & Error \(=2.205 \mathrm{E}-06\)
\end{tabular}

\section*{QDAWS}

Integrates a function with algebraic-logarithmic singularities.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
A- Lower limit of integration. (Input)
\(\boldsymbol{B}\) - Upper limit of integration. (Input) \(B\) must be greater than \(A\)

IWEIGH - Type of weight function used. (Input)
IWEIGH Weight
1
\((\mathrm{X}-\mathrm{A}) * * A L P H A *(\mathrm{~B}-\mathrm{X}) * *\) BETAW

2
\[
(X-A) * * A L P H A *(B-X) * * \text { BETAW } * \operatorname{LOG}(X-A)
\]

3
\[
(X-A) * * A L P H A *(B-X) * * B E T A W * \operatorname{LOG}(B-X)
\]

4
\[
(X-A) * * A L P H A *(B-X) * * B E T A W * \operatorname{LOG}(X-A) * \operatorname{LOG}(B-x)
\]

ALPHA - Parameter in the weight function. (Input)
ALPHA must be greater than -1.0 .
BETAW - Parameter in the weight function. (Input)
BETAW must be greater than -1.0 .
RESULT — Estimate of the integral from A to B of F * WEIGHT. (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . e-3\) for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . \mathrm{e}-3\) for single precision and \(1 . \mathrm{d}-8\) for double precision.
ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAWS (F, A, B, IWEIGH, ALPHA, BETAW, RESULT [,...] )
Specific: The specific interface names are S_QDAWS and D_QDAWS.

\section*{FORTRAN 77 Interface}

Single: CALL QDAWS (F, A, B, IWEIGH, ALPHA, BETAW, ERRABS, ERRREL, RESULT, ERREST)

Double: The double precision name is DQDAWS.

\section*{Description}

The routine QDAWS uses a globally adaptive scheme in an attempt to reduce the absolute error. This routine computes integrals whose integrands have the special form \(w(x) f(x)\), where \(w(x)\) is a weight function described above. A combination of modified Clenshaw-Curtis and Gauss-Kronrod formulas is employed. In addition to the general strategy described for the IMSL routine QDAG, this routine uses an extrapolation procedure known as the \(\varepsilon\)-algorithm. The routine QDAWS is an implementation of the routine QAWS, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AWS/DQ2AWS. The reference is

CALL Q2AWS (F, A, B, IWEIGH, ALPHA, BETAW, ERRABS, ERRREL, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD)

The additional arguments are as follows:
MAXSUB - Maximum number of subintervals allowed. (Input)
A value of 500 is used by QDAWS.
NEVAL - Number of evaluations of \(F\). (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. Let \(k\) be NSUBIN if NSUBIN. LE. (MAXSUB/2 +2 ), MAXSUB \(+1-\) NSUBIN otherwise. The first \(k\) locations contain pointers to the error estimates over the subintervals, such that \(\operatorname{ELIST}(\operatorname{IORD}(1)), \ldots, \operatorname{ELIST}(\operatorname{IORD}(k))\) form a decreasing sequence. (Output)
2. Informational errors

Type Code
41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
3. If EXACT is the exact value, QDAWS attempts to find RESULT such that ABS(EXACT - RESULT).LE.MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The value of
\[
\int_{0}^{1}[(1+x)(1-x)]^{1 / 2} x \ln (x) d x=\frac{3 \ln (2)-4}{9}
\]
is estimated. The values of the actual and estimated error are machine dependent.
```

USE QDAWS_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER IWEIGH, NOUT
REAL A, ABS, ALOG, ALPHA, B, BETAW, ERRABS, ERREST, ERROR, \&
EXACT, F, RESULT
INTRINSIC ABS, ALOG
EXTERNAL F
! Get output unit number
CALL UMACH (2, NOUT)
A = 0.0
B = 1.0
ALPHA = 1.0
BETAW = 0.5
IWEIGH = 2
Set error tolerances
ERRABS = 0.0
CALL QDAWS (F, A, B, IWEIGH, ALPHA, BETAW, RESULT, \&
ERRABS=ERRABS, ERREST=ERREST)

```
```

! Print results
EXACT = (3.*ALOG (2.)-4.)/9.
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
!
REAL FUNCTION F (X)
REAL X
REAL SQRT
INTRINSIC SQRT
F = SQRT (1.0+X)
RETURN
END

```

\section*{Output}

Computed \(=-0.213 \quad\) Exact \(=-0.213\)

Error estimate \(=1.261 \mathrm{E}-08 \quad\) Error \(=2.980 \mathrm{E}-08\)

\section*{QDAWC}

Integrates a function \(f(x) /(x-c)\) in the Cauchy principal value sense.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied function to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
x - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
\(\boldsymbol{A}\) - Lower limit of integration. (Input)
B - Upper limit of integration. (Input)
C - Singular point. (Input)
C must not equal A or B .
\(\boldsymbol{R E S U L T}\) - Estimate of the integral from A to B of \(\mathrm{F}(\mathrm{x}) /(\mathrm{x}-\mathrm{c})\). (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . \mathrm{e}-3\) for single precision and \(1 . \mathrm{d}-8\) for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERREL \(=1 . \mathrm{e}-3\) for single precision and 1.d-8 for double precision.

ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDAWC (F, A, B, C, RESULT [,...])

Specific: The specific interface names are S_QDAWC and D_QDAWC.

\section*{FORTRAN 77 Interface}

Single: CALL QDAWC (F, A, B, C, ERRABS, ERRREL, RESULT, ERREST)
Double: The double precision name is DQDAWC.

\section*{Description}

The routine \(Q D A W C\) uses a globally adaptive scheme in an attempt to reduce the absolute error. This routine computes integrals whose integrands have the special form \(w(x) f(x)\), where \(w(x)=1 /(x-c)\). If \(c\) lies in the interval of integration, then the integral is interpreted as a Cauchy principal value. A combination of modified Clenshaw-Curtis and Gauss-Kronrod formulas are employed. In addition to the general strategy described for the IMSL routine QDAG, this routine uses an extrapolation procedure known as the \(\varepsilon\)-algorithm. The routine QDAWC is an implementation of the subroutine QAWC, which is fully documented by Piessens et al. (1983).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of Q2AWC / DQ2AWC. The reference is:

CALL Q2AWC ( \(\mathrm{F}, \mathrm{A}, \mathrm{B}, \mathrm{C}\), ERRABS, ERRREL, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD)

The additional arguments are as follows:

MAXSUB - Number of subintervals allowed. (Input)
A value of 500 is used by QDAWC.
NEVAL - Number of evaluations of F. (Output)
NSUBIN - Number of subintervals generated. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. (Output)
Let \(k\) be NSUBIN if NSUBIN.LE.(MAXSUB \(/ 2+2\) ), MAXSUB \(+1-\) NSUBIN otherwise. The first \(k\) locations contain pointers to the error estimates over the subintervals, such that \(\operatorname{ELIST}(\operatorname{IORD}(1)), \ldots, \operatorname{ELIST}(\operatorname{IORD}(k))\) form a decreasing sequence.
2. Informational errors
Type Code

41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
3. If EXACT is the exact value, QDAWC attempts to find RESULT such that ABS(EXACT - RESULT) . LE. MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

The Cauchy principal value of
\[
\int_{-1}^{5} \frac{1}{x\left(5 x^{3}+6\right)} d x=\frac{\ln (125 / 631)}{18}
\]
is estimated. The values of the actual and estimated error are machine dependent.
```

USE QDAWC INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
REAL A, ABS, ALOG, B, C, ERRABS, ERREST, ERROR, EXACT, \&
F, RESULT
INTRINSIC ABS, ALOG
EXTERNAL F
CALL UMACH (2, NOUT)
A = -1.0
B = 5.0
C = 0.0
Set error tolerances
ERRABS = 0.0

```
!
!
```

    CALL QDAWC (F, A, B, C, RESULT, ERRABS=ERRABS, ERREST=ERREST)
    !
EXACT = ALOG(125./631.)/18.
ERROR = 2*ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
!
REAL FUNCTION F (X)
REAL X
F=1.0/(5.* X** 3+6.0)
RETURN
END

```

\section*{Output}

Computed \(=-0.090 \quad\) Exact \(=-0.090\)

Error estimate \(=2.022 \mathrm{E}-06 \quad\) Error \(=2.980 \mathrm{E}-08\)

\section*{QDNG}

Integrates a smooth function using a nonadaptive rule.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X})\), where
X - Independent variable. (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
\(\boldsymbol{A}\) - Lower limit of integration. (Input)
B - Upper limit of integration. (Input)

RESULT - Estimate of the integral from A to B of F . (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . e-3\) for single precision and 1.d-8 for double precision.

ERRREL - Relative accuracy desired. (Input)
Default: ERRREL = 1.e-3 for single precision and 1.d-8 for double precision.

ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QDNG (F, A, B, RESULT [,...])
Specific: The specific interface names are S_QDNG and D_QDNG.

\section*{FORTRAN 77 Interface}

Single: CALL QDNG (F, A, B, ERRABS, ERRREL, RESULT, ERREST)

Double: The double precision name is DQDNG.

\section*{Description}

The routine QDNG is designed to integrate smooth functions. This routine implements a nonadaptive quadrature procedure based on nested Paterson rules of order 10, 21, 43, and 87 . These rules are positive quadrature rules with degree of accuracy \(19,31,64\), and 130 , respectively. The routine QDNG applies these rules successively, estimating the error, until either the error estimate satisfies the user-supplied constraints or the last rule is applied. The routine QDNG is based on the routine QNG by Piessens et al. (1983).

This routine is not very robust, but for certain smooth functions it can be efficient. If QDNG should not perform well, we recommend the use of the IMSL routine QDAGS.

\section*{Comments}
1. Informational error

Type Code
41 The maximum number of steps allowed have been taken. The integral is too difficult for QDNG.
2. If EXACT is the exact value, QDNG attempts to find RESULT such that ABS(EXACT - RESULT). LE. MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.
3. This routine is designed for efficiency, not robustness. If the above error is encountered, try QDAGS.

\section*{Example}

The value of
\[
\int_{0}^{2} x e^{x} d x=e^{2}+1
\]
is estimated. The values of the actual and estimated error are machine dependent.
```

USE QDNG_INT
USE UMACH_INT

```
```

IMPLICIT NONE
INTEGER NOUT
REAL A, ABS, B, ERRABS, ERREST, ERROR, EXACT, EXP, \&
F, RESULT
INTRINSIC ABS, EXP
EXTERNAL F
CALL UMACH (2, NOUT)
A = 0.0
B = 2.0
! Set error tolerances
ERRABS = 0.0
CALL QDNG (F, A, B, RESULT, ERRABS=ERRABS, ERREST=ERREST)
EXACT = 1.0 + EXP(2.0)
ERROR = ABS (RESULT-EXACT)
WRITE (NOUT,99999) RESULT, EXACT, ERREST, ERROR
99999 FORMAT (' Computed =', F8.3, 13X, ' Exact =', F8.3, /, /, \&
' Error estimate =', 1PE10.3, 6X, 'Error =', 1PE10.3)
END
!
REAL FUNCTION F (X)
REAL X
REAL EXP
INTRINSIC EXP
F = X* EXP(X)
RETURN
END

```

\section*{Output}
\begin{tabular}{ll} 
Computed \(=8.389\) & Exact \(=8.389\) \\
Error estimate \(=5.000 \mathrm{E}-05\) & Error \(=9.537 \mathrm{E}-07\)
\end{tabular}

\section*{TWODQ}

Computes a two-dimensional iterated integral.

\section*{Required Arguments}

F - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{X}, \mathrm{Y})\), where
\(X-\) First argument of F . (Input)
\(\mathrm{Y}-\) Second argument of F . (Input)
F - The function value. (Output)
F must be declared EXTERNAL in the calling program.
A - Lower limit of outer integral. (Input)
B - Upper limit of outer integral. (Input)
\(\boldsymbol{G}\) - User-supplied FUNCTION to evaluate the lower limits of the inner integral.
The form is \(G(X)\), where
X - Only argument of G . (Input)
G - The function value. (Output)
G must be declared EXTERNAL in the calling program.
\(\boldsymbol{H}\) - User-supplied FUNCTION to evaluate the upper limits of the inner integral. The form is \(\mathrm{H}(\mathrm{X})\), where
\(X\) - Only argument of \(H\). (Input)
H - The function value. (Output)
H must be declared EXTERNAL in the calling program.
RESULT - Estimate of the integral from A to B of F . (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS = 1.e-3 for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input) Default: ERRREL \(=1 . e-3\) for single precision and 1.d-8 for double precision.

IRULE --- Choice of quadrature rule. (Input)
Default: IRULE \(=2\).
The Gauss-Kronrod rule is used with the following points:

\section*{IRULE Points}
1 7-15

2 10-21
3 15-31
\(4 \quad 20-41\)
\(5 \quad 25-51\)
6 30-61
If the function has a peak singularity, use IRULE \(=1\). If the function is oscillatory, use IRULE \(=6\).

ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL TWODQ (F, A, B, G, H, RESULT [, ...])
Specific: The specific interface names are S_TWODQ and D_TWODQ.

\section*{FORTRAN 77 Interface}

Single: CALL TWODQ (F, A, B, G, H, ERRABS, ERRREL, IRULE, RESULT, ERREST)

Double: The double precision name is DTWODQ.

\section*{Description}

The routine TWODQ approximates the two-dimensional iterated integral
\[
\int_{a}^{b} \int_{g(x)}^{h(x)} f(x, y) d y d x
\]
with the approximation returned in RESULT. An estimate of the error is returned in ERREST. The approximation is achieved by iterated calls to QDAG. Thus, this algorithm will share many of the characteristics of the routine QDAG. As in QDAG, several options are available. The absolute and relative error must be specified, and in addition, the Gauss-Kronrod pair must be specified (IRULE). The lower-numbered rules are used for less smooth integrands while the higher-order rules are more efficient for smooth (oscillatory) integrands.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of T2ODQ/DT2ODQ. The reference is:

CALL T2ODQ (F, A, B, G, H, ERRABS, ERRREL, IRULE, RESULT, ERREST, MAXSUB, NEVAL, NSUBIN, ALIST, BLIST, RLIST, ELIST, IORD, WK, IWK)

The additional arguments are as follows:
MAXSUB - Number of subintervals allowed. (Input)
A value of 250 is used by TWODQ.
NEVAL - Number of evaluations of F. (Output)
NSUBIN - Number of subintervals generated in the outer integral. (Output)
ALIST - Array of length MAXSUB containing a list of the NSUBIN left endpoints for the outer integral. (Output)

BLIST - Array of length MAXSUB containing a list of the NSUBIN right endpoints for the outer integral. (Output)

RLIST - Array of length MAXSUB containing approximations to the NSUBIN integrals over the intervals defined by ALIST, BLIST, pertaining only to the outer integral. (Output)

ELIST - Array of length MAXSUB containing the error estimates of the NSUBIN values in RLIST. (Output)

IORD - Array of length MAXSUB. (Output)
Let \(K\) be NSUBIN if NSUBIN.LE.(MAXSUB \(/ 2+2\) ), MAXSUB \(+1-\) NSUBIN otherwise. Then the first K locations contain pointers to the error estimates over the corresponding subintervals, such that \(\operatorname{ELIST}(\operatorname{IORD}(1)), \ldots, \operatorname{ELIST}(\operatorname{IORD}(\mathrm{K}))\) form a decreasing sequence.
\(\boldsymbol{W K}\) - Work array of length 4 * MAXSUB, needed to evaluate the inner integral.
IWK - Work array of length MAXSUB, needed to evaluate the inner integral.
2. Informational errors
Type Code

41 The maximum number of subintervals allowed has been reached.
32 Roundoff error, preventing the requested tolerance from being achieved, has been detected.
33 A degradation in precision has been detected.
3. If EXACT is the exact value, TWODQ attempts to find RESULT such that ABS(EXACT - RESULT). LE.MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example 1}

In this example, we approximate the integral
\[
\int_{0}^{1} \int_{1}^{3} y \cos \left(x+y^{2}\right) d y d x
\]

The value of the error estimate is machine dependent.
```

    USE TWODQ_INT
    USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
    IMPLICIT - NONE
    INTEGER IRULE, NOUT
    REAL A, B, ERRABS, ERREST, ERRREL, F, G, H, RESULT
    EXTERNAL F, G, H
    CALL UMACH (2, NOUT)
    A = 0.0
    B = 1.0
    ! Set error tolerances
ERRABS = 0.0
ERRREL = 0.01
! Parameter for oscillatory function
IRULE = 6
CALL TWODQ (F, A, B, G, H, RESULT, ERRABS, ERRREL, IRULE, ERREST)
WRITE (NOUT,99999) RESULT, ERREST
99999 FORMAT (' Result =', F8.3, 13X, ' Error estimate = ', 1PE9.3)
END

```
```

!
REAL FUNCTION F (X, Y)
REAL X, Y
REAL COS
INTRINSIC
COS
F = Y* COS (X+Y*Y)
RETURN
END
!
REAL FUNCTION G (X)
REAL X
G = 1.0
RETURN
END
!
REAL FUNCTION H (X)
REAL X
H = 3.0
RETURN
END

```

\section*{Output}
```

Result = -0.514 Error estimate = 3.065E-06

```

\section*{Additional Examples}

\section*{Example 2}

We modify the above example by assuming that the limits for the inner integral depend on \(x\) and, in particular, are \(g(x)=-2 x\) and \(h(x)=5 x\). The integral now becomes
\[
\int_{0}^{1} \int_{-2 x}^{5 x} y \cos \left(x+y^{2}\right) d y d x
\]

The value of the error estimate is machine dependent.
```

USE TWODQ_INT
USE UMACH_INT
INTEGER IRULE, NOUT
REAL A, B, ERRABS, ERREST, ERRREL, F, G, H, RESULT
EXTERNAL F, G, H
!
CALL UMACH (2, NOUT)
A = 0.0
B = 1.0
! Set error tolerances
ERRABS = 0.001
ERRREL = 0.0
! Parameter for oscillatory function
IRULE = 6
CALL TWODQ (F, A, B, G, H, RESULT, ERRABS, ERRREL, IRULE, ERREST)
Print results

```
```

    WRITE (NOUT,99999) RESULT, ERREST
    99999 FORMAT (' Computed =', F8.3, 13X, ' Error estimate = ', 1PE9.3)
END
REAL FUNCTION F (X, Y)
REAL X, Y
!
REAL COS
INTRINSIC COS
!
F = Y* COS (X+Y*Y)
RETURN
END
REAL FUNCTION G (X)
REAL X
!
G = -2.0*X
RETURN
END
REAL FUNCTION H (X)
REAL X
!
H = 5.0*X
RETURN
END

```

\section*{Output}

Computed \(=-0.083 \quad\) Error estimate \(=2.095 \mathrm{E}-06\)

\section*{QAND}

Integrates a function on a hyper-rectangle.

\section*{Required Arguments}
\(\boldsymbol{F}\) - User-supplied FUNCTION to be integrated. The form is \(\mathrm{F}(\mathrm{N}, \mathrm{X})\), where N - The dimension of the hyper-rectangle. (Input)
X - The independent variable of dimension N . (Input)
F - The value of the integrand at X . (Output)
F must be declared EXTERNAL in the calling program.
\(N\) - The dimension of the hyper-rectangle. (Input)
N must be less than or equal to 20 .

A - Vector of length N. (Input)
Lower limits of integration.
B-Vector of length \(N\). (Input)
Upper limits of integration.

RESULT - Estimate of the integral from A to B of F . (Output)
The integral of F is approximated over the N -dimensional hyper-rectangle
A.LE.X.LE.B.

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: ERRABS \(=1 . \mathrm{e}-3\) for single precision and 1.d-8 for double precision.
ERRREL - Relative accuracy desired. (Input)
Default: ERRREL \(=1 . e-3\) for single precision and 1.d-8 for double precision.
MAXFCN - Approximate maximum number of function evaluations to be permitted. (Input)
MAXFCN cannot be greater than \(256^{\mathrm{N}}\) or \(\operatorname{IMACH}(5)\) if N is greater than 3 .
Default: MAXFCN \(=32 * *\).

ERREST - Estimate of the absolute value of the error. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL QAND (F, N, A, B, RESULT [, ...])
Specific: The specific interface names are S_QAND and D_QAND.

\section*{FORTRAN 77 Interface}

Single: CALL QAND ( \(\mathrm{F}, \mathrm{N}, \mathrm{A}, \mathrm{B}, \mathrm{ERRABS}, \mathrm{ERRREL}, \mathrm{MAXFCN}, \mathrm{RESULT}\), ERREST)

Double: The double precision name is DQAND.

\section*{Description}

The routine QAND approximates the \(n\)-dimensional iterated integral
\[
\int_{a_{1}}^{b_{1}} \ldots \int_{a_{n}}^{b_{n}} f\left(x_{1}, \ldots, x_{n}\right) d x_{n} \ldots d x_{1}
\]
with the approximation returned in RESULT. An estimate of the error is returned in ERREST. The approximation is achieved by iterated applications of product Gauss formulas. The integral is first estimated by a two-point tensor product formula in each direction. Then for \(i=1, \ldots, n\) the routine calculates a new estimate by doubling the number of points in the \(i\)-th direction, but halving the number immediately afterwards if the new estimate does not change appreciably. This process is repeated until either one complete sweep results in no increase in the number of sample points in any dimension, or the number of Gauss points in one direction exceeds 256 , or the number of function evaluations needed to complete a sweep would exceed MAXFCN.

\section*{Comments}
1. Informational errors

Type Code
31 MAXFCN was set greater than \(256^{N}\).
42 The maximum number of function evaluations has been reached, and convergence has not been attained.
2. If EXACT is the exact value, QAND attempts to find RESULT such that ABS(EXACT - RESULT). LE.MAX(ERRABS, ERRREL * ABS(EXACT)). To specify only a relative error, set ERRABS to zero. Similarly, to specify only an absolute error, set ERRREL to zero.

\section*{Example}

In this example, we approximate the integral of
\[
e^{-\left(x_{1}^{2}+x_{2}^{2}+x_{3}^{2}\right)}
\]
on an expanding cube. The values of the error estimates are machine dependent. The exact integral over
\[
\mathbf{R}^{3} \text { is } \pi^{3 / 2}
\]

USE QAND_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER I, J, MAXFCN, N, NOUT
REAL A(3), B(3), CNST, ERRABS, ERREST, ERRREL, F, RESULT
EXTERNAL F
! Get output unit number
CALL UMACH (2, NOUT)
\(!\)
\(\mathrm{N} \quad=3\)
MAXFCN \(=100000\)
! \(\quad\) ERRABS \(=0.0001\)
ERRREL \(=0.001\)
!
DO 20 I=1, 6 CNST \(=\) I/2.0
```

!
!
Set limits of integration
As CNST approaches infinity, the answer approaches PI**1.5
DO $10 \mathrm{~J}=1,3$

$$
A(J)=-C N S T
$$

$$
\mathrm{B}(\mathrm{~J})=\mathrm{CNST}
$$

10 CONTINUE CALL QAND ( $\mathrm{F}, \mathrm{N}, \mathrm{A}, \mathrm{B}, \mathrm{RESULT}, \mathrm{ERRABS}, \mathrm{ERRREL}, \mathrm{MAXFCN}, ~ E R R E S T) ~$ WRITE (NOUT,99999) CNST, RESULT, ERREST

```

20 CONTINUE
99999 FORMAT (1X, 'For CNST = ', F4.1, ', result = ', F7.3, ' with ', \&
```

        'error estimate ', 1PE10.3)
    END
    !
REAL FUNCTION F (N, X)
INTEGER N
REAL X(N)
REAL EXP
INTRINSIC EXP
F = EXP(-(X(1)*X(1)+X(2)*X(2)+X(3)*X(3)))
RETURN
END

```

\section*{Output}
For CNST \(=0.5\), result \(=0.785\) with error estimate \(3.934 \mathrm{E}-06\)
For CNST \(=1.0\), result \(=3.332\) with error estimate \(2.100 \mathrm{E}-03\)
For CNST \(=1.5\), result \(=5.021\) with error estimate \(1.192 \mathrm{E}-05\)
For CNST \(=2.0\), result \(=5.491\) with error estimate \(2.413 \mathrm{E}-04\)
For CNST \(=2.5\), result \(=5.561\) with error estimate \(4.232 \mathrm{E}-03\)
For CNST \(=3.0\), result \(=5.568\) with error estimate \(2.580 \mathrm{E}-04\)

\section*{QMC}

Integrates a function over a hyper rectangle using a quasi-Monte Carlo method.

\section*{Required Arguments}

FCN - User-supplied FUNCTION to be integrated. The form is FCN (X), where
x - The independent variable. (Input)
FCN - The value of the integrand at X. (Output)
FCN must be declared EXTERNAL in the calling program.
\(\boldsymbol{A}-\quad\) Vector containing lower limits of integration. (Input)
\(\boldsymbol{B}\) - Vector containing upper limits of integration. (Input)
RESULT - The value of
\[
\int_{a_{1}}^{b_{1}} \ldots \int_{a_{n}}^{b_{n}} f\left(x_{1}, \ldots, x_{n}\right) d x_{n} \ldots d x_{1}
\]
is returned, where n is the dimension of x . If no value can be computed, then NaN is returned. (Output)

\section*{Optional Arguments}

ERRABS - Absolute accuracy desired. (Input)
Default: \(1.0 \mathrm{e}-2\).

ERRREL - Relative accuracy desired. (Input)
Default: 1.0e-2.
ERREST - Estimate of the absolute value of the error. (Output)
MAXEVALS - Number of evaluations allowed. (Input) Default: No limit.

BASE - The base of the Faure sequence. (Input)
Default: The smallest prime number greater than or equal to the number of dimensions (length of \(a\) and \(b\) ).

SKIP - The number of points to be skipped at the beginning of the Faure sequence. (Input) Default: \(\left\lfloor\right.\) base \(\left.^{m / 2-1}\right\rfloor\), where \(m=\lfloor\log B / \log\) base \(\rfloor\) and \(B\) is the largest representable integer.

\section*{FORTRAN 90 Interface}

Generic: CALL QMC (FCN, A, B, RESULT [,...])

Specific: The specific interface names are S_QMC and D_QMC.

\section*{Description}

Integration of functions over hyper rectangle by direct methods, such as QAND, is practical only for fairly low dimensional hypercubes. This is because the amount of work required increases exponentially as the dimension increases.
An alternative to direct methods is \(Q M C\), in which the integral is evaluated as the value of the function averaged over a sequence of randomly chosen points. Under mild assumptions on the function, this method will converge like
\[
1 / \sqrt{k}
\]
where \(k\) is the number of points at which the function is evaluated.
It is possible to improve on the performance of QMC by carefully choosing the points at which the function is to be evaluated. Randomly distributed points tend to be non-uniformly distributed. The alternative to a sequence of random points is a low-discrepancy sequence. A low-discrepancy sequence is one that is highly uniform.
This function is based on the low-discrepancy Faure sequence as computed by FAURE_NEXT, see Stat Library, Chapter 18, Random Number Generation.

\section*{Example}

This example evaluates the n -dimensional integral
\[
\int_{0}^{1} \ldots \int_{0}^{1} \sum_{i=1}^{w} \prod_{j=1}^{i}(-1)^{i} x_{j} d x_{1} \ldots d x_{n}=-\frac{1}{3}\left[1-\left(-\frac{1}{2}\right)^{n}\right]
\]
with \(n=10\).
```

use qmc int
implicit none
integer, parameter :: ndim=10
real(kind(1d0)) :: a (ndim)
real(kind(1d0)) :: b(ndim)
real(kind(1d0)) :: result
integer :: I
external fcn
a = 0.d0
b = 1.d0
call qmc(fcn, a, b, result)
write (*,*) 'result = ', result
end
real(kind(1d0)) function fcn(x)
implicit none
real(kind(1d0)), dimension(:) :: x
integer :: i, j
real(kind(1d0)) :: prod, sum, sign
sign = -1.d0
sum = 0.d0
do i=1, size(x)
prod = 1.d0
prod = product(x(1:i))
sum = sum + (sign * prod)
sign = -sign
end do
fcn = sum
end function fcn

```

\section*{Output}
```

result = -0.3334789

```

\section*{GQRUL}

Computes a Gauss, Gauss-Radau, or Gauss-Lobatto quadrature rule with various classical weight functions.

\section*{Required Arguments}
\(\boldsymbol{N}\) - Number of quadrature points. (Input)
\(\boldsymbol{Q} \boldsymbol{X}\) - Array of length N containing quadrature points. (Output)
\(\boldsymbol{Q W}\) - Array of length N containing quadrature weights. (Output)

\section*{Optional Arguments}

IWEIGH - Index of the weight function. (Input)
Default: \(\operatorname{IWEIGH}=1\).
\begin{tabular}{llll} 
IWEIGH & WT \((\mathbf{X})\) & Interval & Name \\
1 & 1 & \((-1,+1)\) & Legendre \\
2 & \(1 / \sqrt{1-X^{2}}\) & \((-1,+1)\) & Chebyshev 1st kind \\
3 & \(\sqrt{1-X^{2}}\) & \((-1,+1)\) & Chebyshev 2nd kind \\
4 & \(e^{-X^{2}}\) & \((-\infty,+\infty)\) & Hermite \\
5 & \((1-X)^{\alpha}(1+X)^{\beta}\) & \((-1,+1)\) & Jacobi \\
6 & \(e^{-X} X^{\alpha}\) & \((0,+\infty)\) & Generalized Laguerre \\
7 & \(1 / \cosh (X)\) & \((-\infty,+\infty)\) & COSH
\end{tabular}

ALPHA - Parameter used in the weight function with some values of IWEIGH, otherwise it is ignored. (Input)
Default: ALPHA \(=2.0\).

BETAW - Parameter used in the weight function with some values of IWEIGH, otherwise it is ignored. (Input)
Default: \(\operatorname{BETAW}=2.0\).

NFIX - Number of fixed quadrature points. (Input)
NFIX \(=0,1\) or 2 . For the usual Gauss quadrature rules, NFIX \(=0\).
Default: NFIX \(=0\).

QXFIX - Array of length NFIX (ignored if NFIX \(=0\) ) containing the preset quadrature point(s). (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL GQRUL (N, QX, QW [,...])

Specific: The specific interface names are S_GQRUL and D_GQRUL.

\section*{FORTRAN 77 Interface}

Single: CALL GQRUL (N, IWEIGH, ALPHA, BETAW, NFIX, QXFIX, QX, QW)
Double: The double precision name is DGQRUL.

\section*{Description}

The routine GQRUL produces the points and weights for the Gauss, Gauss-Radau, or Gauss-Lobatto quadrature formulas for some of the most popular weights. In fact, it is slightly more general than this suggests because the extra one or two points that may be specified do not have to lie at the endpoints of the interval. This routine is a modification of the subroutine GAUSSQUADRULE (Golub and Welsch 1969).

In the simple case when NFIX \(=0\), the routine returns points in \(x=Q \mathrm{X}\) and weights in \(w=Q W\) so that
\[
\int_{a}^{b} f(x) w(x) d x=\sum_{i=1}^{N} f\left(x_{i}\right) w_{i}
\]
for all functions \(f\) that are polynomials of degree less than \(2 N\).
If NFIX \(=1\), then one of the above \(x_{i}\) equals the first component of QXFIX. Similarly, if NFIX \(=2\), then two of the components of \(x\) will equal the first two components of QXFIX. In general, the accuracy of the above quadrature formula degrades when NFIX increases. The quadrature rule will integrate all functions \(f\) that are polynomials of degree less than \(2 N-\) NFIX.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of G2RUL / DG2RUL. The reference is

CALL G2RUL (N, IWEIGH, ALPHA, BETAW, NFIX, QXFIX, QX, QW, WK)
The additional argument is

WK — Work array of length N .
2. If IWEIGH specifies the weight \(\mathrm{WT}(\mathrm{X})\) and the interval \((a, b)\), then approximately
\[
\int_{a}^{b} F(X) * W T(X) d X=\sum_{I=1}^{N} F(Q X(I)) * Q W(I)
\]
3. Gaussian quadrature is always the method of choice when the function \(F(X)\) behaves like a polynomial. Gaussian quadrature is also useful on infinite intervals (with appropriate weight functions), because other techniques often fail.
4. The weight function \(1 / \cosh (X)\) behaves like a polynomial near zero and like \(e^{|X|}\) far from zero.

\section*{Example 1}

In this example, we obtain the classical Gauss-Legendre quadrature formula, which is accurate for polynomials of degree less than \(2 N\), and apply this when \(N=6\) to the function \(x^{8}\) on the interval \([-1,1]\). This quadrature rule is accurate for polynomials of degree less than 12.

USE GQRUL INT
```

    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=6)
    INTEGER I, NOUT
    REAL ANSWER, QW(N), QX(N), SUM
    get output unit number
    Get points and weights from GQRUL
    Write results from GQRUL
    WRITE (NOUT,99998) (I,QX(I),I,QW(I),I=1,N)
    99998 FORMAT (6(6X,'QX(',I1,') = ',F8.4,7X,'QW(',I1,') = ',F8.5,/))
Evaluate the integral from these
points and weights
SUM = 0.0
DO 10 I=1, N
SUM = SUM + QX(I)**8*QW(I)
10 CONTINUE
ANSWER = SUM
WRITE (NOUT,99999) ANSWER
99999 FORMAT (/, ' The quadrature result making use of these ', \&
'points and weights is ', 1PE10.4, '.')
END

```

\section*{Output}
```

QX(1)=-0.9325 QW(1)=0.17132
QX(2) = -0.6612 QW (2) = 0.36076
QX(3) = -0.2386 QW(3) = 0.46791
QX(4)=0.2386 QW(4)=0.46791
QX(5)=0.6612 QW(5)=0.36076
QX(6)=0.9325 QW(6)=0.17132
The quadrature result making use of these points and weights is 2.2222E-01.

```

\section*{Additional Examples}

\section*{Example 2}

We modify Example 1 by requiring that both endpoints be included in the quadrature formulas and again apply the new formulas to the function \(x^{8}\) on the interval \([-1,1]\). This quadrature rule is accurate for polynomials of degree less than 10 .
```

USE GQRUL INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=6)
INTEGER I, IWEIGH, NFIX, NOUT
REAL ALPHA, ANSWER, BETAW, QW (N), QX(N), QXFIX(2), SUM

```
```

! Get output unit number
CALL UMACH (2, NOUT)
!
IWEIGH = 1
ALPHA = 0.0
BETAW = 0.0
NFIX = 2
QXFIX(1) = -1.0
QXFIX(2) = 1.0
! Get points and weights from GQRUL
CALL GQRUL (N, QX, QW, ALPHA=ALPHA, BETAW=BETAW, NFIX=NFIX, \&
QXFIX=QXFIX)
Write results from GQRUL
WRITE (NOUT,99998) (I,QX(I),I,QW(I),I=1,N)
99998 FORMAT (6(6X,'QX(',I1,') = ',F8.4,7X,'QW(',I1,') = ',F8.5,/))
Evaluate the integral from these
points and weights
SUM = 0.0
DO 10 I=1, N
SUM = SUM + QX(I)** **QW (I)
1 0 ~ C O N T I N U E
ANSWER = SUM
WRITE (NOUT,99999) ANSWER
99999 FORMAT (/, ' The quadrature result making use of these ', \&
'points and weights is ', 1PE10.4, '.')
END

```

\section*{Output}
```

QX(1)=-1.0000 QW(1) = 0.06667
QX(2) = -0.7651 QW (2) = 0.37847
QX(3)=-0.2852 QW (3) = 0.55486
QX(4)=0.2852 QW(4)=0.55486
QX(5) = 0.7651 QW (5) = 0.37847
QX(6)=1.0000 QW(6)=0.06667
The quadrature result making use of these points and weights is 2.2222E-01.

```

\section*{GQRCF}

Computes a Gauss, Gauss-Radau or Gauss-Lobatto quadra ture rule given the recurrence coefficients for the monic polynomials orthogonal with respect to the weight function.

\section*{Required Arguments}
\(\boldsymbol{N}\) - Number of quadrature points. (Input)
\(\boldsymbol{B}\) - Array of length N containing the recurrence coefficients. (Input) See Comments for definitions.
\(\boldsymbol{C}\) - Array of length N containing the recurrence coefficients. (Input) See Comments for definitions.
\(\boldsymbol{Q} \boldsymbol{X}\) - Array of length N containing quadrature points. (Output)
\(\boldsymbol{Q W}\) - Array of length N containing quadrature weights. (Output)

\section*{Optional Arguments}

NFIX - Number of fixed quadrature points. (Input) NFIX \(=0,1\) or 2 . For the usual Gauss quadrature rules NFIX \(=0\). Default: NFIX \(=0\).

QXFIX - Array of length NFIX (ignored if NFIX \(=0\) ) containing the preset quadrature point(s). (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL GQRCF (N, B, C, QX, QW [,...])

Specific: The specific interface names are \(S_{-} G Q R C F\) and \(D_{-} G Q R C F\).

\section*{FORTRAN 77 Interface}

Single: CALL GQRCF (N, B, C, NFIX, QXFIX, QX, QW)
Double: The double precision name is DGQRCF.

\section*{Description}

The routine GQRCF produces the points and weights for the Gauss, Gauss-Radau, or Gauss-Lobatto quadrature formulas given the three-term recurrence relation for the orthogonal polynomials. In particular, it is assumed that the orthogonal polynomials are monic, and hence, the three-term recursion may be written as
\[
p_{i}(x)=\left(x-b_{i}\right) p_{i-1}(x)-c_{i} p_{i-2}(x) \text { for } i=1, \ldots, N
\]
where \(p_{0}=1\) and \(p_{-1}=0\). It is obvious from this representation that the degree of \(p_{i}\) is \(i\) and that \(p_{i}\) is monic. In order for the recurrence to give rise to a sequence of orthogonal polynomials (with respect to a nonnegative measure), it is necessary and sufficient that \(c_{i}>0\). This routine is a modification of the subroutine GAUSSQUADRULE (Golub and Welsch 1969). In the simple case when NFIX \(=0\), the routine returns points in \(x=Q X\) and weights in \(w=Q W\) so that
\[
\int_{a}^{b} f(x) w(x) d x=\sum_{i=1}^{N} f\left(x_{i}\right) w_{i}
\]
for all functions \(f\) that are polynomials of degree less than \(2 N\). Here, \(w\) is any weight function for which the above recurrence produces the orthogonal polynomials \(p_{i}\) on the interval \([a, b]\) and \(w\) is normalized by
\[
\int_{a}^{b} w(x) d x=c_{1}
\]

If \(\mathrm{NFIX}=1\), then one of the above \(x_{i}\) equals the first component of QXFIX. Similarly, if NFIX \(=2\), then two of the components of \(x\) will equal the first two components of QXFIX. In general, the accuracy of the above quadrature formula degrades when NFIX increases. The quadrature rule will integrate all functions \(f\) that are polynomials of degree less than \(2 N-\) NFIX.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of G2RCF/DG2RCF. The reference is:

CALL G2RCF (N, B, C, NFIX, QXFIX, QX, QW, WK)
The additional argument is:
\(\boldsymbol{W K}\) - Work array of length N.
2. Informational error

Type Code
41 No convergence in 100 iterations.
3. The recurrence coefficients \(B(I)\) and \(C(I)\) define the monic polynomials via the relation \(P(I)=(X-B(I+1)) * P(I-1)-C(I+1) * P(I-2) . C(1)\) contains the zero-th moment
\[
\int W T(X) d X
\]
of the weight function. Each element of C must be greater than zero.
4. If \(W T(X)\) is the weight specified by the coefficients and the interval is \((a, b)\), then approximately
\[
\int_{a}^{b} F(X) * W T(X) d X=\sum_{I=1}^{N} F(Q X(I)) * Q W(I)
\]
5. Gaussian quadrature is always the method of choice when the function \(F(X)\) behaves like a polynomial. Gaussian quadrature is also useful on infinite intervals (with appropriate weight functions) because other techniques often fail.

\section*{Example}

We compute the Gauss quadrature rule (with \(N=6\) ) for the Chebyshev weight, \(\left(1+x^{2}\right)^{(-1 / 2)}\), from the recurrence coefficients. These coefficients are obtained by a call to the IMSL routine RECCF.
```

USE GQRCF_INT
USE UMACH INT
USE RECCF-
IMPLICIT NONE
INTEGER N

```
```

    PARAMETER (N=6)
    INTEGER I, NFIX, NOUT
    REAL B(N), C(N), QW(N), QX(N), QXFIX(2)
    ! Get output unit number
CALL UMACH (2, NOUT)
Recursion coefficients will come from
routine RECCF.
The call to RECCF finds recurrence
coefficients for Chebyshev
polynomials of the lst kind.
CALL RECCF (N, B, C)
The call to GQRCF will compute the
quadrature rule from the recurrence
coefficients determined above.
CALL GQRCF (N, B, C, QX, QW)
WRITE (NOUT,99999) (I,QX(I),I,QW(I),I=1,N)
99999 FORMAT (6(6X,'QX(',I1,') = ',F8.4,7X,'QW(',I1,') = ',F8.5,/))
!
END

```

\section*{Output}
\begin{tabular}{ll}
\(\mathrm{QX}(1)=-0.9325\) & \(\mathrm{QW}(1)=0.17132\) \\
\(\mathrm{QX}(2)=-0.6612\) & \(\mathrm{QW}(2)=0.36076\) \\
\(\mathrm{QX}(3)=-0.2386\) & \(\mathrm{QW}(3)=0.46791\) \\
\(\mathrm{QX}(4)=0.2386\) & \(\mathrm{QW}(4)=0.46791\) \\
\(\mathrm{QX}(5)=0.6612\) & \(\mathrm{QW}(5)=0.36076\) \\
\(\mathrm{QX}(6)=0.9325\) & \(\mathrm{QW}(6)=0.17132\)
\end{tabular}

\section*{RECCF}

Computes recurrence coefficients for various monic polynomials.

\section*{Required Arguments}
\(N\) - Number of recurrence coefficients. (Input)
\(\boldsymbol{B}\) - Array of length N containing recurrence coefficients. (Output)
\(\boldsymbol{C}\) - Array of length N containing recurrence coefficients. (Output)

\section*{Optional Arguments}

IWEIGH - Index of the weight function. (Input)
Default: \(\operatorname{IWEIGH}=1\).
\begin{tabular}{llll} 
IWEIGH & WT \((\mathbf{X})\) & Interval & Name \\
1 & 1 & \((-1,+1)\) & Legendre \\
2 & \(1 / \sqrt{1-X^{2}}\) & \((-1,+1)\) & Chebyshev 1st kind \\
3 & \(\sqrt{1-X^{2}}\) & \((-1,+1)\) & Chebyshev 2nd kind \\
4 & \(e^{-X^{2}}\) & \((-\infty,+\infty)\) & Hermite \\
5 & \((1-X)^{\alpha}(1+X)^{\beta}\) & \((-1,+1)\) & Jacobi \\
6 & \(e^{-X} X^{\alpha}\) & \((0,+\infty)\) & Generalized Laguerre \\
7 & \(1 / \cosh (X)\) & \((-\infty,+\infty)\) & COSH
\end{tabular}

ALPHA - Parameter used in the weight function with some values of IWEIGH, otherwise it is ignored. (Input)
Default: ALPHA=1.0.
BETAW - Parameter used in the weight function with some values of IWEIGH, otherwise it is ignored. (Input)
Default: BETAW=1.0.

\section*{FORTRAN 90 Interface}

Generic: CALL RECCF (N, B, C [, ...])
Specific: The specific interface names are \(S \_\)RECCF and D_RECCF.

\section*{FORTRAN 77 Interface}

Single: CALL RECCF (N, IWEIGH, ALPHA, BETAW, B, C)
Double: The double precision name is DRECCF.

\section*{Description}

The routine RECCF produces the recurrence coefficients for the orthogonal polynomials for some of the most important weights. It is assumed that the orthogonal polynomials are monic; hence, the three-term recursion may be written as
\[
p_{i}(x)=\left(x-b_{i}\right) p_{i-1}(x)-c_{i} p_{i-2}(x) \text { for } i=1, \ldots, N
\]
where \(p_{0}=1\) and \(p_{-1}=0\). It is obvious from this representation that the degree of \(p_{i}\) is \(i\) and that \(p_{i}\) is monic. In order for the recurrence to give rise to a sequence of orthogonal polynomials (with respect to a nonnegative measure), it is necessary and sufficient that \(c_{i}>0\).

\section*{Comments}

The recurrence coefficients \(B(I)\) and \(C(I)\) define the monic polynomials via the relation \(P(I)=(X-B(I+1)) * P(I-1)-C(I+1) * P(I-2)\). The zero-th moment
\[
\left(\int W T(X) d X\right)
\]
of the weight function is returned in \(C(1)\).

\section*{Example}

Here, we obtain the well-known recurrence relations for the first six monic Legendre polynomials, Chebyshev polynomials of the first kind, and Laguerre polynomials.
```

USE RECCF_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=6)
INTEGER I, IWEIGH, NOUT
REAL ALPHA, B(N), C (N), BETAW
! Get output unit number
CALL UMACH (2, NOUT)
!
CALL RECCF (N, B, C)
WRITE (NOUT,99996)
WRITE (NOUT,99999) (I, B (I), I, C(I), I=1,N)
!
IWEIGH = 2
CALL RECCF (N, B, C, IWEIGH=IWEIGH)
WRITE (NOUT,99997)
WRITE (NOUT,99999) (I, B (I), I, C(I), I=1,N)
!
IWEIGH = 6
ALPHA = 0.0
BETAW = 0.0
CALL RECCF (N, B, C, IWEIGH=IWEIGH, ALPHA=ALPHA)
WRITE (NOUT,99998)
WRITE (NOUT,99999) (I, B (I), I, C (I), I=1,N)
!
99996 FORMAT (1X, 'Legendre')
9 9 9 9 7 ~ F O R M A T ~ ( / , ~ 1 X , ~ ' C h e b y s h e v , ~ f i r s t ~ k i n d ' )
9 9 9 9 8 ~ F O R M A T ~ ( / , ~ 1 X , ~ ' L a g u e r r e ' ) ~
99999 FORMAT (6(6X,'B(',I1,') = ',F8.4,7X,'C(',I1,') = ',F8.5,/))
END

```

\section*{Output}
```

Legendre
B(1)=0.0000 C(1)= 2.00000
B(2)=0.0000 C(2) = 0.33333
B(3)=0.0000 C(3) = 0.26667
B(4)=0.0000 C(4) = 0.25714
B(5)=0.0000 C(5) = 0.25397
B(6)=0.0000 C(6) = 0.25253
Chebyshev, first kind
B(1)=0.0000 C(1)=3.14159

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline B (2) & \(=\) & 0.0000 & C (2) & = & 0.50000 \\
\hline B (3) & \(=\) & 0.0000 & C (3) & = & 0.25000 \\
\hline B (4) & \(=\) & 0.0000 & C (4) & \(=\) & 0.25000 \\
\hline B (5) & \(=\) & 0.0000 & C (5) & = & 0.25000 \\
\hline B (6) & \(=\) & 0.0000 & C (6) & \(=\) & 0.25000 \\
\hline \multicolumn{6}{|l|}{Laguerre} \\
\hline B (1) & \(=\) & 1.0000 & C (1) & \(=\) & 1.00000 \\
\hline B (2) & \(=\) & 3.0000 & C (2) & \(=\) & 1.00000 \\
\hline B (3) & \(=\) & 5.0000 & C (3) & = & 4.00000 \\
\hline B (4) & \(=\) & 7.0000 & C (4) & \(=\) & 9.00000 \\
\hline B (5) & \(=\) & 9.0000 & C (5) & = & 16.00000 \\
\hline B (6) & \(=\) & 11.0000 & C (6) & \(=\) & 25.00000 \\
\hline
\end{tabular}

\section*{RECQR}

Computes recurrence coefficients for monic polynomials given a quadrature rule.

\section*{Required Arguments}
\(\boldsymbol{Q X}\) - Array of length n containing the quadrature points. (Input)
\(\boldsymbol{Q W}\) - Array of length N containing the quadrature weights. (Input)
B - Array of length NTERM containing recurrence coefficients. (Output)
C - Array of length NTERM containing recurrence coefficients. (Output)

\section*{Optional Arguments}
\(\boldsymbol{N}\) - Number of quadrature points. (Input)
Default: \(\mathrm{N}=\operatorname{size}(\mathrm{QX}, 1)\).
NTERM - Number of recurrence coefficients. (Input)
NTERM must be less than or equal to N .
Default: NTERM \(=\operatorname{size}(B, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL RECQR (QX, QW, B, C [, ...])
Specific: The specific interface names are \(S \_R E C Q R\) and \(D \_R E C Q R\).

\section*{FORTRAN 77 Interface}

Single: CALL RECQR ( \(N, ~ Q X, Q W, ~ N T E R M, B, C\) )
Double: The double precision name is \(D R E C Q R\).

\section*{Description}

The routine \(R E C Q R\) produces the recurrence coefficients for the orthogonal polynomials given the points and weights for the Gauss quadrature formula. It is assumed that the orthogonal polynomials are monic; hence the three-term recursion may be written
\[
p_{i}(x)=\left(x-b_{i}\right) p_{i-1}(x)-c_{i} p_{i-2}(x) \text { for } i=1, \ldots, N
\]
where \(p_{0}=1\) and \(p_{-1}=0\). It is obvious from this representation that the degree of \(p_{i}\) is \(i\) and that \(p_{i}\) is monic. In order for the recurrence to give rise to a sequence of orthogonal polynomials (with respect to a nonnegative measure), it is necessary and sufficient that \(c_{i}>0\).

This routine is an inverse routine to \(G Q R C F\). Given the recurrence coefficients, the routine \(G Q R C F\) produces the corresponding Gauss quadrature formula, whereas the routine RECQR produces the recurrence coefficients given the quadrature formula.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(R 2 C Q R / D R 2 C Q R\). The reference is:
```

CALL R2CQR (N, QX, QW, NTERM, B, C,WK)

```

The additional argument is:
WKWK - Work array of length 2 * N .
2. The recurrence coefficients \(B(I)\) and \(C(I)\) define the monic polynomials via the relation \(P(I)=(X-B(I+1)) * P(I-1)-C(I+1) * P(I-2)\). The zero-th moment
\[
\left(\int W T(X) d X\right)
\]
of the weight function is returned in \(C(1)\).

\section*{Example}

To illustrate the use of \(\operatorname{RECQR}\), we will input a simple choice of recurrence coefficients, call GQRCF for the quadrature formula, put this information into \(R E C Q R\), and recover the recurrence coefficients.
```

USE RECQR_INT
USE UMACH INT
USE GQRCF INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=5)
INTEGER I, J, NFIX, NOUT, NTERM
REAL B(N), C(N), FLOAT, QW(N), QX(N), QXFIX(2)
INTRINSIC FLOAT
Get output unit number

```
```

    CALL UMACH (2, NOUT)
    NFIX = 0
    ! Set arrays B and C of recurrence
DO 10 J=1, N
B(J) = FLOAT (J)
C(J) = FLOAT (J)/2.0
10 CONTINUE
WRITE (NOUT,99995)
99995 FORMAT (1X, 'Original recurrence coefficients')
WRITE (NOUT,99996) (I,B(I),I,C(I),I=1,N)
99996 FORMAT (5 (6X,'B(',I1,') = ',F8.4,7X,'C(',I1,') = ',F8.5,/))
!
! The call to GQRCF will compute the
quadrature rule from the recurrence
coefficients given above.
CALL GQRCF (N, B, C, QX, QW)
WRITE (NOUT,99997)
99997 FORMAT (/, 1X, 'Quadrature rule from the recurrence coefficients' \&
)
WRITE (NOUT,99998) (I,QX(I),I,QW(I),I=1,N)
99998 FORMAT (5 (6X,'QX(',I1,') = ',F8.4,7X,'QW(',I1,') = ',F8.5,/))
!
! Call RECQR to recover the original
recurrence coefficients
NTERM = N
CALL RECQR (QX, QW, B, C)
WRITE (NOUT,99999)
9 9 9 9 9 ~ F O R M A T ~ ( / , ~ 1 X , ~ ' R e c u r r e n c e ~ c o e f f i c i e n t s ~ d e t e r m i n e d ~ b y ~ R E C Q R ' )
WRITE (NOUT,99996) (I,B(I),I,C(I),I=1,N)
!
END

```

\section*{Output}
```

Original recurrence coefficients
B(1)=1.0000 C(1) = 0.50000
B(2)=2.0000 C(2) = 1.00000
B(3)=3.0000 C(3) = 1.50000
B(4)=4.0000 C(4)=2.00000
B(5)=5.0000 C(5) = 2.50000

```
Quadrature rule from the recurrence coefficients
\begin{tabular}{ll}
\(\mathrm{QX}(1)=0.1525\) & \(\mathrm{QW}(1)=0.25328\) \\
\(\mathrm{QX}(2)=1.4237\) & \(\mathrm{QW}(2)=0.17172\) \\
\(\mathrm{QX}(3)=2.7211\) & \(\mathrm{QW}(3)=0.06698\) \\
\(\mathrm{QX}(4)=4.2856\) & \(\mathrm{QW}(4)=0.00790\) \\
\(\mathrm{QX}(5)=6.4171\) & \(\mathrm{QW}(5)=0.00012\)
\end{tabular}
Recurrence coefficients determined by RECQR
\begin{tabular}{ll}
\(\mathrm{B}(1)=1.0000\) & \(\mathrm{C}(1)=0.50000\) \\
\(\mathrm{~B}(2)=2.0000\) & \(\mathrm{C}(2)=1.00000\) \\
\(\mathrm{~B}(3)=3.0000\) & \(\mathrm{C}(3)=1.50000\)
\end{tabular}
```

B(4)=4.0000 C(4) = 2.00000
B(5)=5.0000 C(5) = 2.50000

```

\section*{FQRUL}

Computes a Fejér quadrature rule with various classical weight functions.

\section*{Required Arguments}
\(N\) - Number of quadrature points. (Input)
\(\boldsymbol{A}-\) Lower limit of integration. (Input)
\(\boldsymbol{B}\) - Upper limit of integration. (Input) B must be greater than A.
\(\boldsymbol{Q} \boldsymbol{X}\) - Array of length n containing quadrature points. (Output)
\(\mathbf{Q W}\) - Array of length N containing quadrature weights. (Output)

\section*{Optional Arguments}

IWEIGH - Index of the weight function. (Input)
Default: \(\operatorname{IWEIGH}=1\).
```

IWEIGH WT(X)
1
2 1/(X - ALPHA)
3 (B-X)}\mp@subsup{)}{}{\alpha}(X-A\mp@subsup{)}{}{\beta
4 (B-X) }(\textrm{X}-\textrm{A}\mp@subsup{)}{}{\beta}\operatorname{log}(\textrm{X}-\textrm{A}
5 (B-X) }(\textrm{X}-\textrm{A}\mp@subsup{)}{}{\beta}\operatorname{log}(\textrm{B}-\textrm{X}

```

ALPHA - Parameter used in the weight function (except if IWEIGH \(=1\), it is ignored).
(Input)
If \(\operatorname{IWEIGH}=2\), then it must satisfy A. LT.ALPHA. LT.B. If IWEIGH \(=3,4\), or 5, then
ALPHA must be greater than -1 .
Default: ALPHA= 0.0.

BETAW - Parameter used in the weight function (ignored if \(I W E I G H=1\) or 2). (Input)
BETAW must be greater than -1.0 .
Default: BETAW= 0.0.

\section*{FORTRAN 90 Interface}

Generic: CALL FQRUL (N, A, B, QX, QW [, ...])
Specific: The specific interface names are S_FQRUL and D_FQRUL.

\section*{FORTRAN 77 Interface}

Single: CALL FQRUL (N, A, B, IWEIGH, ALPHA, BETAW, QX, QW)
Double: The double precision name is DFQRUL.

\section*{Description}

The routine FQRUL produces the weights and points for the Fejér quadrature rule. Since this computation is based on a quarter-wave cosine transform, the computations are most efficient when \(N\), the number of points, is a product of small primes. These quadrature formulas may be an intermediate step in a more complicated situation, see for instance Gautschi and Milovanofic (1985).

The Fejér quadrature rules are based on polynomial interpolation. First, choose classical abscissas (in our case, the Gauss points for the Chebyshev weight function \(\left(1-x^{2}\right)^{-1 / 2}\) ), then derive the quadrature rule for a different weight. In order to keep the presentation simple, we will describe the case where the interval of integration is \([-1,1]\) even though \(F Q R U L\) allows rescaling to an arbitrary interval \([a, b]\).
We are looking for quadrature rules of the form
\[
Q(f):=\sum_{j=1}^{N} w_{j} f\left(x_{j}\right)
\]
where the
\[
\left\{x_{j}\right\}_{j=1}^{N}
\]
are the zeros of the \(N\)-th Chebyshev polynomial (of the first kind) \(T_{N}(x)=\cos (N \arccos x)\). The weights in the quadrature rule \(Q\) are chosen so that, for all polynomials \(p\) of degree less than \(N\),
\[
Q(p)=\sum_{j=1}^{N} w_{j} p\left(x_{j}\right)=\int_{-1}^{1} p(x) w(x) d x
\]
for some weight function \(w\). In \(F Q R U L\), the user has the option of choosing \(w\) from five families of functions with various algebraic and logarithmic endpoint singularities.
These Fejér rules are important because they can be computed using specialized FFT quarter-wave transform routines. This means that rules with a large number of abscissas may be computed efficiently. If we insert \(\mathrm{T}_{l}\) for \(p\) in the above formula, we obtain
\[
Q\left(T_{l}\right)=\sum_{j=1}^{N} w_{j} T_{l}\left(x_{j}\right)=\int_{-1}^{1} T_{l}(x) w(x) d x
\]
for \(1=0, \ldots, \mathrm{~N}-1\). This is a system of linear equations for the unknown weights \(w_{j}\) that can be simplified by noting that
\[
x_{j}=\cos \frac{(2 j-1) \pi}{2 N} \quad j=1, \ldots, N
\]
and hence,
\[
\begin{aligned}
\int_{-1}^{1} T_{l}(x) w(x) d x= & \sum_{j=1}^{N} w_{j} T_{l}\left(x_{j}\right) \\
& =\sum_{j=1}^{N} w_{j} \cos \frac{l(2 j-1) \pi}{2 N}
\end{aligned}
\]

The last expression is the cosine quarter-wave forward transform for the sequence
\[
\left\{w_{j}\right\}_{j=1}^{N}
\]
that is implemented in Chapter 6, Transforms under the name QCOSF. More importantly, QCOSF has an inverse QCOSB. It follows that if the integrals on the left in the last expression can be computed, then the Fejer rule can be derived efficiently for highly composite integers \(N\) utilizing QCOSB. For more information on this topic, consult Davis and Rabinowitz (1984, pages 84-86) and Gautschi (1968, page 259).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of F2RUL/DF2RUL. The reference is:

CALL F2RUL (N, A, B, IWEIGH, ALPHA, BETAW, QX, QW, WK)
The additional argument is:
\(W K\) - Work array of length \(3 * N+15\).
2. If IWEIGH specifies the weight \(\mathrm{WT}(\mathrm{X})\) and the interval ( \(\mathrm{A}, \mathrm{B}\) ), then approximately
\[
\int_{A}^{B} F(X) * W T(X) d X=\sum_{I=1}^{N} F(Q X(I)) * Q W(I)
\]
3. The routine \(F Q R\) RL uses an \(f f t\), so it is most efficient when \(N\) is the product of small primes.

\section*{Example}

Here, we obtain the Fejér quadrature rules using 10,100 , and 200 points. With these rules, we get successively better approximations to the integral
\[
\int_{0}^{1} x \sin \left(41 \pi x^{2}\right) d x=\frac{1}{41 \pi}
\]
```

    USE FQRUL_INT
    USE UMACH INT
    USE CONST }\mp@subsup{}{}{-}\mathrm{ INT
    IMPLICIT NONE
    INTEGER NMAX
    PARAMETER (NMAX=200)
    INTEGER I, K, N, NOUT
    REAL A, ANSWER, B, F, QW (NMAX), &
        QX(NMAX), SIN, SUM, X, PI, ERROR
    INTRINSIC SIN, ABS
    !
F(X) = X*SIN(41.0*PI*X**2)
CALL UMACH (2, NOUT)
!
PI = CONST('PI')
DO 20 K=1, 3
IF (K .EQ. 1) N = 10
IF (K .EQ. 2) N = 100
IF (K .EQ. 3) N = 200
A = 0.0
B = 1.0
! Get points and weights from FQRUL
CALL FQRUL (N, A, B, QX, QW)
Evaluate the integral from these
points and weights
SUM = 0.0
DO 10 I=1, N
SUM = SUM + F(QX(I))*QW(I)
1 0
CONTINUE
ANSWER = SUM
ERROR = ABS (ANSWER - 1.0/(41.0*PI))
WRITE (NOUT,99999) N, ANSWER, ERROR
20 CONTINUE
!
99999 FORMAT (/, 1X, 'When N = ', I3, ', the quadrature result making ' \&
, 'use of these points ', /, ' and weights is ', 1PE11.4, \&
', with error ', 1PE9.2, '.')
END

```

\section*{Output}

When \(N=10\), the quadrature result making use of these points and weights is \(-1.6523 \mathrm{E}-01\), with error \(1.73 \mathrm{E}-01\).

When \(N=100\), the quadrature result making use of these points and weights is 7.7637E-03, with error 2.79E-08.

When \(N=200\), the quadrature result making use of these points and weights is \(7.7636 \mathrm{E}-03\), with error \(1.40 \mathrm{E}-08\).

\section*{DERIV}

This function computes the first, second or third derivative of a user-supplied function.

\section*{Function Return Value}

DERIV - Estimate of the first (KORDER \(=1)\), second \((K O R D E R=2)\) or third \((K O R D E R=3)\) derivative of FCN at X . (Output)

\section*{Required Arguments}

FCN - User-supplied FUNCTION whose derivative at X will be computed. The form is \(\operatorname{FCN}(\mathrm{X})\), where

X - Independent variable. (Input)
FCN - The function value. (Output)
ECN must be declared EXTERNAL in the calling program.
\(\boldsymbol{X}-\) Point at which the derivative is to be evaluated. (Input)

\section*{Optional Arguments}

KORDER - Order of the derivative desired (1, 2 or 3). (Input)
Default: \(K O R D E R=1\).

BGSTEP - Beginning value used to compute the size of the interval used in computing the derivative. (Input)
The interval used is the closed interval ( \(\mathrm{X}-4\) * BGSTEP, \(\mathrm{x}+4\) * BGSTEP). BGSTEP
must be positive.
Default: BGSTEP \(=.01\).
TOL - Relative error desired in the derivative estimate. (Input)
Default: TOL \(=1 . \mathrm{e}-2\) for single precision and 1.d-4 for double precision.

\section*{FORTRAN 90 Interface}

Generic: DERIV (FCN, X [, ...])
Specific: The specific interface names are S_DERIV and D_DERIV.

\section*{FORTRAN 77 Interface}

Single: DERIV (FCN, KORDER, X, BGSTEP, TOL)
Double: The double precision function name is DDERIV.

\section*{Description}

DERIV produces an estimate to the first, second, or third derivative of a function. The estimate originates from first computing a spline interpolant to the input function using values within the interval ( \(\mathrm{X}-4.0 *\) BGSTEP, \(\mathrm{X}+4.0 *\) BGSTEP ), then differentiating the spline at X .

\section*{Comments}
1. Informational errors
Type Code

32 Roundoff error became dominant before estimates converged. Increase precision and/or increase BGSTEP.
41 Unable to achieve desired tolerance in derivative estimation. Increase precision, increase TOL and/or change BGSTEP. If this error continues, the function may not have a derivative at X .
2. Convergence is assumed when
\[
\frac{2}{3}|\mathrm{D} 2-\mathrm{D} 1|<\mathrm{TOL}
\]
for two successive derivative estimates D1 and D2.
3. The initial step size, BGSTEP, must be chosen small enough that FCN is defined and reasonably smooth in the interval ( \(\mathrm{X}-4\) * BGSTEP, \(\mathrm{X}+4\) * BGSTEP), yet large enough to avoid roundoff problems.

\section*{Example 1}

In this example, we obtain the approximate first derivative of the function
\[
f(x)=-2 \sin (3 x / 2)
\]
at the point \(x=2\).
```

USE DERIV_INT

```
USE UMACH_INT
IMPLICIT NONE
INTEGER KORDER, NCOUNT, NOUT
REAL BGSTEP, DERV, TOL, X
EXTERNAL FCN
```

!

```
CALL UMACH (2, NOUT)
\(!\)
\(\mathrm{X}=2.0\)
BGSTEP \(=0.2\)
NCOUNT = 1
DERV = DERIV (FCN, X, BGSTEP=BGSTEP)
WRITE (NOUT, 99999) DERV
99999 FORMAT (/, 1X, 'First derivative of FCN is ', 1PE10.3)
END
!
```

REAL FUNCTION FCN (X)
REAL X
REAL SIN
INTRINSIC SIN
FCN = -2.0*SIN(1.5*X)
RETURN
END

```

\section*{Output}
```

First derivative of FCN is 2.970E+00

```

\section*{Additional Example}

\section*{Example 2}

In this example, we attempt to approximate in single precision the third derivative of the function
\[
f(x)=2 x^{4}+3 x
\]
at the point \(x=0.75\). Although the function is well-behaved near \(x=0.75\), finding derivatives is often computationally difficult on 32-bit machines. The difficulty is overcome in double precision.
```

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER KORDER, NOUT
REAL BGSTEP, DERV, X, TOL
DOUBLE PRECISION DBGSTE, DDERV, DFCN, DTOL, DX
EXTERNAL DFCN, FCN
CALL UMACH (2, NOUT)
CALL ERSET (0, -1, 0)
X = 0.75
BGSTEP = 0.1
KORDER = 3
In single precision, on a 32-bit
machine, the following attempt
produces an error message
DERV = DERIV(FCN, X, KORDER, BGSTEP,TOL)
In double precision, we get good
results
DX = 0.75D0
DBGSTE = 0.1D0
DTOL = 0.01D0
KORDER = 3
DDERV = DERIV(DFCN, DX,KORDER, DBGSTE, DTOL)
WRITE (NOUT,99999) DDERV
99999 FORMAT (/, 1X, 'The third derivative of DFCN is ', 1PD10.4)
END
REAL FUNCTION FCN (X)

```
\(!\)
```

        REAL X
        FCN = 2.0*X**4 + 3.0*X
        RETURN
        END
    !
DOUBLE PRECISION FUNCTION DFCN (X)
DOUBLE PRECISION X
DFCN = 2.0DO*X**4 + 3.0DO*X
RETURN
END

```

\section*{Output}
```

*** FATAL
***
*** BGSTEP = 1.000000E-01. If this error continues the function
*** may not have a derivative at X = 7.500000E-01
Increase precision, increase TOL = 1.000000E-02 and/or change
The third derivative of DFCN is 3.6000D+01

```

\section*{Chapter 5: Differential Equations}

\section*{Routines}
5.1. First-Order Ordinary Differential Equations
5.1.1 Solution of the Initial-Value Problem for ODEs
IVPRK ..... 927
Runge-Kutta method
IVMRK ..... 934
Runge-Kutta method, various orders
IVPAG ..... 944
5.1.2 Solution of the Boundary-Value Problem for ODEs
Finite-difference method ..... BVPFD ..... 961
Multiple-shooting method BVPMS ..... 973
5.1.3 Solution of Differential-Algebraic Systems Petzold-Gear method DASPG ..... 980
5.2 Partial Differential Equations
5.2.1 Solution of Systems of PDEs in One Dimension Method of lines with Variable Griddings PDE_1D_MG ..... 1004
Method of lines with a Hermite cubic basis MOLCH ..... 1038
5.2.2 Solution of a PDE in Two and Three Dimensions
Two-dimensional fast Poisson solver FPS2H ..... 1053
Three-dimensional fast Poisson solver. FPS3H ..... 1059
5.3. Sturm-Liouville ProblemsEigenvalues, eigenfunctions,and spectral density functionsSLEIG1066
Indices of eigenvalues SLCNT ..... 1078

\section*{Usage Notes}

A differential equation is an equation involving one or more dependent variables (called \(y_{i}\) or \(u_{i}\) ), their derivatives, and one or more independent variables (called \(t, x\), and \(y\) ). Users will typically need to relabel their own model variables so that they correspond to the variables used in the solvers described here. A differential equation with one independent variable is called an ordinary differential equation (ODE). A system of equations involving derivatives in one independent variable and other dependent variables is called a differential-algebraic system. A differential equation with more than one independent variable is called a partial differential equation (PDE).

The order of a differential equation is the highest order of any of the derivatives in the equation. Some of the routines in this chapter require the user to reduce higher-order problems to systems of first-order differential equations.

\section*{Ordinary Differential Equations}

It is convenient to use the vector notation below. We denote the number of equations as the value \(N\). The problem statement is abbreviated by writing it as a system of first-order ODEs
\[
y(t)=\left[y_{1}(t), \ldots, y_{N}(t)\right]^{T}, f(t, y)=\left[f_{1}(t, y), \ldots, f_{N}(t, y)\right]^{T}
\]

The problem becomes
\[
y^{\prime}=\frac{d y(t)}{d t}=f(t, y)
\]
with initial values \(y\left(t_{0}\right)\). Values of \(y(t)\) for \(t>t_{0}\) or \(t<t_{0}\) are required. The routines IVPRK, IVMRK, and IVPAG, solve the IVP for systems of ODEs of the form \(y^{\prime}=f(t, y)\) with \(y\left(t=t_{0}\right)\) specified. Here, \(f\) is a user supplied function that must be evaluated at any set of values \(\left(t, y_{1}, \ldots, y_{N}\right)\); \(i=1, \ldots, N\). The routines IVPAG, and DASPG, will also solve implicit systems of the form \(A y^{\prime}=f\) \((t, y)\) where \(A\) is a user supplied matrix. For IVPAG, the matrix \(A\) must be nonsingular.

The system \(y^{\prime}=f(t, y)\) is said to be stiff if some of the eigenvalues of the Jacobian matrix \(\left\{\partial f_{i} / \partial y_{j}\right\}\) have large, negative real parts. This is often the case for differential equations representing the behavior of physical systems such as chemical reactions proceeding to equilibrium where subspecies effectively complete their reaction in different epochs. An alternate model concerns discharging capacitors such that different parts of the system have widely varying decay rates (or time constants). This definition of stiffness, based on the eigenvalues of the Jacobian matrix, is not satisfactory. Users typically identify stiff systems by the fact that numerical differential equation solvers such as IVPRK, are inefficient, or else they fail. The most common inefficiency is that a large number of evaluations of the functions \(f_{i}\) are required. In such cases, use routine IVPAG, or DASPG. For more about stiff systems, see Gear (1971, Chapter 11) or Shampine and Gear (1979).

In the boundary value problem (BVP) for ODEs, constraints on the dependent variables are given at the endpoints of the interval of interest, \([a, b]\). The routines BVPFD and BVPMS solve the BVP for systems of the form \(y^{\prime}(t)=f(t, y)\), subject to the conditions
\[
h_{i}\left(y_{1}(a), \ldots, y_{N}(a), y_{1}(b), \ldots, y_{N}(b)\right)=0 \quad i=1, \ldots, N
\]

Here, \(f\) and \(h=\left[h_{1}, \ldots, h_{N}\right]^{T}\) are user-supplied functions.

\section*{Differential-algebraic Equations}

Frequently, it is not possible or not convenient to express the model of a dynamical system as a set of ODEs. Rather, an implicit equation is available in the form
\[
g_{i}\left(t, y, \ldots, y_{N}, y_{1}^{\prime}, \ldots, y_{N}^{\prime}\right)=0 \quad i=1, \ldots, N
\]

The \(g_{i}\) are user-supplied functions. The system is abbreviated as
\[
g\left(t, y, y^{\prime}\right)=\left[g_{1}\left(t, y, y^{\prime}\right), \ldots, g_{N}\left(t, y, y^{\prime}\right)\right]^{T}=0
\]

With initial value \(y\left(t_{0}\right)\). Any system of ODEs can be trivially written as a differential-algebraic system by defining
\[
g\left(t, y, y^{\prime}\right)=f(t, y)-y^{\prime}
\]

The routine DASPG solves differential-algebraic systems of index 1 or index 0 . For a definition of index of a differential-algebraic system, see (Brenan et al. 1989). Also, see Gear and Petzold (1984) for an outline of the computing methods used.

\section*{Partial Differential Equations}

The routine MOLCH solves the IVP problem for systems of the form
\[
\frac{\partial u_{i}}{\partial t}=f_{i}\left(x, t, u_{1}, \ldots, u_{N}, \frac{\partial u_{1}}{\partial x}, \ldots, \frac{\partial u_{N}}{\partial x}, \frac{\partial^{2} u_{1}}{\partial x^{2}}, \ldots, \frac{\partial^{2} u_{N}}{\partial x^{2}}\right)
\]
subject to the boundary conditions
\[
\begin{aligned}
& \alpha_{1}^{(i)} u_{i}(a)+\beta_{1}^{(i)} \frac{\partial u_{i}}{\partial x}(a)=\gamma_{1}(t) \\
& \alpha_{2}^{(i)} u_{i}(b)+\beta_{2}^{(i)} \frac{\partial u_{i}}{\partial x}(b)=\gamma_{2}(t)
\end{aligned}
\]
and subject to the initial conditions
\[
u_{i}\left(x, t=t_{0}\right)=g_{i}(x)
\]
for \(i=1, \ldots, N\). Here, \(f_{i}, g_{i}\),
\[
\alpha_{j}^{(i)} \text {, and } \beta_{j}^{(i)}
\]
are user-supplied, \(j=1,2\).
The routines FPS2H and FPS 3H solve Laplace's, Poisson's, or Helmholtz's equation in two or three dimensions. FPS2H uses a fast Poisson method to solve a PDE of the form
\[
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}+c u=f(x, y)
\]
over a rectangle, subject to boundary conditions on each of the four sides. The scalar constant \(c\) and the function \(f\) are user specified. FPS3H solves the three-dimensional analogue of this problem.

Users wishing to solve more general PDE's, in more general 2-d and 3-d regions are referred to Visual Numerics' partner PDE2D (www.pde2d.com).

\section*{Summary}

The following table summarizes the types of problems handled by the routines in this chapter. With the exception of FPS2H and FPS3H, the routines can handle more than one differential equation.
\begin{tabular}{|c|c|c|}
\hline Problem & Consideration & Routine \\
\hline \[
\begin{aligned}
& A y^{\prime}=f(t, y) \\
& y\left(t_{0}\right)=y_{0}
\end{aligned}
\] & \(A\) is a general, symmetric positive definite, band or symmetric positive definite band matrix. & IVPAG \\
\hline & Stiff or expensive to evaluate \(f(t, y)\), banded Jacobian or finely spaced output needed. & IVPAG \\
\hline \[
\begin{aligned}
& y^{\prime}=f(t, y), \\
& y\left(t_{0}\right)=y_{0},
\end{aligned}
\] & High accuracy needed and not stiff. (Uses Adams methods) & IVPAG \\
\hline & Moderate accuracy needed and not stiff. & IVPRK \\
\hline \[
\begin{aligned}
& y^{\prime}=f(t, y) \\
& h(y(a), y(b))=0
\end{aligned}
\] & BVP solver using finite differences & BVPFD \\
\hline & BVP solver using multiple shooting & BVPMS \\
\hline \[
\begin{aligned}
& g\left(t, y, y^{\prime}\right)=0 \\
& y\left(t_{0}\right), y^{\prime}\left(t_{0}\right) \text { given }
\end{aligned}
\] & \begin{tabular}{l}
Stiff, differential-algebraic solver for systems of index 1 or 0 . \\
Note: DASPG uses the user-supplied \(y^{\prime}\left(t_{0}\right)\) only as an initial guess to help it find the correct initial \(y^{\prime}\left(t_{0}\right)\) to get started.
\end{tabular} & DASPG \\
\hline \[
\begin{aligned}
& u_{t}=f\left(x, t, u, u_{x}, u_{x x}\right) \\
& \alpha_{1} u(a)+\beta_{1} u_{x}(a)=\gamma_{1}(t) \\
& \alpha_{2} u(b)+\beta_{2} u_{x}(b)=\gamma_{2}(t)
\end{aligned}
\] & Method of lines using cubic splines and ODEs. & MOLCH \\
\hline \(u_{x x}+u_{y y}+c u=f(x, y)\) on a rectangle, given \(u\) or \(u_{n}\) on each edge. & Fast Poisson solver & FPS2H \\
\hline \begin{tabular}{l}
\[
u_{x x}+u_{y y}+u_{z z}+c u=f(x, y, z)
\] \\
on a box, given \(u\) or \(u_{n}\) on each face
\end{tabular} & Fast Poisson solver & FPS3H \\
\hline \[
\begin{aligned}
& -\left(p u^{\prime}\right)^{\prime}+q u=\lambda r u, \\
& \alpha_{1} u(a)-\alpha_{2}\left(p u^{\prime}(a)\right) \\
& =\lambda\left(\alpha_{1}^{\prime} u(a)-\alpha_{2}^{\prime}\left(p u^{\prime}(a)\right)\right) \\
& \beta_{1} u(b)+\beta_{2}\left(p u^{\prime}(b)\right)=0
\end{aligned}
\] & Sturm-Liouville problems & SLEIG \\
\hline
\end{tabular}

\section*{IVPRK}

Solves an initial-value problem for ordinary differential equations using the Runge-Kutta-Verner fifth-order and sixth-order method.

\section*{Required Arguments}

IDO - Flag indicating the state of the computation. (Input/Output)

\section*{IDO State}

1 Initial entry
2 Normal re-entry
3 Final call to release workspace
4 Return because of interrupt 1
5 Return because of interrupt 2 with step accepted
6 Return because of interrupt 2 with step rejected
Normally, the initial call is made with \(\operatorname{IDO}=1\). The routine then sets \(\operatorname{IDO}=2\), and this value is used for all but the last call that is made with \(\operatorname{IDO}=3\). This final call is used to release workspace, which was automatically allocated by the initial call with \(\operatorname{IDO}=1\). No integration is performed on this final call. See Comment 3 for a description of the other interrupts.

FCN - User-supplied SUBRout Ine to evaluate functions. The usage is
CALL \(\operatorname{FCN}(\mathbb{N}, \mathrm{T}, \mathrm{Y}, \mathrm{YPRIME}\) ), where
N - Number of equations. (Input)
T - Independent variable, \(t\). (Input)
Y - Array of size N containing the dependent variable values, \(y\).
(Input)
YPRIME - Array of size N containing the values of the vector \(y^{\prime}\)
evaluated at \((t, y)\). (Output)
FCN must be declared EXTERNAL in the calling program.
\(\boldsymbol{T}\) - Independent variable. (Input/Output)
On input, \(T\) contains the initial value. On output, \(T\) is replaced by TEND unless error conditions have occurred. See IDO for details.

TEND - Value of \(t\) where the solution is required. (Input)
The value tend may be less than the initial value of \(t\).
\(\boldsymbol{Y}\) - Array of size NEQ of dependent variables. (Input/Output)
On input, y contains the initial values. On output, Y contains the approximate solution.

\section*{Optional Arguments}

NEQ - Number of differential equations. (Input)
Default: \(\mathrm{NEQ}=\operatorname{size}(\mathrm{Y}, 1)\).
TOL - Tolerance for error control. (Input)
An attempt is made to control the norm of the local error such that the global error is proportional to TOL.
Default: TOL = machine precision.
PARAM - A floating-point array of size 50 containing optional parameters. (Input/ Output) If a parameter is zero, then a default value is used. These default values are given below. Parameters that concern values of step size are applied in the direction of integration. The following parameters may be set by the user:

PARAM Meaning
1 HINIT Initial value of the step size. Default: 10.0 * MAX (AMACH (1), \(\operatorname{AMACH}(4) * \operatorname{mAX}(\operatorname{ABS}(T E N D), \operatorname{ABS}(T)))\)
2 HMIN Minimum value of the step size. Default: 0.0
3 HMAX Maximum value of the step size. Default: 2.0
4 MXSTEP Maximum number of steps allowed. Default: 500
5 MXFCN Maximum number of function evaluations allowed. Default: No enforced limit.

6
Not used.
7 INTRP1 If nonzero, then return with IDO \(=4\) before each step. See Comment 3. Default: 0 .

8
INTRP2 If nonzero, then return with IDO \(=5\) after every successful step and with IDO \(=6\) after every unsuccessful step. See Comment 3. Default: 0 .
9 SCALE A measure of the scale of the problem, such as an approximation to the average value of a norm of the Jacobian matrix along the solution. Default: 1.0
10 INORM Switch determining error norm. In the following, \(e_{i}\) is the absolute value of an estimate of the error in \(y_{i}(t)\).
Default: \(0.0-\min\) (absolute error, relative error) \(=\max \left(e_{i} / w_{i}\right)\); \(i=1, \ldots, N E Q\), where \(w_{i}=\max \left(\left|y_{i}(t)\right|, 1.0\right)\).
\(1-\) absolute error \(=\max \left(e_{i}\right), i=1 \ldots\), NEQ.
\(2-\max \left(e_{i} / w_{i}\right), i=1 \ldots, N E Q\) where \(w_{i}=\max \left(\left|y_{i}(t)\right|\right.\), FLOOR \()\), and FLOOR is PARAM(11).
3 - Scaled Euclidean norm defined as
where \(w_{i}=\max \left(\left|y_{i}(t)\right|, 1.0\right)\). Other definitions of YMAX can be specified by the user, as explained in Comment 1.
\begin{tabular}{ll}
\(11 \quad\) FLOOR & \begin{tabular}{l} 
Used in the norm computation associated with parameter \\
INORM. Default: 1.0.
\end{tabular} \\
\(12-30\) & Not used.
\end{tabular}

The following entries in PARAM are set by the program.
\begin{tabular}{lll} 
& PARAM & Meaning \\
31 & HTRIAL & Current trial step size. \\
32 & HMINC & Computed minimum step size allowed. \\
33 & HMAXC & Computed maximum step size allowed. \\
34 & NSTEP & Number of steps taken. \\
35 & NFCN & Number of function evaluations used. \\
\(36-50\) & & Not used.
\end{tabular}

\section*{FORTRAN 90 Interface}

Generic: CALL IVPRK (IDO, FCN, T, TEND, Y [,...])
Specific: The specific interface names are S_IVPRK and D_IVPRK.

\section*{FORTRAN 77 Interface}

Single: CALL IVPRK (IDO, NEQ, FCN, T, TEND, TOL, PARAM, Y)
Double: The double precision name is DIVPRK.

\section*{Description}

Routine IVPRK finds an approximation to the solution of a system of first-order differential equations of the form \(y_{0}=f(t, y)\) with given initial data. The routine attempts to keep the global error proportional to a user-specified tolerance. This routine is efficient for nonstiff systems where the derivative evaluations are not expensive.

The routine IVPRK is based on a code designed by Hull, Enright and Jackson (1976, 1977). It uses Runge-Kutta formulas of order five and six developed by J. H. Verner.

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of I2 PRK/DI2 PRK. The reference is:

CALL I2PRK (IDO, NEQ, FCN, T, TEND, TOL, PARAM, Y, VNORM, WK)
The additional arguments are as follows: YMAX \(=\sqrt{\sum_{i=1}^{N E Q} e_{i}^{2} / w_{i}^{2}}\)

VNORM - A Fortran SUBROUTINE to compute the norm of the error. (Input)
The routine may be provided by the user, or the IMSL routine I3PRK/DI3PRK may be used. In either case, the name must be declared in a Fortran EXTERNAL statement. If usage of the IMSL routine is intended, then the name I3PRK/DI3PRK should be used. The usage of the error norm routine is CALL VNORM ( \(\mathrm{N}, \mathrm{V}, \mathrm{Y}, \mathrm{YMAX}, ~ E N O R M\) ), where

\section*{Arg Definition}

N Number of equations. (Input)
V Array of size N containing the vector whose norm is to be computed. (Input)

Y Array of size N containing the values of the dependent variable. (Input)
YMAX Array of size \(n\) containing the maximum values of \(|y(t)|\). (Input)
ENORM Norm of the vector V. (Output)
VNORM must be declared EXTERNAL in the calling program.
WK - Work array of size 10 N using the working precision. The contents of wk must not be changed from the first call with \(I D O=1\) until after the final call with \(I D O=3\).
2. Informational errors

Type Code
41 Cannot satisfy error condition. The value of TOL may be too small.
42 Too many function evaluations needed.
43 Too many steps needed. The problem may be stiff.
3. If \(\operatorname{PARAM}(7)\) is nonzero, the subroutine returns with \(I D O=4\) and will resume calculation at the point of interruption if re-entered with IDO \(=4\). If \(\operatorname{PARAM}(8)\) is nonzero, the subroutine will interrupt the calculations immediately after it decides whether or not to accept the result of the most recent trial step. The values used are \(I D O=5\) if the routine plans to accept, or IDO \(=6\) if it plans to reject the step. The values of IDO may be changed by the user (by changing IDO from 6 to 5) in order to force acceptance of a step that would otherwise be rejected. Some parameters the user might want to examine after return from an interrupt are IDO, HTRIAL, NSTEP, NFCN, \(T\), and \(Y\). The array Y contains the newly computed trial value for \(y(t)\), accepted or not.

\section*{Example 1}

Consider a predator-prey problem with rabbits and foxes. Let \(r\) be the density of rabbits and let \(f\) be the density of foxes. In the absence of any predator-prey interaction, the rabbits would increase at a rate proportional to their number, and the foxes would die of starvation at a rate proportional to their number. Mathematically,
\[
\begin{gathered}
r^{\prime}=2 r \\
f^{\prime}=-f
\end{gathered}
\]

The rate at which the rabbits are eaten by the foxes is \(2 r f\), and the rate at which the foxes increase, because they are eating the rabbits, is \(r f\). So, the model to be solved is
\[
\begin{aligned}
& r^{\prime}=2 r-2 r f \\
& f^{\prime}=-f+r f
\end{aligned}
\]

The initial conditions are \(r(0)=1\) and \(f(0)=3\) over the interval \(0 \leq t \leq 10\).
In the program \(\mathrm{Y}(1)=r\) and \(\mathrm{Y}(2)=f\). Note that the parameter vector PARAM is first set to zero with IMSL routine SSET (Chapter 9, Basic Matrix/Vector Operations). Then, absolute error control is selected by setting \(\operatorname{PARAM}(10)=1.0\).

The last call to IVPRK with IDO \(=3\) deallocates IMSL workspace allocated on the first call to IVPRK. It is not necessary to release the workspace in this example because the program ends after solving a single problem. The call to release workspace is made as a model of what would be needed if the program included further calls to IMSL routines.
```

USE IVPRK_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER MXPARM, N
PARAMETER (MXPARM=50, N=2)
INTEGER IDO, ISTEP, NOUT
REAL PARAM(MXPARM), T, TEND, TOL, Y(N)
EXTERNAL FCN
CALL UMACH (2, NOUT)
T = 0.0
Y(1) = 1.0
Y(2)=3.0
TOL = 0.0005
PARAM = 0.EO
PARAM(10) = 1.0
! Print header
WRITE (NOUT,99999)
IDO = 1
ISTEP = 0
10 CONTINUE
ISTEP = ISTEP + 1
TEND = ISTEP
CALL IVPRK (IDO, FCN, T, TEND, Y, TOL=TOL, PARAM=PARAM)
IF (ISTEP .LE. 10) THEN
WRITE (NOUT,'(I6,3F12.3)') ISTEP, T, Y
Final call to release workspace

```
```

            IF (ISTEP .EQ. 10) IDO = 3
            GO TO 10
        END IF
    9 9 9 9 9 ~ F O R M A T ~ ( 4 X , ~ ' I S T E P ' , ~ 5 X , ~ ' T i m e ' , ~ 9 X , ~ ' Y 1 ' , ~ 1 1 X , ~ ' Y 2 ' ) ~
END
SUBROUTINE FCN (N, T, Y, YPRIME)
!
INTEGER N
REAL T, Y(N), YPRIME (N)
YPRIME(1) = 2.0*Y(1) - 2.0*Y(1)*Y(2)
YPRIME (2) = -Y(2) + Y(1)*Y(2)
RETURN
END

```

\section*{Output}
\begin{tabular}{lrrr} 
ISTEP & Time & Y1 & Y2 \\
1 & 1.000 & 0.078 & 1.465 \\
2 & 2.000 & 0.085 & 0.578 \\
3 & 3.000 & 0.292 & 0.250 \\
4 & 4.000 & 1.449 & 0.187 \\
5 & 5.000 & 4.046 & 1.444 \\
6 & 6.000 & 0.176 & 2.256 \\
7 & 7.000 & 0.066 & 0.908 \\
8 & 8.000 & 0.148 & 0.367 \\
9 & 9.000 & 0.655 & 0.188 \\
10 & 10.000 & 3.157 & 0.352
\end{tabular}

\section*{Additional Examples}

\section*{Example 2}

This is a mildly stiff problem (F2) from the test set of Enright and Pryce (1987). It is included here because it illustrates the inefficiency of requiring more function evaluations with a nonstiff solver, for a requested accuracy, than would be required using a stiff solver. Also, see IVPAG Example 2, where the problem is solved using a BDF method. The number of function evaluations may vary, depending on the accuracy and other arithmetic characteristics of the computer. The test problem has \(n=2\) equations:
\[
\begin{array}{ll}
y_{1}^{\prime} & =-y_{1}-y_{1} y_{2}+k_{1} y_{2} \\
y_{2}^{\prime} & =-k_{2} y_{2}+k_{3}\left(1-y_{2}\right) y_{1} \\
y_{1}(0) & =1 \\
y_{2}(0) & =0 \\
k_{1} & =294 \\
k_{2} & =3 \\
k_{3} & =0.01020408 \\
\text { tend } & =240
\end{array}
\]

USE IVPRK_INT
```

    USE UMACH_INT
    IMPLICIT NONE
    INTEGER MXPARM, N
    PARAMETER (MXPARM=50, N=2)
                SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER IDO, ISTEP, NOUT
    REAL PARAM(MXPARM), T, TEND, TOL, Y (N)
                                    SPECIFICATIONS FOR SUBROUTINES
                                    SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL FCN
    CALL UMACH (2, NOUT)
        Set initial conditions
    T = 0.0
    Y(1) = 1.0
    Y(2) = 0.0
    TOL = 0.001
    Set PARAM to default
    PARAM = 0.0EO
    PARAM(10) = 1.0
    WRITE (NOUT,99998)
    IDO = 1
    ISTEP = 0
    10 CONTINUE
    ISTEP = ISTEP + 24
    TEND = ISTEP
    CALL IVPRK (IDO, FCN, T, TEND, Y, TOL=TOL, PARAM=PARAM)
    IF (ISTEP .LE. 240) THEN
        WRITE (NOUT,'(I6,3F12.3)') ISTEP/24, T, Y
                            Final call to release workspace
        IF (ISTEP .EQ. 240) IDO = 3
        GO TO 10
    END IF
    WRITE (NOUT,99999) PARAM(35)
    99998 FORMAT (4X, 'ISTEP', 5X, 'Time', 9X, 'Y1', 11X, 'Y2')
99999 FORMAT (4X, 'Number of fcn calls with IVPRK =', F6.0)
END
SUBROUTINE FCN (N, T, Y, YPRIME)
SPECIFICATIONS FOR ARGUMENTS
INTEGER N
REAL T, Y(N), YPRIME (N)
SPECIFICATIONS FOR DATA VARIABLES
REAL AK1, AK2, AK3
DATA AK1, AK2, AK3/294.0E0, 3.0E0, 0.01020408E0/
YPRIME (1) = -Y(1) - Y(1)*Y(2) + AK1*Y(2)
YPRIME (2) = -AK2*Y(2) + AK3*(1.0E0-Y(2))*Y(1)
RETURN
END

```

\section*{Output}
\begin{tabular}{lccr} 
ISTEP & Time & Y1 & Y2 \\
1 & 24.000 & 0.688 & 0.002 \\
2 & 48.000 & 0.634 & 0.002 \\
3 & 72.000 & 0.589 & 0.002 \\
4 & 96.000 & 0.549 & 0.002 \\
5 & 120.000 & 0.514 & 0.002 \\
6 & 144.000 & 0.484 & 0.002 \\
7 & 168.000 & 0.457 & 0.002 \\
8 & 192.000 & 0.433 & 0.001 \\
9 & 216.000 & 0.411 & 0.001 \\
10 & 240.000 & 0.391 & 0.001 \\
Number of fcn calls with IVPRK \(=2153\).
\end{tabular}

\section*{IVMRK}

Solves an initial-value problem \(y^{\prime}=f(t, y)\) for ordinary differential equations using Runge-Kutta pairs of various orders.

\section*{Required Arguments}

IDO - Flag indicating the state of the computation. (Input/Output)
IDO State
1 Initial entry
2 Normal re-entry
3 Final call to release workspace
4 Return after a step
5 Return for function evaluation (reverse communication)
Normally, the initial call is made with \(I D O=1\). The routine then sets \(I D O=2\), and this value is used for all but the last call that is made with IDO \(=3\). This final call is used to release workspace, which was automatically allocated by the initial call with \(I D O=1\).

FCN - User-supplied SUBROUTINE to evaluate functions. The usage is
CALL FCN (N, T, Y, YPRIME), where
N - Number of equations. (Input)
T - Independent variable. (Input)
Y - Array of size N containing the dependent variable values, \(y\). (Input)
YPRIME - Array of size N containing the values of the vector \(y^{\prime}\) evaluated at \((t, y)\).
(Output)
ECN must be declared EXTERNAL in the calling program.
\(\boldsymbol{T}\) - Independent variable. (Input/Output)
On input, \(T\) contains the initial value. On output, \(T\) is replaced by TEND unless error conditions have occurred.

TEND - Value of \(t\) where the solution is required. (Input)
The value of TEND may be less than the initial value of \(t\).
\(\boldsymbol{Y}\) - Array of size N of dependent variables. (Input/Output)
On input, Y contains the initial values. On output, Y contains the approximate solution.
YPRIME - Array of size N containing the values of the vector \(y^{\prime}\) evaluated at \((t, y)\).
(Output)

\section*{Optional Arguments}
\(N\) - Number of differential equations. (Input)
Default: \(N=\operatorname{size}(Y, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL IVMRK (IDO, FCN, T, TEND, Y, YPRIME [,...])

Specific: The specific interface names are S_IVMRK and D_IVMRK.

\section*{FORTRAN 77 Interface}

Single: CALL IVMRK (IDO, N, FCN, T, TEND, Y, YPRIME)
Double: The double precision name is DIVMRK.

\section*{Description}

Routine IVMRK finds an approximation to the solution of a system of first-order differential equations of the form \(y^{\prime}=f(t, y)\) with given initial data. Relative local error is controlled according to a user-supplied tolerance. For added efficiency, three Runge-Kutta formula pairs, of orders 3, 5, and 8 , are available.

Optionally, the values of the vector \(y^{\prime}\) can be passed to IVMRK by reverse communication, avoiding the user-supplied subroutine FCN. Reverse communication is especially useful in applications that have complicated algorithmic requirement for the evaluations of \(f(t, y)\). Another option allows assessment of the global error in the integration.

The routine IVMRK is based on the codes contained in RKSUITE, developed by R. W. Brankin, I. Gladwell, and L. F. Shampine (1991).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of I2MRK/DI2MRK. The reference is:

CALL I2MRK (IDO, N, FCN, T, TEND, Y, YPRIME, TOL, THRES, PARAM, YMAX, RMSERR, WORK, IWORK)

The additional arguments are as follows:

TOL - Tolerance for error control. (Input)
THRES - Array of size N. (Input)
THRES (I) is a threshold for solution component \(Y\) (I). It is chosen so that the value of \(Y(L)\) is not important when \(Y(L)\) is smaller in magnitude than THRES (L) . THRES (L) must be greater than or equal to sqrt (amach (4)).

PARAM - A floating-point array of size 50 containing optional parameters. (Input/Output)
If a parameter is zero, then a default value is used. These default values are given below. The following parameters must be set by the user:

PARAM Meaning
1 HINIT Initial value of the step size. Must be chosen such that \(0.01 \geq\) HINIT \(\geq 10.0\) amach(4).
Default: automatic selection of stepsize.
2 METHOD Specify which Runge-Kutta pair is to be used.
1 - use the \((2,3)\) pair
2 - use the \((4,5)\) pair
3 - use the \((7,8)\) pair.
Default: METHOD \(=1\) if \(1 . e-2 \geq\) tol \(>1 . e-4\)
METHOD \(=2\) if \(1 . \mathrm{e}-4 \geq\) tol \(>1 . \mathrm{e}-6\)
METHOD \(=3\) if \(1 . e-6 \geq\) tol
3 ERREST ERREST \(=1\) attempts to assess the true error, the difference between the numerical solution and the true solution. The cost of this is roughly twice the cost of the integration itself with METHOD \(=2\) or METHOD \(=3\), and three times with METHOD \(=1\).
Default: ERREST \(=0\).
4 INTRP If nonzero, then return the IDO \(=4\) before each step. See Comment 3. Default: 0

5 RCSTAT If nonzero, then reverse communication is used to get derivative information. See Comment 4. Default: 0.

6-30 Not used
The following entries are set by the program:
31 HTRIAL Current trial step size.
32 nstep Number of steps taken.
33 NFCN Number of function evaluations.
34 ERRMAX The maximum approximate weighted true error taken over all solution components and all steps from T through the current integration point.
35 TERRMX First value of the independent variable where an
approximate true error attains the maximum value ERRMAX.

YMAX Array of size N, where \(\operatorname{ymAx}(\mathrm{L})\) is the largest value of \(\operatorname{ABS}(\mathrm{Y}(\mathrm{L})\) ) computed at any step in the integration so far.

RMSERR - Array of size \(N\) where RMSERR(L) approximates the RMS average of the true error of the numerical solution for the L-th solution component, \(L=1, \ldots, N\). The average is taken over all steps from \(T\) through the current integration point. RMSERR is accessed and set only if \(\operatorname{PARAM}(3)=1\).

WORK - Floating point work array of size 39N using the working precision. The contents of WORK must not be changed from the first call with IDO \(=1\) until after the final call with \(\operatorname{IDO}=3\).

IWORK - Length of array work. (Input)
2. Informational errors

Type Code
41 It does not appear possible to achieve the accuracy specified by TOL and THRES(*) using the current precision and METHOD. A larger value for METHOD, if possible, will permit greater accuracy with this precision. The integration must be restarted.
42 The global error assessment may not be reliable beyond the current integration point T . This may occur because either too little or too much accuracy has been requested or because \(f(t, y)\) is not smooth enough for values of \(t\) just past TEND and current values of the solution \(y\). This return does not mean that you cannot integrate past TEND, rather that you cannot do it with \(\operatorname{PARAM}(3)=1\).

3 If PARAM (4) is nonzero, the subroutine returns with IDO \(=4\) and will resume calculation at the point of interruption if re-entered with \(I D O=4\). Some parameters the user might want to examine are IDO, HTRIAL, NSTEP, NFCN, T, and Y. The array Y contains the newly computed trial value for \(y(t)\), accepted or not.

4 If PARAM (5) is nonzero, the subroutine will return with IDO \(=5\). At this time, evaluate the derivatives at T , place the result in YPRIME, and call IVMRK again. The dummy function I40RK/DI40RK may be used in place of FCN.

\section*{Example 1}

This example integrates the small system (A.2.B2) from the test set of Enright and Pryce (1987):
\[
\begin{aligned}
y_{1}^{\prime} & =-y_{1}+y_{2} \\
y_{2}^{\prime} & =y_{1}-2 y_{2}+y_{3} \\
y_{3}^{\prime} & =y_{2}-y_{3} \\
y_{1}(0) & =2 \\
y_{2}(0) & =0 \\
y_{3}(0) & =1
\end{aligned}
\]
```

USE IVMRK_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=3)
INTEGER IDO
REAL T, TEND, Y(N), YPRIME(N)
EXTERNAL FCN
T = 0.0
TEND = 20.0
Y(1) = 2.0
Y(2) = 0.0
Y(3) = 1.0
IDO = 1
CALL IVMRK (IDO, FCN, T, TEND, Y, YPRIME)
Final call to release workspace
IDO = 3
CALL IVMRK (IDO, FCN, T, TEND, Y, YPRIME)
CALL WRRRN ('Y', Y)
END
SUBROUTINE FCN (N, T, Y, YPRIME)
Specifications for arguments
INTEGER N
REAL T, Y(*), YPRIME(*)
YPRIME (1) = -Y(1) + Y(2)
YPRIME (2) = Y(1) - 2.0*Y(2) + Y(3)
YPRIME (3) = Y(2) - Y(3)
RETURN
END

```

\section*{Output}

Y
\(1 \quad 1.000\)
21.000
31.000

\section*{Additional Examples}

\section*{Example 2}

This problem is the same mildly stiff problem (A.1.F2) from the test set of Enright and Pryce as Example 2 for IVPRK.
\[
\begin{aligned}
y_{1}^{\prime} & =-y_{1}-y_{1} y_{2}+k_{1} y_{2} \\
y_{2}^{\prime} & =-k_{2} y_{2}+k_{3}\left(1-y_{2}\right) y_{1} \\
y_{1}(0) & =1 \\
y_{2}(0) & =0 \\
k_{1} & =294 \\
k_{2} & =3 \\
k_{3} & =0.01020408 \\
\text { tend } & =240
\end{aligned}
\]

Although not a stiff solver, one notes the greater efficiency of IVMRK over IVPRK, in terms of derivative evaluations. Reverse communication is also used in this example. Users will find this feature particularly helpful if their derivative evaluation scheme is difficult to isolate in a separate subroutine.
```

USE I2MRK INT
USE UMACH INT
USE AMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=2)
INTEGER IDO, ISTEP, LWORK, NOUT
REAL PARAM(50), PREC, RMSERR(N), T, TEND, THRES (N), TOL, \&
WORK(1000), Y(N), YMAX(N), YPRIME (N)
REAL AK1, AK2, AK3
SAVE AK1, AK2, AK3
INTRINSIC SQRT
REAL SQRT
EXTERNAL I4ORK
DATA AK1, AK2, AK3/294.0, 3.0, 0.01020408/
CALL UMACH (2, NOUT)
T = 0.0
Y(1) = 1.0
Y(2) = 0.0
Set tolerance for error control,
threshold vector and parameter

```
```

! vector
TOL = . 001
PREC = AMACH(4)
THRES = SQRT (PREC)
PARAM = 0.0E0
LWORK = 1000
! Turn on derivative evaluation by
! reverse communication
PARAM(5) = 1
IDO = 1
ISTEP = 24
! Print header
WRITE (NOUT,99998)
1 0 ~ C O N T I N U E
TEND = ISTEP
CALL I2MRK (IDO, N, I4ORK, T, TEND, Y, YPRIME, TOL, THRES, PARAM,\&
YMAX, RMSERR, WORK, LWORK)
IF (IDO .EQ. 5) THEN
Evaluate derivatives
YPRIME (1) = -Y(1) - Y(1)*Y(2) + AK1*Y(2)
YPRIME (2) = -AK2*Y(2) + AK3*(1.0-Y(2))*Y(1)
GO TO 10
ELSE IF (ISTEP .LE. 240) THEN
!
!
Integrate to 10 equally spaced points
WRITE (NOUT,'(I6,3F12.3)') ISTEP/24, T, Y
IF (ISTEP .EQ. 240) IDO = 3
ISTEP = ISTEP + 24
GO TO 10
END IF
! Show number of derivative evaluations
WRITE (NOUT,99999) PARAM(33)
99998 FORMAT (3X, 'ISTEP', 5X, 'TIME', 9X, 'Y1', 10X, 'Y2')
99999 FORMAT (/, 4X, 'NUMBER OF DERIVATIVE EVALUATIONS WITH IVMRK =', \&
F6.0)
END
! DUMMY FUNCTION TO TAKE THE PLACE OF DERIVATIVE EVALUATOR
SUBROUTINE I4ORK (N, T, Y, YPRIME)
INTEGER N
REAL T, y(*), YPRIME(*)
RETURN
END

```

\section*{Output}
\begin{tabular}{lrrr} 
ISTEP & TIME & Y1 & Y2 \\
1 & 24.000 & 0.688 & 0.002 \\
2 & 48.000 & 0.634 & 0.002 \\
3 & 72.000 & 0.589 & 0.002 \\
4 & 96.000 & 0.549 & 0.002
\end{tabular}
\begin{tabular}{llll}
5 & 120.000 & 0.514 & 0.002 \\
6 & 144.000 & 0.484 & 0.002 \\
7 & 168.000 & 0.457 & 0.002 \\
8 & 192.000 & 0.433 & 0.001 \\
9 & 216.000 & 0.411 & 0.001 \\
10 & 240.000 & 0.391 & 0.001 \\
NUMBER & OF & DERIVATIVE & EVALUATIONS
\end{tabular}

\section*{Example 3}

This example demonstrates how exceptions may be handled. The problem is from Enright and Pryce (A.2.F1), and has discontinuities. We choose this problem to force a failure in the global error estimation scheme, which requires some smoothness in \(y\). We also request an initial relative error tolerance which happens to be unsuitably small in this precision.

If the integration fails because of problems in global error assessment, the assessment option is turned off, and the integration is restarted. If the integration fails because the requested accuracy is not achievable, the tolerance is increased, and global error assessment is requested. The reason error assessment is turned on is that prior assessment failures may have been due more in part to an overly stringent tolerance than lack of smoothness in the derivatives.

When the integration is successful, the example prints the final relative error tolerance, and indicates whether or not global error estimation was possible.
\[
\begin{aligned}
& y_{1}^{\prime}=y_{2} \\
& y_{2}^{\prime}=\left\{\begin{array}{l}
2 a y_{2}-\left(\pi^{2}+a^{2}\right) y_{1}+1,\lfloor x\rfloor \text { even } \\
2 a y_{2}-\left(\pi^{2}+a^{2}\right) y_{1}-1,\lfloor x\rfloor \mathrm{odd}
\end{array}\right. \\
& y_{1}(0)=0 \\
& y_{2}(0)=0 \\
& \quad a=0.1 \\
& \quad\lfloor x\rfloor=\text { largest integer } \leq x
\end{aligned}
\]

USE IMSL_LIBRARIES
IMPLICIT NONE

INTEGER N
PARAMETER ( \(\mathrm{N}=2\) )
! INTEGER IDO, LWORK, NOUT Specifications for local variables
INTEGER IDO, LWORK, NOUT
REAL PARAM (50), PREC, RMSERR(N), T, TEND, THRES (N), TOL, \& WORK (100), Y(N), YMAX(N), YPRIME(N)
\(!\)
Specifications for intrinsics
INTRINSIC SQRT
REAL SQRT
\(\begin{array}{ll}\text { ! } & \text { Specifications for subroutines } \\ \text { ! } & \\ \text { ! Specifications for functions }\end{array}\)
EXTERNAL FCN
```

!
CALL UMACH (2, NOUT)
! Turn off stopping for FATAL errors
CALL ERSET (4, -1, 0)
! Initialize input, turn on global
LWORK = 100
PREC = AMACH(4)
TOL = SQRT(PREC)
PARAM = 0.0E01
THRES = TOL
TEND = 20.0E0
PARAM(3) = 1
!
1 0 ~ C O N T I N U E
!
T = 0.0E0
Y(1) = 0.0E0
Y(2) = 0.0E0
IDO = 1
CALL I2MRK (IDO, N, FCN, T, TEND, Y, YPRIME, TOL, THRES, PARAM,\&
YMAX, RMSERR, WORK, LWORK)
IF (IERCD() .EQ. 32) THEN
Unable to achieve requested
accuracy, so increase tolerance.
Activate global error assessment
TOL = 10.0*TOL
PARAM(3) = 1
WRITE (NOUT,99995) TOL
GO TO 10
ELSE IF (IERCD() .EQ. 34) THEN
Global error assessment has failed,
cannot continue from this point,
so restart integration
WRITE (NOUT,99996)
PARAM(3) = 0
GO TO 10
END IF
!
Final call to release workspace
IDO = 3
CALL I2MRK (IDO, N, FCN, T, TEND, Y, YPRIME, TOL, THRES, PARAM,\&
YMAX, RMSERR, WORK, LWORK)
Summarize status
WRITE (NOUT,99997) TOL
IF (PARAM(3) .EQ. 1) THEN
WRITE (NOUT,99998)
ELSE
WRITE (NOUT,99999)
END IF
CALL WRRRN ('Y', Y)
!
99995 FORMAT (/, 'CHANGING TOLERANCE TO ', E9.3, ' AND RESTARTING ...'\&

```
```

        , /, 'ALSO (RE) ENABLING GLOBAL ERROR ASSESSMENT', /)
    99996 FORMAT (/, 'DISABLING GLOBAL ERROR ASSESSMENT AND RESTARTING ...'\&
, /)
99997 FORMAT (/, 72('-'), //, 'SOLUTION OBTAINED WITH TOLERANCE = ',\&
E9.3)
99998 FORMAT ('GLOBAL ERROR ASSESSMENT IS AVAILABLE')
99999 FORMAT ('GLOBAL ERROR ASSESSMENT IS NOT AVAILABLE')
!
END
!
SUBROUTINE FCN (N, T, Y, YPRIME)
USE CONST_INT
! Specifications for arguments
INTEGER N
REAL T, Y(*), YPRIME(*)
REAL PI
LOGICAL FIRST
SAVE FIRST, PI
Specifications for intrinsics
INIRINSIC INT, MOD
INTEGER INT, MOD
!
!
DATA FIRST/.TRUE./
!
IF (FIRST) THEN
PI = CONST('PI')
FIRST = .FALSE.
END IF
!
A = 0.1E0
YPRIME(1) = Y(2)
IF (MOD(INT (T), 2) .EQ. 0) THEN
YPRIME (2) = 2.0E0*A*Y(2) - (PI*PI+A*A)*Y(1) + 1.0E0
ELSE
YPRIME (2) = 2.0EO*A*Y(2) - (PI*PI+A*A)*Y(1) - 1.0E0
END IF
RETURN
END

```

\section*{Output}
```

*** FATAL ERROR 34 from i2mrk. The global error assessment may not
*** be reliable for T past 9.994749E-01. The integration is
*** being terminated.
DISABLING GLOBAL ERROR ASSESSMENT AND RESTARTING ...
*** FATAL ERROR 32 from i2mrk. In order to satisfy the error
*** requirement IGMRK would have to use a step size of
*** 3.647129E- 06 at TNOW = 9.999932E-01. This is too small

```
```

    *** for the current precision.
    CHANGING TOLERANCE TO 0.345E-02 AND RESTARTING ...
ALSO (RE) ENABLING GLOBAL ERROR ASSESSMENT
*** FATAL ERROR 34 from i2mrk. The global error assessment may
*** not be reliable for T past 9.986024E-01. The integration
*** is being terminated.
DISABLING GLOBAL ERROR ASSESSMENT AND RESTARTING ...

```
```

SOLUTION OBTAINED WITH TOLERANCE = 0.345E-02

```
SOLUTION OBTAINED WITH TOLERANCE = 0.345E-02
GLOBAL ERROR ASSESSMENT IS NOT AVAILABLE
    Y
    1 -12.30
2 0.95
```


## IVPAG

Solves an initial-value problem for ordinary differential equations using either Adams-Moulton's or Gear's BDF method.

## Required Arguments

IDO - Flag indicating the state of the computation. (Input/Output)

## IDO State

1 Initial entry
2 Normal re-entry

3 Final call to release workspace
4 Return because of interrupt 1
5 Return because of interrupt 2 with step accepted
6 Return because of interrupt 2 with step rejected

7 Return for new value of matrix A.

Normally, the initial call is made with $\operatorname{IDO}=1$. The routine then sets $I D O=2$, and this value is then used for all but the last call that is made with $I D O=3$. This final call is
only used to release workspace, which was automatically allocated by the initial call with $I D O=1$. See Comment 5 for a description of the interrupts.

When IDO $=7$, the matrix $A$ at $t$ must be recomputed and IVPAG/DIVPAG called again. No other argument (including IDO) should be changed. This value of IDO is returned only if $\operatorname{PARAM}(19)=2$.

FCN - User-supplied SUBROUTINE to evaluate functions. The usage is CALL FCN (N, T, Y, YPRIME), where
N - Number of equations. (Input)
T - Independent variable, $t$. (Input)
Y - Array of size N containing the dependent variable values, $y$. (Input)
YPRIME - Array of size N containing the values of the vector $y^{\prime}$
evaluated at $(t, y)$. (Output)
See Comment 3 .
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{F C N J}$ - User-supplied SUBROUTINE to compute the Jacobian. The usage is
CALL FCNJ ( $\mathrm{N}, \mathrm{T}, \mathrm{Y}, \mathrm{DYPDY}$ ) where
N - Number of equations. (Input)
T - Independent variable, $t$. (Input)
$\mathrm{Y}-$ Array of size N containing the dependent variable values, $y(t)$. (Input)
DYPDY - An array, with data structure and type determined by
$\operatorname{PARAM}(14)=$ MTYPE, containing the required partial derivatives $\partial f_{i} / \partial \mathrm{y}_{j}$. (Output)
These derivatives are to be evaluated at the current values of $(t, y)$. When the
Jacobian is dense, MTYPE $=0$ or $=2$, the leading dimension of DYPDY has the value $n$. When the Jacobian matrix is banded, $M T Y P E=1$, and the leading dimension of DYPDY has the value $2 *$ NLC + NUC +1 . If the matrix is banded positive definite symmetric, $\operatorname{MTYPE}=3$, and the leading dimension of DYPDY has the value NUC +1 .
FCNJ must be declared EXTERNAL in the calling program. If PARAM(19) = IATYPE is nonzero, then FCNJ should compute the Jacobian of the righthand side of the equation $A y^{\prime}=f(t, y)$. The subroutine FCNJ is used only if $\operatorname{PARAM}(13)=\operatorname{MITER}=1$.
$\boldsymbol{T}$ - Independent variable, $t$. (Input/Output)
On input, $T$ contains the initial independent variable value. On output, $T$ is replaced by TEND unless error or other normal conditions arise. See IDO for details.

TEND - Value of $t=$ tend where the solution is required. (Input)
The value tend may be less than the initial value of $t$.
$\boldsymbol{Y}$ - Array of size NEQ of dependent variables, $y(t)$. (Input/Output)
On input, Y contains the initial values, $y\left(t_{0}\right)$. On output, Y contains the approximate solution, $y(t)$.

## Optional Arguments

NEQ - Number of differential equations. (Input)
Default: $\mathrm{NEQ}=\operatorname{size}(\mathrm{Y}, 1)$
$\boldsymbol{A}$ - Matrix structure used when the system is implicit. (Input)
The matrix $A$ is referenced only if $\operatorname{PARAM}(19)=$ IATYPE is nonzero. Its data structure is determined by PARAM(14) = мTYPE. The matrix A must be nonsingular and MITER must be 1 or 2 . See Comment 3 .

TOL - Tolerance for error control. (Input)
An attempt is made to control the norm of the local error such that the global error is proportional to TOL.
Default: TOL = 001
PARAM - A floating-point array of size 50 containing optional parameters. (Input/Output) If a parameter is zero, then the default value is used. These default values are given below. Parameters that concern values of the step size are applied in the direction of integration. The following parameters may be set by the user:

|  | PARAM | Meaning |
| :---: | :---: | :---: |
| 1 | Hinit | Initial value of the step size H . Always nonnegative. Default: $0.001 \mid$ tend $-t_{0} \mid$. |
| 2 | HMIN | Minimum value of the step size H. Default: 0.0 . |
| 3 | HMAX | Maximum value of the step size H. Default: No limit, beyond the machine scale, is imposed on the step size. |
| 4 | MXSTEP | Maximum number of steps allowed. Default: 500 . |
| 5 | MXFCN | Maximum number of function evaluations allowed. Default: No enforced limit. |
| 6 | MAXORD | Maximum order of the method. Default: If Adams-Moulton method is used, then 12. If Gear's or BDF method is used, then 5 . The defaults are the maximum values allowed. |
| 7 | INTRP1 | If this value is set nonzero, the subroutine will return before every step with IDO $=4$. See Comment 5 . Default: 0 . |
| 8 | INTRP2 | If this value is nonzero, the subroutine will return after every successful step with IDO $=5$ and return with IDO $=6$ after every unsuccessful step. See Comment 5. Default: 0 |
| 9 | SCALE | A measure of the scale of the problem, such as an approximation to the average value of a norm of the Jacobian along the solution. Default: 1.0 |

Switch determining error norm. In the following, $e_{i}$ is the absolute value of an estimate of the error in $y_{i}(t)$. Default: 0.
$0-\min ($ absolute error, relative error $)=\max \left(e_{i} / w_{i}\right) ; i=1$, $\ldots, N$, where $w_{i}=\max \left(\left|y_{i}(t)\right|, 1.0\right)$.
$1-$ absolute error $=\max \left(e_{i}\right), i=1 \ldots$, NEQ.
$2-\max \left(e_{i} / w_{i}\right), i=1 \ldots, N$ where $w_{i}=\max \left(\left|y_{i}(t)\right|\right.$, FLOOR), and FLOOR is the value PARAM(11).
3 - Scaled Euclidean norm defined as

$$
\mathrm{YMAX}=\sqrt{\sum_{i=1}^{N E Q} e_{i}^{2} / w_{i}^{2}}
$$

where $w_{i}=\max \left(\left|y_{i}(t)\right|, 1.0\right)$. Other definitions of YMAX can be specified by the user, as explained in Comment 1.
11 FLOOR Used in the norm computation associated the parameter INORM. Default: 1.0.

Integration method indicator.
$1=$ METH selects the Adams-Moulton method.
$2=$ METH selects Gear's BDF method.
Default: 1.
13
MITER Nonlinear solver method indicator.
Note: If the problem is stiff and a chord or modified Newton method is most efficient, use MITER $=1$ or $=2$.
$0=$ MITER selects functional iteration. The value IATYPE must be set to zero with this option.

1 = MITER selects a chord method with a user-provided Jacobian.

2 = MITER selects a chord method with a divideddifference Jacobian.
$3=$ MITER selects a chord method with the Jacobian replaced by a diagonal matrix based on a directional derivative. The value IATYPE must be set to zero with this option.

Default: 0.

MTYPE Matrix type for $A$ (if used) and the Jacobian (if MITER = 1 or $=2$ ). When both are used, $A$ and the Jacobian must be of the same type.
$0=$ MTYPE selects full matrices.
$1=$ MTYPE selects banded matrices .
$2=$ MTYPE selects symmetric positive definite matrices.
$3=$ MTYPE selects banded symmetric positive definite matrices.

Default: 0 .
15 NLC Number of lower codiagonals, used if MTYPE $=1$.
Default: 0 .
16 NUC Number of upper codiagonals, used if MTYPE $=1$ or $\mathrm{MTYPE}=3$.
Default: 0 .
Not used.
EPS $\quad$ Relative tolerance used in computing divided difference Jacobians.

Default: SQRT(AMACH(4))
IAtype Type of the matrix $A$.
$0=\operatorname{IATYPE}$ implies $A$ is not used (the system is explicit).
$1=\operatorname{IATYPE}$ if $A$ is a constant matrix.
$2=$ IATYPE if $A$ depends on $t$.
Default: 0 .
20
LDA Leading dimension of array $A$ exactly as specified in the dimension statement in the calling program. Used if IATYPE is not zero.

Default:
$\mathrm{N} \quad$ if $\mathrm{MTYPE}=0$ or $=2$
$\mathrm{NUC}+\mathrm{NLC}+1 \quad$ if MTYPE $=1$
$\mathrm{NUC}+1 \quad$ if MTYPE $=3$
21-30
Not used.

The following entries in the array PARAM are set by the program:

## PARAM Meaning

HTRIAL Current trial step size.
HMINC Computed minimum step size.
HMAXC Computed maximum step size.

| 34 | NSTEP | Number of steps taken. |
| :--- | :--- | :--- |
| 35 | NFCN | Number of function evaluations used. |
| 36 | NJE | Number of Jacobian evaluations. |
| $37-50$ |  | Not used. |

## FORTRAN 90 Interface

Generic: CALL IVPAG (IDO, FCN, FCNJ, T, TEND, Y [,...])

Specific: The specific interface names are S_IVPAG and D_IVPAG.

## FORTRAN 77 Interface

Single: CALL IVPAG (IDO, NEQ, FCN, FCNJ, A, T, TEND, TOL, PARAM, Y)
Double: The double precision name is DIVPAG.

## Description

The routine IVPAG solves a system of first-order ordinary differential equations of the form $y^{\prime}=f(t, y)$ or $A y^{\prime}=f(t, y)$ with initial conditions where $A$ is a square nonsingular matrix of order $N$. Two classes of implicit linear multistep methods are available. The first is the implicit AdamsMoulton method (up to order twelve); the second uses the backward differentiation formulas BDF (up to order five). The BDF method is often called Gear's stiff method. In both cases, because basic formulas are implicit, a system of nonlinear equations must be solved at each step. The deriviative matrix in this system has the form $L=A+\eta J$ where $\eta$ is a small number computed by IVPAG and $J$ is the Jacobian. When it is used, this matrix is computed in the user-supplied routine FCNJ or else it is approximated by divided differences as a default. Using defaults, $A$ is the identity matrix. The data structure for the matrix $L$ may be identified to be real general, real banded, symmetric positive definite, or banded symmetric positive definite. The default structure for $L$ is real general.

## Comments

1. Workspace and a user-supplied error norm subroutine may be explicitly provided, if desired, by use of I2PAG/DI2PAG. The reference is:
```
CALL I2PAG (IDO, NEQ, FCN, FCNJ, A, T, TEND, TOL, PARAM, Y, YTEMP, YMAX,
ERROR, SAVE1, SAVE2, PW, IPVT, VNORM)
```

None of the additional array arguments should be changed from the first call with $I D O=1$ until after the final call with $I D O=3$. The additional arguments are as follows:

YTEMP - Array of size NMETH. (Workspace)
YMAX - Array of size NEQ containing the maximum Y-values computed so far. (Output)

ERROR - Array of size NEQ containing error estimates for each component of Y. (Output)

SAVE1 - Array of size NEQ. (Workspace)
SAVE2 - Array of size NEQ. (Workspace)
$\boldsymbol{P W}$ - Array of size NPW. (Workspace)
IPVT — Array of size NEQ. (Workspace)
VNORM - A Fortran SUBROUTINE to compute the norm of the error. (Input) The routine may be provided by the user, or the IMSL routine I3PRK/DI3PRK may be used. In either case, the name must be declared in a Fortran EXternal statement. If usage of the IMSL routine is intended, then the name I3PRK/DI3PRK should be specified. The usage of the error norm routine is CALL VNORM (NEQ, V, Y, YMAX, ENORM) where

## Arg. Definition

NEQ Number of equations. (Input)
V Array of size N containing the vector whose norm is to be computed. (Input)

Y Array of size N containing the values of the dependent variable. (Input)
YMAX Array of size N containing the maximum values of $|y(t)|$. (Input)
ENORM Norm of the vector V. (Output)
VNORM must be declared EXTERNAL in the calling program.
2. Informational errors

Type Code
41 After some initial success, the integration was halted by repeated error-test failures.
42 The maximum number of function evaluations have been used.
43 The maximum number of steps allowed have been used. The problem may be stiff.
44 On the next step $T+$ н will equal $T$. Either $T O L$ is too small, or the problem is stiff.
Note: If the Adams-Moulton method is the one used in the integration, then users can switch to the BDF methods. If the BDF methods are being used, then these comments are gratuitous and indicate that the problem is too stiff for this combination of method and value of TOL.

5 After some initial success, the integration was halted by a test on TOL.
6 Integration was halted after failing to pass the error test even after dividing the initial step size by a factor of $1.0 \mathrm{E}+10$. The value TOL may be too small.
7 Integration was halted after failing to achieve corrector convergence even after dividing the initial step size by a factor of $1.0 \mathrm{E}+10$. The value TOL may be too small.
8 IATYPE is nonzero and the input matrix $A$ multiplying $y^{\prime}$ is singular.
3. Both explicit systems, of the form $y^{\prime}=f(t, y)$, and implicit systems, $A y^{\prime}=f(t, y)$, can be solved. If the system is explicit, then $\operatorname{PARAM}(19)=0$; and the matrix $A$ is not referenced. If the system is implicit, then PARAM(14) determines the data structure of the array A. If $\operatorname{PARAM}(19)=1$, then $A$ is assumed to be a constant matrix. The value of $A$ used on the first call (with IDO $=1$ ) is saved until after a call with $I D O=3$. The value of A must not be changed between these calls.
If $\operatorname{PARAM}(19)=2$, then the matrix is assumed to be a function of $t$.
4. If MTYPE is greater than zero, then MITER must equal 1 or 2 .
5. If $\operatorname{PARAM}(7)$ is nonzero, the subroutine returns with IDO $=4$ and will resume calculation at the point of interruption if re-entered with IDO $=4$. If PARAM $(8)$ is nonzero, the subroutine will interrupt immediately after decides to accept the result of the most recent trial step. The value $I D O=5$ is returned if the routine plans to accept, or $I D O=6$ if it plans to reject. The value IDO may be changed by the user (by changing IDO from 6 to 5) to force acceptance of a step that would otherwise be rejected. Relevant parameters to observe after return from an interrupt are IDO, HTRIAL, NSTEP, NFCN, NJE, $T$ and Y. The array Y contains the newly computed trial value $y(t)$.

## Example 1

Euler's equation for the motion of a rigid body not subject to external forces is

$$
\begin{array}{ll}
y_{1}^{\prime}=y_{2} y_{3} & y_{1}(0)=0 \\
y_{2}^{\prime}=-y_{1} y_{3} & y_{2}(0)=1 \\
y_{3}^{\prime}=-0.51 y_{1} y_{2} & y_{3}(0)=1
\end{array}
$$

Its solution is, in terms of Jacobi elliptic functions, $y_{1}(t)=\operatorname{sn}(t ; k), y_{2}(t)=\operatorname{cn}(t ; k), y_{3}(t)=\operatorname{dn}(t ; k)$ where $k^{2}=0.51$. The Adams-Moulton method of IVPAG is used to solve this system, since this is the default. All parameters are set to defaults.

The last call to IVPAG with IDO $=3$ releases IMSL workspace that was reserved on the first call to IVPAG. It is not necessary to release the workspace in this example because the program ends after solving a single problem. The call to release workspace is made as a model of what would be needed if the program included further calls to IMSL routines.
Because $\operatorname{PARAM}(13)=\operatorname{MITER}=0$, functional iteration is used and so subroutine $\operatorname{FCNJ}$ is never called. It is included only because the calling sequence for IVPAG requires it.

```
USE IVPAG_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N, NPARAM
PARAMETER (N=3, NPARAM=50)
!
    INTEGER IDO, IEND, NOUT
    REAL A(1,1), T, TEND, TOL, Y(N)
                                    SPECIFICATIONS FOR SUBROUTINES
                                    SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL FCN, FCNJ
        Initialize
!
IDO = 1
T=0.0
Y(1) = 0.0
Y(2) = 1.0
Y(3) = 1.0
TOL = 1.0E-6
! Write title
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
!
    IEND = 0
    10 CONTINUE
    IEND = IEND + 1
    TEND = IEND
    The array a(*,*) is not used.
    CALL IVPAG (IDO, FCN, FCNJ, T, TEND, Y, TOL=TOL)
    IF (IEND .LE. 10) THEN
        WRITE (NOUT,99999) T, Y
                        Finish up
        IF (IEND .EQ. 10) IDO = 3
        GO TO 10
    END IF
99998 FORMAT (11X, 'T', 14X, 'Y(1)', 11X, 'Y(2)', 11X, 'Y(3)')
99999 FORMAT (4F15.5)
    END
!
    SUBROUTINE FCN (N, X, Y, YPRIME)
    INTEGER 
    YPRIME (1) = Y(2)*Y(3)
    YPRIME (2) = -Y(1)*Y(3)
    YPRIME (3) = -0.51*Y(1)*Y(2)
    RETURN
    END
!
    SUBROUTINE FCNJ (N, X, Y, DYPDY)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL X, Y(N), DYPDY (N, *)
```

$!$
This subroutine is never called
RETURN
END

## Output

| T | $\mathrm{Y}(1)$ | $\mathrm{Y}(2)$ | $\mathrm{Y}(3)$ |
| :---: | ---: | ---: | ---: |
| 1.00000 | 0.80220 | 0.59705 | 0.81963 |
| 2.00000 | 0.99537 | -0.09615 | 0.70336 |
| 3.00000 | 0.64141 | -0.76720 | 0.88892 |
| 4.00000 | -0.26961 | -0.96296 | 0.98129 |
| 5.00000 | -0.91173 | -0.41079 | 0.75899 |
| 6.00000 | -0.95751 | 0.28841 | 0.72967 |
| 7.00000 | -0.42877 | 0.90342 | 0.95197 |
| 8.00000 | 0.51092 | 0.85963 | 0.93106 |
| 9.00000 | 0.97567 | 0.21926 | 0.71730 |
| 10.00000 | 0.87790 | -0.47884 | 0.77906 |

## Additional Examples

## Example 2

The BDF method of IVPAG is used to solve Example 2 of IVPRK. We set PARAM(12) $=2$ to designate the BDF method. A chord or modified Newton method, with the Jacobian computed by divided differences, is used to solve the nonlinear equations. Thus, we set $\operatorname{PARAM}(13)=2$. The number of evaluations of $y^{\prime}$ is printed after the last output point, showing the efficiency gained when using a stiff solver compared to using IVPRK on this problem. The number of evaluations may vary, depending on the accuracy and other arithmetic characteristics of the computer.

```
USE IVPAG INT
USE UMACH INT
IMPLICIT NONE
INTEGER MXPARM, N
PARAMETER (MXPARM=50, N=2)
SPECIFICATIONS FOR PARAMETERS
INTEGER MABSE, MBDF, MSOLVE
PARAMETER (MABSE=1, MBDF=2, MSOLVE=2)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER IDO, ISTEP, NOUT
REAL A(1,1), PARAM(MXPARM), T, TEND, TOL, Y(N)
    SPECIFICATIONS FOR SUBROUTINES
    SPECIFICATIONS FOR FUNCTIONS
EXTERNAL FCN, FCNJ
CALL UMACH (2, NOUT)
T = 0.0
Y(1) = 1.0
Y(2) = 0.0
TOL = 0.001
Set error tolerance
Set PARAM to defaults
```

```
    PARAM = 0.0E0
!
    PARAM(10) = MABSE
    PARAM(12) = MBDF
! Select chord method and
! PARAM(13) = MSOLVE
! Print header
    WRITE (NOUT,99998)
    IDO = 1
    ISTEP = 0
    10 CONTINUE
    ISTEP = ISTEP + 24
    TEND = ISTEP
! The array a(*,*) is not used.
    CALL IVPAG (IDO, FCN, FCNJ, T, TEND, Y, TOL=TOL, &
                PARAM=PARAM)
    IF (ISTEP .LE. 240) THEN
        WRITE (NOUT,'(I6,3F12.3)') ISTEP/24, T, Y
                    Final call to release workspace
        IF (ISTEP .EQ. 240) IDO = 3
        GO TO 10
    END IF
! Show number of function calls.
    WRITE (NOUT,99999) PARAM(35)
99998 FORMAT (4X, 'ISTEP', 5X, 'Time', 9X, 'Y1', 11X, 'Y2')
99999 FORMAT (4X, 'Number of fcn calls with IVPAG =', F6.0)
    END
    SUBROUTINE FCN (N, T, Y, YPRIME)
                                    SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(N), YPRIME (N)
                                    SPECIFICATIONS FOR SAVE VARIABLES
    REAL AK1, AK2, AK3
    SAVE AK1, AK2, AK3
    DATA AK1, AK2, AK3/294.0E0, 3.0E0, 0.01020408E0/
!
    YPRIME(1) = -Y(1) - Y(1)*Y(2) + AK1*Y(2)
    YPRIME (2) = -AK2*Y(2) + AK3*(1.0E0-Y(2))*Y(1)
    RETURN
    END
    SUBROUTINE FCNJ (N, T, Y, DYPDY)
! SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(N), DYPDY (N,*)
!
    RETURN
    END
```


## Output

| ISTEP | Time | Y1 | Y2 |
| :---: | ---: | ---: | ---: |
| 1 | 24.000 | 0.689 | 0.002 |


| 2 | 48.000 | 0.636 | 0.002 |
| :---: | ---: | ---: | ---: |
| 3 | 72.000 | 0.590 | 0.002 |
| 4 | 96.000 | 0.550 | 0.002 |
| 5 | 120.000 | 0.515 | 0.002 |
| 6 | 144.000 | 0.485 | 0.002 |
| 7 | 168.000 | 0.458 | 0.002 |
| 8 | 192.000 | 0.434 | 0.001 |
| 9 | 216.000 | 0.412 | 0.001 |
| 10 | 240.000 | 0.392 | 0.001 |
| Number of fcn calls with IVPAG $=$ | 73. |  |  |

## Example 3

The BDF method of IVPAG is used to solve the so-called Robertson problem:

$$
\begin{array}{ll}
y_{1}^{\prime}=-c_{1} y_{1}+c_{2} y_{2} y_{3} & y_{1}(0)=1 \\
y_{2}^{\prime}=-y_{1}^{\prime}-y_{3}^{\prime} & y_{2}(0)=0 \\
y_{3}^{\prime}=c_{3} y_{2}^{2} & y_{3}(0)=0 \\
c_{1}=0.04, c_{2}=10^{4}, c_{3}=3 \times 10^{7} & 0 \leq t \leq 10
\end{array}
$$

Output is obtained after each unit of the independent variable. A user-provided subroutine for the Jacobian matrix is used. An absolute error tolerance of $10^{-5}$ is required.

```
USE IVPAG INT
USE UMACH_INT
IMPLICIT NONE
INTEGER MXPARM, N
PARAMETER (MXPARM=50, N=3)
INTEGER MABSE, MBDF, MSOLVE
PARAMETER (MABSE=1, MBDF=2, MSOLVE=1)
INTEGER IDO, ISTEP, NOUT
REAL A(1,1), PARAM(MXPARM), T, TEND, TOL, Y(N)
    SPECIFICATIONS FOR SUBROUTINES
    SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL FCN, FCNJ
CALL UMACH (2, NOUT)
T = 0.0
Y(1) = 1.0
Y(2) = 0.0
Y(3) = 0.0
TOL = 1.OE-5
    PARAM = 0.0E0
    PARAM(10) = MABSE
        Select absolute error control
        Select BDF method
```

!

```
    PARAM(12) = MBDF
                                Select chord method and
!
! WRITE (NOUT,99998)
IDO = 1
ISTEP = 0
    10 CONTINUE
ISTEP = ISTEP + 1
TEND = ISTEP
The array a(*,*) is not used.
    CALL IVPAG (IDO, FCN, FCNJ, T, TEND, Y, TOL=TOL, PARAM=PARAM)
    IF (ISTEP .LE. 10) THEN
        WRITE (NOUT,'(I6,F12.2,3F13.5)') ISTEP, T, Y
        Final call to release workspace
        IF (ISTEP .EQ. 10) IDO = 3
        GO TO 10
    END IF
99998 FORMAT (4X, 'ISTEP', 5X, 'Time', 9X, 'Y1', 11X, 'Y2', 11X, &
        'Y3')
    END
    SUBROUTINE FCN (N, T, Y, YPRIME)
                                    SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(N), YPRIME (N)
                                    SPECIFICATIONS FOR SAVE VARIABLES
    REAL C1, C2, C3
    SAVE C1, C2, C3
    DATA C1, C2, C3/0.04E0, 1.0E4, 3.0E7/
    YPRIME(1) = -C1*Y(1) + C2*Y(2)*Y(3)
    YPRIME (3) = C3*Y(2)**2
    YPRIME (2) = -YPRIME (1) - YPRIME (3)
    RETURN
    END
    SUBROUTINE FCNJ (N, T, Y, DYPDY)
                        SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(N), DYPDY(N,*)
        SPECIFICATIONS FOR SAVE VARIABLES
    REAL C1, C2, C3
    SAVE C1, C2, C3
    EXTERNAL SSET
    DATA C1, C2, C3/0.04E0, 1.0E4, 3.0E7/
        Clear array to zero
    CALL SSET (N**2, 0.0, DYPDY, 1)
        Compute partials
    DYPDY (1,1) = -C1
    DYPDY(1,2) = C2*Y(3)
    DYPDY(1,3) = C2*Y(2)
    DYPDY (3,2) = 2.0*C3*Y(2)
```

```
DYPDY (2,1) = - DYPDY (1,1)
DYPDY (2,2) = - DYPDY (1,2) - DYPDY (3,2)
DYPDY (2,3) = -DYPDY (1,3)
RETURN
END
```


## Output

| ISTEP | Time | Y1 | Y2 | Y3 |
| :--- | :---: | :---: | :---: | :---: |
| 1 | 1.00 | 0.96647 | 0.00003 | 0.03350 |
| 2 | 2.00 | 0.94164 | 0.00003 | 0.05834 |
| 3 | 3.00 | 0.92191 | 0.00002 | 0.07806 |
| 4 | 4.00 | 0.90555 | 0.00002 | 0.09443 |
| 5 | 5.00 | 0.89153 | 0.00002 | 0.10845 |
| 6 | 6.00 | 0.87928 | 0.00002 | 0.12070 |
| 7 | 7.00 | 0.86838 | 0.00002 | 0.13160 |
| 8 | 8.00 | 0.85855 | 0.00002 | 0.14143 |
| 9 | 9.00 | 0.84959 | 0.00002 | 0.15039 |
| 10 | 10.00 | 0.84136 | 0.00002 | 0.15862 |

## Example 4

Solve the partial differential equation

$$
e^{-t} \frac{\partial u}{\partial t}=\frac{\partial^{2} u}{\partial x^{2}}
$$

with the initial condition

$$
u(t=0, x)=\sin x
$$

and the boundary conditions

$$
u(t, x=0)=u(t, x=\pi)=0
$$

on the square $[0,1] \times[0, \pi]$, using the method of lines with a piecewise-linear Galerkin discretization. The exact solution is $u(t, x)=\exp \left(1-e^{t}\right) \sin x$. The interval $[0, \pi]$ is divided into equal intervals by choosing breakpoints $x_{k}=k \pi /(N+1)$ for $k=0, \ldots, N+1$. The unknown function $u(t, x)$ is approximated by

$$
\sum_{k=1}^{N} c_{k}(t) \phi_{k}(x)
$$

where $\phi_{k}(x)$ is the piecewiselinear function that equals 1 at $x_{k}$ and is zero at all of the other breakpoints. We approximate the partial differential equation by a system of $N$ ordinary differential equations, $A d c / d t=R c$ where $A$ and $R$ are matrices of order $N$. The matrix $A$ is given by

$$
\begin{aligned}
& e^{-t} 2 h / 3 \quad \text { if } i=j \\
& A_{i j}=e^{-t} \int_{0}^{\pi} \phi_{i}(x) \phi_{j}(x) d x=e^{-t} h / 6 \quad \text { if } i=j \pm 1 \\
& 0 \quad \text { otherwise }
\end{aligned}
$$

where $h=1 /(N+1)$ is the mesh spacing. The matrix $R$ is given by

$$
R_{i j}=\int_{0}^{\pi} \phi_{i}^{\prime \prime}(x) \phi_{j}(x) d x=-\int_{0}^{\pi} \phi_{i}^{\prime}(x) \phi_{\mathrm{j}}^{\prime}(x) d x=1 / h \quad \text { if } i=j, ~ \begin{array}{ll}
-2 / h & \text { if } i=j \pm 1
\end{array}
$$

0 otherwise
The integrals involving

$$
\phi_{i}^{\prime \prime}
$$

are assigned the values of the integrals on the right-hand side, by using the boundary values and integration by parts. Because this system may be stiff, Gear's BDF method is used.

In the following program, the array $\mathrm{Y}(1: \mathrm{N})$ corresponds to the vector of coefficients, $c$. Note that Y contains $\mathrm{N}+2$ elements; $\mathrm{Y}(0)$ and $\mathrm{Y}(\mathrm{N}+1)$ are used to store the boundary values. The matrix $A$ depends on $t$ so we set $\operatorname{PARAM}(19)=2$ and evaluate $A$ when IVPAG returns with IDO $=7$. The subroutine FCN computes the vector $R c$, and the subroutine FCNJ computes $R$. The matrices $A$ and $R$ are stored as band-symmetric positive-definite structures having one upper co-diagonal.

```
USE IVPAG INT
USE CONST - INT
USE WRRRN INT
USE SSET_İNT
IMPLICIT NONE
INTEGER LDA, N, NPARAM, NUC
PARAMETER (N=9, NPARAM=50, NUC=1, LDA=NUC+1)
                                    SPECIFICATIONS FOR PARAMETERS
INTEGER NSTEP
PARAMETER (NSTEP=4)
INTEGER I, IATYPE, IDO, IMETH, INORM, ISTEP, MITER, MTYPE
REAL A(LDA,N), C, HINIT, PARAM(NPARAM), PI, T, TEND, TMAX, &
    TOL, XPOINT(0:N+1), Y(0:N+1)
CHARACTER TITLE*10
COMMON /COMHX/ HX
REAL HX
INTRINSIC EXP, REAL, SIN
REAL EXP, REAL, SIN
EXTERNAL FCN, FCNJ
HINIT = 1.0E-3
INORM = 1
IMETH = 2
MITER = 1
MTYPE = 3
IATYPE = 2
PARAM = 0.0E0
```

```
    PARAM(1) = HINIT
    PARAM(10) = INORM921
    PARAM(12) = IMETH
    PARAM(13) = MITER
    PARAM(14) = MTYPE
    PARAM(16) = NUC
    PARAM(19) = IATYPE
!
                                    Initialize other arguments
    PI = CONST('PI')
    HX = PI/REAL (N+1)
    CALL SSET (N-1, HX/6., A(1:,2), LDA)
    CALL SSET (N, 2.*HX/3., A(2:,1), LDA)
    DO 10 I=0, N + 1
        XPOINT(I) = I*HX
        Y(I) = SIN(XPOINT(I))
    1 0 ~ C O N T I N U E ~
    TOL = 1.0E-6
    T = 0.0
    TMAX = 1.0
! Integrate ODE
    IDO = 1
    ISTEP = 0
    20 CONTINUE
    ISTEP = ISTEP + 1
    TEND = TMAX*REAL(ISTEP)/REAL (NSTEP)
    3 0 ~ C A L L ~ I V P A G ~ ( I D O , ~ F C N , ~ F C N J , ~ T , ~ T E N D , ~ Y ( 1 : ) , ~ N E Q = N , ~ A = A , ~ \& ~
        TOL=TOL, PARAM=PARAM)
                                    Set matrix A
    IF (IDO .EQ. 7) THEN
        C = EXP(-T)
        CALL SSET (N-1, C*HX/6., A(1:,2), LDA)
        CALL SSET (N, 2.*C*HX/3., A(2:,1), LDA)
        GO TO 30
    END IF
    IF (ISTEP .LE. NSTEP) THEN
        WRITE (TITLE,'(A,F5.3,A)') 'U(T=', T, ')'
        CALL WRRRN (TITLE, Y, 1, N+2, 1)
                                    Final call to release workspace
        IF (ISTEP .EQ. NSTEP) IDO = 3
        GO TO 20
    END IF
    END
    SUBROUTINE FCN (N, T, Y, YPRIME)
                        SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(*), YPRIME(N)
        SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER I
        SPECIFICATIONS FOR COMMON /COMHX/
    COMMON /COMHX/ HX
    REAL HX
    SPECIFICATIONS FOR SUBROUTINES
```

```
    EXTERNAL SSCAL
!
    YPRIME (1) = -2.0*Y(1) + Y(2)
    DO 10 I=2, N - 1
        YPRIME (I) = -2.0*Y(I) + Y(I-1) + Y(I+1)
    1 0 \text { CONTINUE}
    YPRIME (N) = -2.0*Y(N) + Y(N-1)
    CALL SSCAL (N, 1.0/HX, YPRIME, 1)
    RETURN
    END
!
    SUBROUTINE FCNJ (N, T, Y, DYPDY)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(*), DYPDY (2,*)
    SPECIFICATIONS FOR COMMON /COMHX/
    COMMON /COMHX/ HX
    REAL HX
    EXTERNAL SSET
!
    CALL SSET (N-1, 1.0/HX, DYPDY(1,2), 2)
    CALL SSET (N, -2.0/HX, DYPDY(2,1), 2)
    RETURN
    END
```


## Output

| $\mathrm{U}(\mathrm{T}=0.250)$ |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 0.0000 | 0.2321 | 0.4414 | 0.6076 | 0.7142 | 0.7510 | 0.7142 | 0.6076 |
| 9 | 10 | 11 |  |  |  |  |  |
| 0.4414 | 0.2321 | 0.0000 |  |  |  |  |  |
| $\mathrm{U}(\mathrm{T}=0.500)$ |  |  |  |  |  |  |  |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 0.0000 | 0.1607 | 0.3056 | 0.4206 | 0.4945 | 0.5199 | 0.4945 | 0.4206 |
| 9 | 10 | 11 |  |  |  |  |  |
| 0.3056 | 0.1607 | 0.0000 |  |  |  |  |  |
| $\mathrm{U}(\mathrm{T}=0.750)$ |  |  |  |  |  |  |  |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 0.0000 | 0.1002 | 0.1906 | 0.2623 | 0.3084 | 0.3243 | 0.3084 | 0.2623 |
| 9 | 10 | 11 |  |  |  |  |  |
| 0.1906 | 0.1002 | 0.0000 |  |  |  |  |  |


| $\mathrm{U}(\mathrm{T}=1.000)$ |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 0.0000 | 0.0546 | 0.1039 | 0.1431 | 0.1682 | 0.1768 | 0.1682 | 0.1431 |
| 9 | 10 | 11 |  |  |  |  |  |
| 0.1039 | 0.0546 | 0.0000 |  |  |  |  |  |

## BVPFD

Solves a (parameterized) system of differential equations with boundary conditions at two points, using a variable order, variable step size finite difference method with deferred corrections.

## Required Arguments

FCNEQN - User-supplied SUBROUTINE to evaluate derivatives. The usage is CALL
FCNEQN (N, T, Y, P, DYDT), where
N - Number of differential equations. (Input)
T - Independent variable, $t$. (Input)
Y - Array of size N containing the dependent variable values, $y(\mathrm{t})$.
(Input)
P - Continuation parameter, $p$. (Input)
See Comment 3.
DYDT - Array of size N containing the derivatives $y^{\prime}(t)$. (Output)
The name FCNEQN must be declared EXTERNAL in the calling program.
FCNJAC - User-supplied SUBROUTINE to evaluate the Jacobian. The usage is CALL
FCNJAC ( $\mathrm{N}, \mathrm{T}, \mathrm{Y}, \mathrm{P}$, DYPDY), where
N - Number of differential equations. (Input)
T - Independent variable, $t$. (Input)
Y - Array of size N containing the dependent variable values. (Input)
P - Continuation parameter, $p$. (Input)
See Comments 3 .
DYPDY -N by N array containing the partial derivatives $a_{i, j}=\partial f_{i} / \partial y_{j}$
evaluated at $(t, y)$. The values $a_{i, j}$ are returned in $\operatorname{DYPDY}(i, j)$.
(Output)
The name FCNJAC must be declared EXTERNAL in the calling program.

FCNBC - User-supplied SUBROUTINE to evaluate the boundary conditions. The usage is
CALL FCNBC (N, YLEFT, YRIGHT, P, H), where
N - Number of differential equations. (Input)
YLEFT - Array of size N containing the values of the dependent
variable at the left endpoint. (Input)
YRIGHT - Array of size N containing the values of the dependent
variable at the right endpoint. (Input)
P - Continuation parameter, $p$. (Input)
See Comment 3.

H - Array of size N containing the boundary condition residuals.
(Output)

The boundary conditions are defined by $h_{i}=0$; for $i=1, \ldots, N$. The left endpoint conditions must be defined first, then, the conditions involving both endpoints, and finally the right endpoint conditions.

The name FCNBC must be declared EXTERNAL in the calling program.
FCNPEQ - User-supplied SUBROUTINE to evaluate the derivative of $y^{\prime}$ with respect to the parameter $p$. The usage is

CALL FCNPEQ ( $\mathrm{N}, \mathrm{T}, \mathrm{Y}, \mathrm{P}, \mathrm{DYPDP}$ ), where
N - Number of differential equations. (Input)
$T$ - Dependent variable, $t$. (Input)
Y - Array of size N containing the dependent variable values. (Input)
P - Continuation parameter, $p$. (Input)
See Comment 3.
DYPDP - Array of size N containing the derivative of $y^{\prime}$
evaluated at $(t, y)$.
(Output)
The name FCNPEQ must be declared EXTERNAL in the calling program.
FCNPBC - User-supplied SUBROUTINE to evaluate the derivative of the boundary conditions with respect to the parameter $p$. The usage is
CALL FCNPBC (N, YLEFT, YRIGHT, $\mathrm{P}, \mathrm{H}$ ), where
N - Number of differential equations. (Input)
YLEFT - Array of size n containing the values of the dependent
variable at the left endpoint. (Input)
YRIGHT - Array of size N containing the values of the dependent
variable at the right endpoint. (Input)
P - Continuation parameter, $p$. (Input)
See Comment 3.
$\mathrm{H}-$ Array of size N containing the derivative of $f_{i}$ with respect to $p$. (Output)

The name FCNPBC must be declared EXTERNAL in the calling program.
NLEFT - Number of initial conditions. (Input)
The value NLEFT must be greater than or equal to zero and less than $N$.
NCUPBC - Number of coupled boundary conditions. (Input)
The value NLEFT + NCUPBC must be greater than zero and less than or equal to N .
TLEFT - The left endpoint. (Input)
TRIGHT - The right endpoint. (Input)

PISTEP - Initial increment size for $p$. (Input)
If this value is zero, continuation will not be used in this problem. The routines FCNPEQ and FCNPBC will not be called.

TOL - Relative error control parameter. (Input)
The computations stop when $\operatorname{ABS}(\operatorname{ERROR}(\mathrm{J}, \mathrm{I})) / \operatorname{MAX}(\operatorname{ABS}(\mathrm{Y}(\mathrm{J}, \mathrm{I})), 1.0) . \operatorname{LT} . \mathrm{TOL}$ for all $J=1, \ldots, N$ and $I=1, \ldots, \operatorname{NGRID}$. Here, $\operatorname{ERROR}(J, I)$ is the estimated error in $Y(J, I)$.

TINIT - Array of size NINIT containing the initial grid points. (Input)
YINIT - Array of size N by NINIT containing an initial guess for the values of Y at the points in TINIT. (Input)

LINEAR - Logical . TRUE. if the differential equations and the boundary conditions are linear. (Input)

MXGRID - Maximum number of grid points allowed. (Input)
NFINAL — Number of final grid points, including the endpoints. (Output)
TFINAL - Array of size MXGRID containing the final grid points. (Output) Only the first NFINAL points are significant.

YFINAL - Array of size N by MXGRID containing the values of Y at the points in TFINAL. (Output)

ERREST - Array of size N. (Output)
$\operatorname{ERREST}(\mathrm{J})$ is the estimated error in $Y(\mathrm{~J})$.

## Optional Arguments

$N$ - Number of differential equations. (Input)
Default: $\mathrm{N}=\operatorname{size}$ (YINIT,1).
NINIT - Number of initial grid points, including the endpoints. (Input)
It must be at least 4 .
Default: NINIT $=\operatorname{size}(T I N I T, 1)$.
LDYINI - Leading dimension of YINIT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDYINI $=\operatorname{size}($ YINIT, 1$)$.
PRINT - Logical . TRUE. if intermediate output is to be printed. (Input) Default: PRINT = .FALSE.

LDYFIN - Leading dimension of YFINAL exactly as specified in the dimension statement of the calling program. (Input)
Default: LDYFIN = size (YFINAL,1).

## FORTRAN 90 Interface

```
Generic: CALL BVPFD (FCNEQN, FCNJAC, FCNBC, FCNPEQ, FCNPBC, NLEFT, NCUPBC, TLEFT, TRIGHT, PISTEP, TOL, TINIT, YINIT, LINEAR, MXGRID, NFINAL, TFINAL, YFINAL, ERREST [,...])
```

Specific: The specific interface names are S_BVPFD and D_BVPFD.

## FORTRAN 77 Interface

Single:

```
CALL BVPFD (FCNEQN, FCNJAC, FCNBC, FCNPEQ, FCNPBC, N,
NLEFT, NCUPBC, TLEFT, TRIGHT, PISTEP, TOL, NINIT, TINIT,
YINIT, LDYINI, LINEAR, PRINT, MXGRID, NFINAL, TFINAL,
YFINAL, LDYFIN, ERREST)
```

Double: The double precision name is DBVPFD.

## Description

The routine BVPFD is based on the subprogram PASVA3 by M. Lentini and V. Pereyra (see Pereyra 1978). The basic discretization is the trapezoidal rule over a nonuniform mesh. This mesh is chosen adaptively, to make the local error approximately the same size everywhere. Higher-order discretizations are obtained by deferred corrections. Global error estimates are produced to control the computation. The resulting nonlinear algebraic system is solved by Newton's method with step control. The linearized system of equations is solved by a special form of Gauss elimination that preserves the sparseness.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2PFD/DB2PFD. The reference is:
```
CALL B2PFD (FCNEQN, FCNJAC, FCNBC, FCNPEQ, FCNPBC, N, NLEFT, NCUPBC,
TLEFT, TRIGHT, PISTEP, TOL, NINIT, TINIT, YINIT, LDYINI, LINEAR, PRINT,
MXGRID, NFINAL, TFINAL, YFINAL, LDYFIN, ERREST, RWORK, IWORK)
```

The additional arguments are as follows:
$\boldsymbol{R W O R K}$ - Floating-point work array of size N (3N * MXGRID + 4N + 1) + MXGRID * $(7 \mathrm{~N}+2)$.

IWORK - Integer work array of size 2 N * MXGRID + N + MXGRID.
2. Informational errors

Type Code
41 More than MXGRID grid points are needed to solve the problem.
42 Newton's method diverged.
33 Newton's method reached roundoff error level.
3. If the value of PISTEP is greater than zero, then the routine BVPFD assumes that the user has embedded the problem into a one-parameter family of problems:

$$
\begin{gathered}
y^{\prime}=y^{\prime}(t, y, p) \\
h\left(y_{\text {tleft }}, y_{\text {tright }}, p\right)=0
\end{gathered}
$$

such that for $p=0$ the problem is simple. For $p=1$, the original problem is recovered. The routine BVPFD automatically attempts to increment from $p=0$ to $p=1$. The value PISTEP is the beginning increment used in this continuation. The increment will usually be changed by routine BVPFD, but an arbitrary minimum of 0.01 is imposed.
4. The vectors TINIT and TFINAL may be the same.
5. The arrays Yinit and YFINAL may be the same.

## Example 1

This example solves the third-order linear equation

$$
y^{\prime \prime \prime}-2 y^{\prime \prime}+y^{\prime}-y=\sin t
$$

subject to the boundary conditions $y(0)=y(2 \pi)$ and $y^{\prime}(0)=y^{\prime}(2 \pi)=1$. (Its solution is $y=\sin t$.) To use BVPFD, the problem is reduced to a system of first-order equations by defining $y_{1}=y, y_{2}=y^{\prime}$ and $y_{3}=y^{\prime \prime}$. The resulting system is

$$
\begin{array}{ll}
y_{1}^{\prime}=y_{2} & y_{2}(0)-1=0 \\
y_{2}^{\prime}=y_{3} & y_{1}(0)-y_{1}(2 \pi)=0 \\
y_{3}^{\prime}=2 y_{3}-y_{2}+y_{1}+\sin t & y_{2}(2 \pi)-1=0
\end{array}
$$

Note that there is one boundary condition at the left endpoint $t=0$ and one boundary condition coupling the left and right endpoints. The final boundary condition is at the right endpoint. The total number of boundary conditions must be the same as the number of equations (in this case 3 ).

Note that since the parameter $p$ is not used in the call to BVPFD, the routines FCNPEQ and FCNPBC are not needed. Therefore, in the call to BVPFD, FCNEQN and FCNBC were used in place of FCNPEQ and FCNPBC.

```
USE BVPFD INT
USE UMACH_INT
USE CONST_INT
IMPLICIT NONE
! SPECIFICATIONS FOR PARAMETERS
INTEGER LDYFIN, LDYINI, MXGRID, NEQNS, NINIT
PARAMETER (MXGRID=45, NEQNS=3, NINIT=10, LDYFIN=NEQNS, &
            LDYINI=NEQNS)
                            SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, J, NCUPBC, NFINAL, NLEFT, NOUT
REAL ERREST(NEQNS), PISTEP, TFINAL(MXGRID), TINIT(NINIT), &
```

```
        TLEFT, TOL, TRIGHT, YFINAL(LDYFIN,MXGRID), &
        YINIT(LDYINI,NINIT)
    LOGICAL LINEAR, PRINT
! SPECIFICATIONS FOR INTRINSICS
    INTRINSIC FLOAT
    REAL FLOAT
! SPECIFICATIONS FOR SUBROUTINES
                                SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL FCNBC, FCNEQN, FCNJAC
    NLEFT = 1
    NCUPBC = 1
    TOL = .001
    TLEFT = 0.0
    TRIGHT = CONST('PI')
    TRIGHT = 2.0*TRIGHT
    PISTEP = 0.0
    PRINT = .FALSE.
    LINEAR = .TRUE.
! Define TINIT
    DO 10 I=1, NINIT
    TINIT(I) = TLEFT + (I-1)*(TRIGHT-TLEFT)/FLOAT(NINIT-1)
    10 CONTINUE
    YINIT = 0.0EO
    CALL BVPFD (FCNEQN, FCNJAC, FCNBC, FCNEQN, FCNBC, NLEFT, &
        NCUPBC, TLEFT, TRIGHT, PISTEP, TOL, TINIT, &
        YINIT, LINEAR, MXGRID, NFINAL, &
        TFINAL, YFINAL, ERREST)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99997)
    WRITE (NOUT,99998) (I,TFINAL(I),(YFINAL(J,I),J=1,NEQNS),I=1, &
        NFINAL)
    WRITE (NOUT,99999) (ERREST(J),J=1,NEQNS)
99997 FORMAT (4X, 'I', 7X, 'T', 14X, 'Y1', 13X, 'Y2', 13X, 'Y3')
99998 FORMAT (I5, 1P4E15.6)
99999 FORMAT (' Error estimates', 4X, 1P3E15.6)
    END
    SUBROUTINE FCNEQN (NEQNS, T, Y, P, DYDX)
        SPECIFICATIONS FOR ARGUMENTS
    INTEGER NEQNS
    REAL T, P, Y(NEQNS), DYDX(NEQNS)
    INTRINSIC SIN
    REAL SIN
! Define PDE
    DYDX(1) = Y(2)
    DYDX(2) = Y(3)
    DYDX(3) = 2.0*Y(3) - Y(2) + Y(1) + SIN(T)
    RETURN
    END
    SUBROUTINE FCNJAC (NEQNS, T, Y, P, DYPDY)
        SPECIFICATIONS FOR ARGUMENTS
```

```
        INTEGER NEQNS
        REAL T, P, Y(NEQNS), DYPDY(NEQNS,NEQNS)
                                Define d(DYDX)/dY
        DYPDY (1,1) = 0.0
        DYPDY (1,2) = 1.0
        DYPDY (1,3) = 0.0
        DYPDY (2,1) = 0.0
        DYPDY (2,2) = 0.0
        DYPDY (2,3) = 1.0
        DYPDY (3,1) = 1.0
        DYPDY (3,2) = -1.0
        DYPDY (3,3) = 2.0
        RETURN
        END
        SUBROUTINE FCNBC (NEQNS, YLEFT, YRIGHT, P, F)
    SPECIFICATIONS FOR ARGUMENTS
        INTEGER NEQNS
        REAL P, YLEFT (NEQNS), YRIGHT (NEQNS), F(NEQNS)
    Define boundary conditions
        F(1) = YLEFT(2) - 1.0
        F(2) = YLEFT(1) - YRIGHT(1)
        F(3) = YRIGHT (2) - 1.0
        RETURN
        END
```


## Output

| $I$ | T | Y 1 | Y 2 | Y 3 |
| :--- | :---: | ---: | :---: | ---: |
| 1 | $0.000000 \mathrm{E}+00$ | $-1.123191 \mathrm{E}-04$ | $1.000000 \mathrm{E}+00$ | $6.242319 \mathrm{E}-05$ |
| 2 | $3.490659 \mathrm{E}-01$ | $3.419107 \mathrm{E}-01$ | $9.397087 \mathrm{E}-01$ | $-3.419580 \mathrm{E}-01$ |
| 3 | $6.981317 \mathrm{E}-01$ | $6.426908 \mathrm{E}-01$ | $7.660918 \mathrm{E}-01$ | $-6.427230 \mathrm{E}-01$ |
| 4 | $1.396263 \mathrm{E}+00$ | $9.847531 \mathrm{E}-01$ | $1.737333 \mathrm{E}-01$ | $-9.847453 \mathrm{E}-01$ |
| 5 | $2.094395 \mathrm{E}+00$ | $8.660529 \mathrm{E}-01$ | $-4.998747 \mathrm{E}-01$ | $-8.660057 \mathrm{E}-01$ |
| 6 | $2.792527 \mathrm{E}+00$ | $3.421830 \mathrm{E}-01$ | $-9.395474 \mathrm{E}-01$ | $-3.420648 \mathrm{E}-01$ |
| 7 | $3.490659 \mathrm{E}+00$ | $-3.417234 \mathrm{E}-01$ | $-9.396111 \mathrm{E}-01$ | $3.418948 \mathrm{E}-01$ |
| 8 | $4.188790 \mathrm{E}+00$ | $-8.656880 \mathrm{E}-01$ | $-5.000588 \mathrm{E}-01$ | $8.658733 \mathrm{E}-01$ |
| 9 | $4.886922 \mathrm{E}+00$ | $-9.845794 \mathrm{E}-01$ | $1.734571 \mathrm{E}-01$ | $9.847518 \mathrm{E}-01$ |
| 10 | $5.585054 \mathrm{E}+00$ | $-6.427721 \mathrm{E}-01$ | $7.658258 \mathrm{E}-01$ | $6.429526 \mathrm{E}-01$ |
| 11 | $5.934120 \mathrm{E}+00$ | $-3.420819 \mathrm{E}-01$ | $9.395434 \mathrm{E}-01$ | $3.423986 \mathrm{E}-01$ |
| 12 | $6.283185 \mathrm{E}+00$ | $-1.123186 \mathrm{E}-04$ | $1.000000 \mathrm{E}+00$ | $6.743190 \mathrm{E}-04$ |
| Error | estimates | $2.840430 \mathrm{E}-04$ | $1.792939 \mathrm{E}-04$ | $5.588399 \mathrm{E}-04$ |

## Additional Examples

## Example 2

In this example, the following nonlinear problem is solved:

$$
y^{\prime \prime}-y^{3}+\left(1+\sin ^{2} t\right) \sin t=0
$$

with $y(0)=y(\pi)=0$. Its solution is $y=\sin t$. As in Example 1, this equation is reduced to a system of first-order differential equations by defining $y_{1}=y$ and $y_{2}=y^{\prime}$. The resulting system is

$$
\begin{array}{ll}
y_{1}^{\prime}=y_{2} & y_{1}(0)=0 \\
y_{2}^{\prime}=y_{1}^{3}-\left(1+\sin ^{2} t\right) \sin t \quad y_{1}(\pi)=0
\end{array}
$$

In this problem, there is one boundary condition at the left endpoint and one at the right endpoint; there are no coupled boundary conditions.

Note that since the parameter $p$ is not used, in the call to BVPFD the routines FCNPEQ and FCNPBC are not needed. Therefore, in the call to BVPFD, FCNEQN and FCNBC were used in place of FCNPEQ and $\operatorname{FCNPBC}$.

```
USE BVPFD INT
USE UMACH INT
USE CONST_INT
IMPLICIT NONE
! SPECIFICATIONS FOR PARAMETERS
INTEGER LDYFIN, LDYINI, MXGRID, NEQNS, NINIT
PARAMETER (MXGRID=45, NEQNS=2, NINIT=12, LDYFIN=NEQNS, &
            LDYINI=NEQNS)
                            SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, J, NCUPBC, NFINAL, NLEFT, NOUT
REAL ERREST (NEQNS), PISTEP, TFINAL(MXGRID), TINIT(NINIT), &
            TLEFT, TOL, TRIGHT, YFINAL(LDYFIN,MXGRID) , &
            YINIT(LDYINI,NINIT)
LOGICAL LINEAR, PRINT
SPECIFICATIONS FOR INTRINSICS
INTRINSIC FLOAT
REAL FLOAT
EXTERNAL FCNBC, FCNEQN, FCNJAC
NLEFT = 1
NCUPBC = 0
TOL = .001
TLEFT = 0.0
TRIGHT = CONST('PI')
PISTEP = 0.0
PRINT = .FALSE.
LINEAR = .FALSE.
    DO 10 I=1, NINIT
        TINIT(I) = TLEFT + (I-1)*(TRIGHT-TLEFT)/FLOAT (NINIT-1)
        YINIT(1,I) = 0.4*(TINIT(I) -TLEFT)*(TRIGHT-TINIT(I))
        YINIT(2,I) = 0.4*(TLEFT-TINIT(I) +TRIGHT-TINIT(I))
    CONTINUE
CALL BVPFD (FCNEQN, FCNJAC, FCNBC, FCNEQN, FCNBC, NLEFT, &
        NCUPBC, TLEFT, TRIGHT, PISTEP, TOL, TINIT, &
        YINIT, LINEAR, MXGRID, NFINAL, &
        TFINAL, YFINAL, ERREST)
            Print results
CALL UMACH (2, NOUT)
```

!

```
        WRITE (NOUT,99997)
        WRITE (NOUT,99998) (I,TFINAL(I), (YFINAL(J,I),J=1,NEQNS),I=1, &
                                    NFINAL)
    WRITE (NOUT,99999) (ERREST (J),J=1,NEQNS)
99997 FORMAT (4X, 'I', 7X, 'T', 14X, 'Y1', 13X, 'Y2')
99998 FORMAT (I5, 1P3E15.6)
99999 FORMAT (' Error estimates', 4X, 1P2E15.6)
        END
        SUBROUTINE FCNEQN (NEQNS, T, Y, P, DYDT)
! SPECIFICATIONS FOR ARGUMENTS
    INTEGER NEQNS
        REAL T, P, Y(NEQNS), DYDT (NEQNS)
                                SPECIFICATIONS FOR INTRINSICS
        INTRINSIC SIN
        REAL SIN
        DYDT(1) = Y(2)
        DYDT(2) = Y(1)**3 - SIN(T)*(1.0+SIN(T)**2)
        RETURN
        END
        SUBROUTINE FCNJAC (NEQNS, T, Y, P, DYPDY)
        SPECIFICATIONS FOR ARGUMENTS
        INTEGER NEQNS
        REAL T, P, Y(NEQNS), DYPDY(NEQNS,NEQNS)
                                Define d(DYDT)/dY
        DYPDY (1,1) = 0.0
        DYPDY (1,2) = 1.0
        DYPDY (2,1) = 3.0*Y(1)**2
        DYPDY (2,2) = 0.0
        RETURN
        END
        SUBROUTINE FCNBC (NEQNS, YLEFT, YRIGHT, P, F)
        SPECIFICATIONS FOR ARGUMENTS
        INTEGER NEQNS
        REAL P, YLEFT (NEQNS), YRIGHT (NEQNS), F(NEQNS)
        F(1) = YLEFT(1)
        F(2) = YRIGHT (1)
        RETURN
        END
```


## Output

```
\begin{tabular}{rccc}
\(I\) & \(T\) & Y 1 & Y 2 \\
1 & \(0.000000 \mathrm{E}+00\) & \(0.000000 \mathrm{E}+00\) & \(9.999277 \mathrm{E}-01\) \\
2 & \(2.855994 \mathrm{E}-01\) & \(2.817682 \mathrm{E}-01\) & \(9.594315 \mathrm{E}-01\) \\
3 & \(5.711987 \mathrm{E}-01\) & \(5.406458 \mathrm{E}-01\) & \(8.412407 \mathrm{E}-01\) \\
4 & \(8.567980 \mathrm{E}-01\) & \(7.557380 \mathrm{E}-01\) & \(6.548904 \mathrm{E}-01\) \\
5 & \(1.142397 \mathrm{E}+00\) & \(9.096186 \mathrm{E}-01\) & \(4.154530 \mathrm{E}-01\) \\
6 & \(1.427997 \mathrm{E}+00\) & \(9.898143 \mathrm{E}-01\) & \(1.423307 \mathrm{E}-01\) \\
7 & \(1.713596 \mathrm{E}+00\) & \(9.898143 \mathrm{E}-01\) & \(-1.423307 \mathrm{E}-01\) \\
8 & \(1.999195 \mathrm{E}+00\) & \(9.096185 \mathrm{E}-01\) & \(-4.154530 \mathrm{E}-01\) \\
9 & \(2.284795 \mathrm{E}+00\) & \(7.557380 \mathrm{E}-01\) & \(-6.548903 \mathrm{E}-01\) \\
10 & \(2.570394 \mathrm{E}+00\) & \(5.406460 \mathrm{E}-01\) & \(-8.412405 \mathrm{E}-01\) \\
11 & \(2.855994 \mathrm{E}+00\) & \(2.817683 \mathrm{E}-01\) & \(-9.594313 \mathrm{E}-01\)
\end{tabular}
```

```
12 3.141593E+00 0.000000E+00 -9.999274E-01
Error estimates 3.906105E-05 7.124186E-05
```


## Example 3

In this example, the following nonlinear problem is solved:

$$
y^{\prime \prime}-y^{3}=\frac{40}{9}\left(t-\frac{1}{2}\right)^{2 / 3}-\left(t-\frac{1}{2}\right)^{8}
$$

with $y(0)=y(1)=\pi / 2$. As in the previous examples, this equation is reduced to a system of firstorder differential equations by defining $y_{1}=y$ and $y_{2}=y^{\prime}$. The resulting system is

$$
\begin{aligned}
& y_{1}^{\prime}=y_{2} \\
& y_{2}^{\prime}=y_{1}^{3}-\frac{40}{9}\left(t-\frac{1}{2}\right)^{2 / 3}+\left(t-\frac{1}{2}\right)^{8} \quad y_{1}(1)=\pi / 2
\end{aligned}
$$

The problem is embedded in a family of problems by introducing the parameter $p$ and by changing the second differential equation to

$$
y_{2}^{\prime}=p y_{1}^{3}+\frac{40}{9}\left(t-\frac{1}{2}\right)^{2 / 3}-\left(t-\frac{1}{2}\right)^{8}
$$

At $p=0$, the problem is linear; and at $p=1$, the original problem is recovered. The derivatives $\partial y^{\prime} / \partial p$ must now be specified in the subroutine FCNPEQ. The derivatives $\partial f / \partial p$ are zero in FCNPBC.

```
    USE BVPFD_INT
    USE UMACH_INT
    IMPLICIT NONE
! SPECIFICATIONS FOR PARAMETERS
    INTEGER LDYFIN, LDYINI, MXGRID, NEQNS, NINIT
    PARAMETER (MXGRID=45, NEQNS=2, NINIT=5, LDYFIN=NEQNS, &
        LDYINI=NEQNS)
    INTEGER NCUPBC, NFINAL, NLEFT, NOUT
    REAL ERREST(NEQNS), PISTEP, TFINAL(MXGRID), TLEFT, TOL, &
        XRIGHT, YFINAL(LDYFIN,MXGRID)
    LOGICAL LINEAR, PRINT
    INTEGER I, J
    REAL TINIT(NINIT), YINIT(LDYINI,NINIT)
    SAVE I, J, TINIT, YINIT
    SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL FCNBC, FCNEQN, FCNJAC, FCNPBC, FCNPEQ
!
    DATA TINIT/0.0, 0.4, 0.5, 0.6, 1.0/
    DATA ((YINIT (I,J),J=1,NINIT),I=1,NEQNS)/0.15749, 0.00215, 0.0, &
        0.00215, 0.15749, -0.83995, -0.05745, 0.0, 0.05745, 0.83995/
        Set parameters
    NLEFT = 1
```

```
NCUPBC = 0
TOL = .001
TLEFT = 0.0
XRIGHT = 1.0
PISTEP = 0.1
PRINT = .FALSE.
LINEAR = .FALSE.
!
!
CALL UMACH (2, NOUT)
WRITE (NOUT,99997)
WRITE (NOUT,99998) (I,TFINAL (I) , (YFINAL (J,I),J=1,NEQNS),I=1, &
                                    NFINAL)
WRITE (NOUT, 99999) (ERREST (J),J=1,NEQNS)
99997 FORMAT (4X, 'I', 7X, 'T', 14X, 'Y1', 13X, 'Y2')
99998 FORMAT (I5, 1P3E15.6)
99999 FORMAT (' Error estimates', 4X, 1P2E15.6)
END
SUBROUTINE FCNEQN (NEQNS, T, Y, P, DYDT)
SPECIFICATIONS FOR ARGUMENTS
INTEGER NEQNS
REAL T, P, Y(NEQNS), DYDT (NEQNS)
DYDT(1) = Y(2)
DYDT (2) = P*Y(1)**3 + 40./9.* ((T-0.5)**2)** (1./3.) - (T-0.5)**8
RETURN
END
SUBROUTINE FCNJAC (NEQNS, T, Y, P, DYPDY)
SPECIFICATIONS FOR ARGUMENTS
INTEGER NEQNS
REAL T, P, Y(NEQNS), DYPDY(NEQNS,NEQNS)
                                    Define d(DYDT)/dY
DYPDY (1,1) = 0.0
DYPDY (1,2) = 1.0
DYPDY (2,1) = P*3.*Y(1)**2
DYPDY (2,2) = 0.0
RETURN
END
SUBROUTINE FCNBC (NEQNS, YLEFT, YRIGHT, P, F)
USE CONST_INT
INTEGER NEQNS
REAL P, YLEFT(NEQNS), YRIGHT(NEQNS), F(NEQNS)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
REAL PI
PI = CONST('PI')
F(1) = YLEFT(1) - PI/2.0
F(2) = YRIGHT(1) - PI/2.0
RETURN
END
SUBROUTINE FCNPEQ (NEQNS, T, Y, P, DYPDP)
```

```
!
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NEQNS
    REAL T, P, Y(NEQNS), DYPDP(NEQNS)
    Define d(DYDT)/dP
    DYPDP(1) = 0.0
    DYPDP(2) = Y(1)**3
    RETURN
    END
    SUBROUTINE FCNPBC (NEQNS, YLEFT, YRIGHT, P, DFDP)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NEQNS
    REAL P, YLEFT (NEQNS), YRIGHT (NEQNS), DFDP (NEQNS)
    SPECIFICATIONS FOR SUBROUTINES
    EXTERNAL SSET
    CALL SSET (NEQNS, 0.0, DFDP, 1)
        RETURN
        END
```


## Output

| I | T | Y1 | Y2 |
| :---: | :---: | :---: | :---: |
| 1 | $0.000000 \mathrm{E}+00$ | $1.570796 \mathrm{E}+00$ | -1.949336E+00 |
| 2 | $4.444445 \mathrm{E}-02$ | $1.490495 \mathrm{E}+00$ | -1.669567E+00 |
| 3 | $8.888889 \mathrm{E}-02$ | $1.421951 \mathrm{E}+00$ | $-1.419465 \mathrm{E}+00$ |
| 4 | 1.333333E-01 | 1.363953E+00 | -1.194307E+00 |
| 5 | $2.000000 \mathrm{E}-01$ | $1.294526 \mathrm{E}+00$ | -8.958461E-01 |
| 6 | $2.666667 \mathrm{E}-01$ | $1.243628 \mathrm{E}+00$ | -6.373191E-01 |
| 7 | $3.333334 \mathrm{E}-01$ | $1.208785 \mathrm{E}+00$ | -4.135206E-01 |
| 8 | $4.000000 \mathrm{E}-01$ | $1.187783 \mathrm{E}+00$ | -2.219351E-01 |
| 9 | $4.250000 \mathrm{E}-01$ | $1.183038 \mathrm{E}+00$ | -1.584200E-01 |
| 10 | $4.500000 \mathrm{E}-01$ | $1.179822 \mathrm{E}+00$ | -9.973146E-02 |
| 11 | $4.625000 \mathrm{E}-01$ | $1.178748 \mathrm{E}+00$ | -7.233893E-02 |
| 12 | $4.750000 \mathrm{E}-01$ | $1.178007 \mathrm{E}+00$ | -4.638248E-02 |
| 13 | $4.812500 \mathrm{E}-01$ | $1.177756 \mathrm{E}+00$ | -3.399763E-02 |
| 14 | $4.875000 \mathrm{E}-01$ | $1.177582 \mathrm{E}+00$ | -2.205547E-02 |
| 15 | $4.937500 \mathrm{E}-01$ | $1.177480 \mathrm{E}+00$ | -1.061177E-02 |
| 16 | $5.000000 \mathrm{E}-01$ | $1.177447 \mathrm{E}+00$ | -1.479182E-07 |
| 17 | $5.062500 \mathrm{E}-01$ | $1.177480 \mathrm{E}+00$ | 1.061153E-02 |
| 18 | $5.125000 \mathrm{E}-01$ | $1.177582 \mathrm{E}+00$ | $2.205518 \mathrm{E}-02$ |
| 19 | $5.187500 \mathrm{E}-01$ | $1.177756 \mathrm{E}+00$ | $3.399727 \mathrm{E}-02$ |
| 20 | $5.250000 \mathrm{E}-01$ | $1.178007 \mathrm{E}+00$ | $4.638219 \mathrm{E}-02$ |
| 21 | $5.375000 \mathrm{E}-01$ | $1.178748 \mathrm{E}+00$ | 7.233876E-02 |
| 22 | $5.500000 \mathrm{E}-01$ | $1.179822 \mathrm{E}+00$ | $9.973124 \mathrm{E}-02$ |
| 23 | $5.750000 \mathrm{E}-01$ | $1.183038 \mathrm{E}+00$ | $1.584199 \mathrm{E}-01$ |
| 24 | $6.000000 \mathrm{E}-01$ | $1.187783 \mathrm{E}+00$ | $2.219350 \mathrm{E}-01$ |
| 25 | $6.666667 \mathrm{E}-01$ | $1.208786 \mathrm{E}+00$ | 4.135205E-01 |
| 26 | 7.333333E-01 | $1.243628 \mathrm{E}+00$ | $6.373190 \mathrm{E}-01$ |
| 27 | $8.000000 \mathrm{E}-01$ | $1.294526 \mathrm{E}+00$ | $8.958461 \mathrm{E}-01$ |
| 28 | $8.666667 \mathrm{E}-01$ | 1.363953E+00 | $1.194307 \mathrm{E}+00$ |
| 29 | $9.111111 \mathrm{E}-01$ | $1.421951 \mathrm{E}+00$ | 1.419465E+00 |
| 30 | $9.555556 \mathrm{E}-01$ | $1.490495 \mathrm{E}+00$ | 1. $669566 \mathrm{E}+00$ |
| 31 | $1.000000 \mathrm{E}+00$ | $1.570796 \mathrm{E}+00$ | 1.949336E+00 |
| Error estimates |  | $3.448358 \mathrm{E}-06$ | $5.549869 \mathrm{E}-05$ |

## BVPMS

Solves a (parameterized) system of differential equations with boundary conditions at two points, using a multiple-shooting method.

## Required Arguments

FCNEQN - User-supplied SUBROUTINE to evaluate derivatives. The usage is CALL FCNEQN (NEQNS, T, Y, P, DYDT), where

NEQNS - Number of equations. (Input)
T - Independent variable, $t$. (Input)
Y - Array of length NEQNS containing the dependent variable. (Input)
P - Continuation parameter used in solving highly nonlinear problems. (Input) See Comment 4.
DYDT - Array of length NEQNS containing $y^{\prime}$ at T. (Output)
The name FCNEQN must be declared EXTERNAL in the calling program.
FCNJAC - User-supplied SUBROUTINE to evaluate the Jacobian. The usage is CALL
FCNJAC (NEQNS, T, Y, P, DYPDY), where
NEQNS - Number of equations. (Input)
$T$ - Independent variable. (Input)
Y - Array of length NEQNS containing the dependent variable. (Input)
P - Continuation parameter used in solving highly nonlinear problems. (Input) See Comment 4.
DYPDY - Array of size NEQNS by NEQNS containing the Jacobian. (Output)
The entry $\operatorname{DYPDY}(i, j)$ contains the partial derivative $\partial f_{i} / \partial y_{j}$ evaluated at $(t, y)$.
The name FCNJAC must be declared EXtERNAL in the calling program.
FCNBC - User-supplied SUBROUTINE to evaluate the boundary conditions. The usage is CALL FCNBC (NEQNS, YLEFT, YRIGHT, P, H), where

NEQNS - Number of equations. (Input)
YLEFT - Array of length NEQNS containing the values of Y at TLEFT. (Input)
YRIGHT - Array of length NEQNS containing the values of Y at TRIGHT. (Input)
$\mathrm{P}-$ Continuation parameter used in solving highly nonlinear problems. (Input) See Comment 4.
H - Array of length NEQNS containing the boundary function values. (Output) The computed solution satisfies (within BTOL) the conditions $h_{i}=0, i=1, \ldots$, NEQNS.

The name FCNBC must be declared EXTERNAL in the calling program.
TLEFT - The left endpoint. (Input)

TRIGHT - The right endpoint. (Input)
NMAX - Maximum number of shooting points to be allowed. (Input)
If NINIT is nonzero, then NMAX must equal NINIT. It must be at least 2.
NFINAL - Number of final shooting points, including the endpoints. (Output)
TFINAL - Vector of length NMAX containing the final shooting points. (Output)
Only the first NFINAL points are significant.

YFINAL - Array of size NEQNS by NMAX containing the values of Y at the points in TFINAL. (Output)

## Optional Arguments

NEQNS - Number of differential equations. (Input)
DTOL - Differential equation error tolerance. (Input) An attempt is made to control the local error in such a way that the global error is proportional to DTOL.
Default: $\mathrm{DTOL}=1.0 \mathrm{e}-4$.

BTOL - Boundary condition error tolerance. (Input)
The computed solution satisfies the boundary conditions, within BTOL tolerance.
Default: $\mathrm{BTOL}=1.0 \mathrm{e}-4$.
MAXIT — Maximum number of Newton iterations allowed. (Input)
Iteration stops if convergence is achieved sooner. Suggested values are MAXIT $=2$ for linear problems and MAXIT $=9$ for nonlinear problems.
Default: MAXIT $=9$.
NINIT - Number of shooting points supplied by the user. (Input) It may be 0 . A suggested value for the number of shooting points is 10 . Default: NINIT $=0$.

TINIT - Vector of length NINIT containing the shooting points supplied by the user.
(Input)
If NINIT $=0$, then TINIT is not referenced and the routine chooses all of the shooting points. This automatic selection of shooting points may be expensive and should only be used for linear problems. If NINIT is nonzero, then the points must be an increasing sequence with $\operatorname{TINIT}(1)=$ TLEFT and $\operatorname{TINIT}(\operatorname{NINIT})=$ TRIGHT. By default, TINIT is not used.

YINIT - Array of size NEQNS by NINIT containing an initial guess for the values of Y at the points in TINIT. (Input)
YINIT is not referenced if NINIT $=0$. By default, YINIT is not used.

LDYINI - Leading dimension of YINIT exactly as specified in the dimension statement of the calling program. (Input)
Default: LDYINI $=\operatorname{size}($ YINIT, 1$)$.

LDYFIN - Leading dimension of YFINAL exactly as specified in the dimension statement of the calling program. (Input)
Default: LDYFIN = size (YFINAL, 1 ).

## FORTRAN 90 Interface

Generic: CALL BVPMS (FCNEQN, FCNJAC, FCNBC, TLEFT, TRIGHT, NMAX, NFINAL, TFINAL, YFINAL [,...])

Specific: The specific interface names are S_BVPMS and D_BVPMS.

## FORTRAN 77 Interface

Single:

```
CALL BVPMS (FCNEQN, FCNJAC, FCNBC, NEQNS, TLEFT, TRIGHT,
DTOL, BTOL, MAXIT, NINIT, TINIT, YINIT, LDYINI, NMAX,
NFINAL, TFINAL, YFINAL, LDYFIN)
```

Double: The double precision name is DBVPMS.

## Description

Define $N=$ NEQNS, $M=$ NFINAL, $t_{a}=$ TLEFT and $t_{b}=$ TRIGHt. The routine BVPMS uses a multipleshooting technique to solve the differential equation system $y^{\prime}=f(t, y)$ with boundary conditions of the form

$$
h_{k}\left(y_{1}\left(t_{a}\right), \ldots, y_{N}\left(t_{a}\right), y_{1}\left(t_{b}\right), \ldots, y_{N}\left(t_{b}\right)\right)=0 \quad \text { for } k=1, \ldots, N
$$

A modified version of IVPRK is used to compute the initial-value problem at each "shot." If there are $M$ shooting points (including the endpoints $t_{a}$ and $t_{b}$ ), then a system of $N M$ simultaneous nonlinear equations must be solved. Newton's method is used to solve this system, which has a Jacobian matrix with a "periodic band" structure. Evaluation of the $N M$ functions and the $N M \times N M$ (almost banded) Jacobian for one iteration of Newton's method is accomplished in one pass from $t_{a}$ to $t_{b}$ of the modified IVPRK, operating on a system of $N(N+1)$ differential equations. For most problems, the total amount of work should not be highly dependent on $M$. Multiple shooting avoids many of the serious ill-conditioning problems that plague simple shooting methods. For more details on the algorithm, see Sewell (1982).

The boundary functions should be scaled so that all components $h_{k}$ are of comparable magnitude since the absolute error in each is controlled.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2 PMS / DB2 PMS. The reference is:
```
CALL B2PMS (FCNEQN, FCNJAC, FCNBC, NEQNS,TLEFT,TRIGHT, DTOL, BTOL,
MAXIT, NINIT, TINIT, YINIT, LDYINI, NMAX, NFINAL, TFINAL, YFINAL, LDYFIN,
WORK, IWK)
```

The additional arguments are as follows:

```
WORK - Work array of length NEQNS * (NEQNS + 1) (NMAX + 12) +
    NEQNS + 30 .
```

IWK - Work array of length NEQNS.
2. Informational errors

Type Code
15 Convergence has been achieved; but to get acceptably accurate approximations to $y(t)$, it is often necessary to start an initial-value solver, for example IVPRK, at the nearest TFINAL $(i)$ point to $t$ with $t$ $\geq$ TFINAL $(i)$. The vectors $\operatorname{YFINAL}(j, i), j=1, \ldots$, NEQNS are used as the initial values.

4

42 More than NMAX shooting points are needed for stability.
43 Newton's iteration did not converge in MAXIT iterations. If the problem is linear, do an extra iteration. If this error still occurs, check that the routine FCNJAC is giving the correct derivatives. If this does not fix the problem, see Comment 3.
4
4 Linear-equation solver failed. The problem may not have a unique solution, or the problem may be highly nonlinear. In the latter case, see Comment 3.
3. Many linear problems will be successfully solved using program-selected shooting points. Nonlinear problems may require user effort and input data. If the routine fails, then increase NMAX or parameterize the problem. With many shooting points the program essentially uses a finite-difference method, which has less trouble with nonlinearities than shooting methods. After a certain point, however, increasing the number of points will no longer help convergence. To parameterize the problem, see Comment 4.
4. If the problem to be solved is highly nonlinear, then to obtain convergence it may be necessary to embed the problem into a one-parameter family of boundary value problems, $y^{\prime}=f(t, y, p), h\left(y\left(t_{a}, t_{b}, p\right)\right)=0$ such that for $p=0$, the problem is simple, e.g., linear; and for $p=1$, the stated problem is solved. The routine BVPMS/DBVPMS automatically moves the parameter from $p=0$ toward $p=1$.
5. This routine is not recommended for stiff systems of differential equations.

## Example

The differential equations that model an elastic beam are (see Washizu 1968, pages 142-143):

$$
\begin{aligned}
& \mathbf{M}_{x x}-\frac{\mathbf{N M}}{\mathbf{E I}}+\mathbf{L}(x)=0 \\
& \mathbf{E I W}_{x x}+\mathbf{M}=0 \\
& \mathbf{E A}_{0}\left(\mathbf{U}_{x}+\mathbf{W}_{x}^{2} / 2\right)-\mathbf{N}=0 \\
& \mathbf{N}_{x}=0
\end{aligned}
$$

where $\mathbf{U}$ is the axial displacement, $\mathbf{W}$ is the transverse displacement, $\mathbf{N}$ is the axial force, $\mathbf{M}$ is the bending moment, $\mathbf{E}$ is the elastic modulus, $\mathbf{I}$ is the moment of inertia, $\mathbf{A}_{0}$ is the cross-sectional area, and $\mathbf{L}(x)$ is the transverse load.
Assume we have a clamped cylindrical beam of radius 0.1 in , a length of 10 in , and an elastic modulus $\mathbf{E}=10.6 \times 10^{6} \mathrm{lb} / \mathrm{in}^{2}$. Then, $\mathbf{I}=0.784 \times 10^{-4}$, and $\mathbf{A}_{0}=\pi 10^{-2} \mathrm{in}^{2}$, and the boundary conditions are $\mathbf{U}=\mathbf{W}=\mathbf{W}_{x}=0$ at each end. If we let $y_{1}=\mathbf{U}, y_{2}=\mathbf{N} / \mathbf{E} A_{0}, y_{3}=\mathbf{W}, y_{4}=\mathbf{W}_{x}$, $y_{5}=\mathbf{M} / \mathbf{E I}$, and $y_{6}=\mathbf{M}_{x} / \mathbf{E I}$, then the above nonlinear equations can be written as a system of six first-order equations.

$$
\begin{aligned}
& y_{1}^{\prime}=y_{2}-\frac{y_{4}^{2}}{2} \\
& y_{2}^{\prime}=0 \\
& y_{3}^{\prime}=y_{4} \\
& y_{4}^{\prime}=-y_{5} \\
& y_{5}^{\prime}=y_{6} \\
& y_{6}^{\prime}=\frac{\mathbf{A}_{0} y_{2} y_{5}}{\mathbf{I}}-\frac{\mathbf{L}(x)}{\mathbf{E I}}
\end{aligned}
$$

The boundary conditions are $y_{1}=y_{3}=y_{4}=0$ at $x=0$ and at $x=10$. The loading function is $\mathbf{L}(x)=-2$, if $3 \leq x \leq 7$, and is zero elsewhere.
The material parameters, $\mathbf{A}_{0}=A 0, \mathbf{I}=A I$, and $\mathbf{E}$, are passed to the evaluation subprograms using the common block PARAM.

```
USE BVPMS INT
USE UMACH_INT
IMPLICIT NONE
INTEGER LDY, NEQNS, NMAX
PARAMETER (NEQNS=6, NMAX=21, LDY=NEQNS)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, MAXIT, NFINAL, NINIT, NOUT
REAL TOL, X(NMAX), XLEFT, XRIGHT, Y(LDY,NMAX)
                                    SPECIFICATIONS FOR COMMON /PARAM/
COMMON /PARAM/ A0, A1, E
REAL A0, A1, E
    SPECIFICATIONS FOR INTRINSICS
INTRINSIC REAL
REAL REAL
    SPECIFICATIONS FOR SUBROUTINES
EXTERNAL FCNBC, FCNEQN, FCNJAC
    Set material parameters
```

```
    AO = 3.14E-2
    A1 = 0.784E-4
    E=10.6E6
!
    XLEFT = 0.0
    XRIGHT = 10.0
    MAXIT = 19
    NINIT = NMAX
    Y = 0.0EO
!
    DO 10 I=1, NINIT
        X(I) = XLEFT + REAL(I-1)/REAL(NINIT-1)*(XRIGHT-XLEFT)
    10 CONTINUE
!
    CALL BVPMS (FCNEQN, FCNJAC, FCNBC, XLEFT, XRIGHT, NMAX, NFINAL, &
    X, Y, MAXIT=MAXIT, NINIT=NINIT, TINIT=X, YINIT=Y)
!
    CALL UMACH (2, NOUT)
    WRITE (NOUT,'(26X,A/12X,A,10X,A,7X,A)') 'Displacement', &
                                    'X', 'Axial', 'Transvers'// &
                            'e'
WRITE (NOUT,'(F15.1,1P2E15.3)') (X (I),Y(1,I),Y(3,I),I=1,NFINAL)
END
SUBROUTINE FCNEQN (NEQNS, X, Y, P, DYDX)
    INTEGER NEQNS
    REAL X, P, Y(NEQNS), DYDX (NEQNS)
    SPECIFICATIONS FOR LOCAL VARIABLES
    REAL FORCE
    COMMON /PARAM/ A0, A1, E
    REAL A0, A1, E
    FORCE = 0.0
    IF (X.GT.3.0 .AND. X.LT.7.0) FORCE = -2.0
    DYDX(1) = Y(2) - P*0.5*Y(4)**2
    DYDX(2) = 0.0
    DYDX(3) = Y(4)
    DYDX(4) = -Y(5)
    DYDX(5) = Y(6)
    DYDX(6) = P*A0*Y(2)*Y(5)/A1 - FORCE/E/A1
    RETURN
    END
    SUBROUTINE FCNBC (NEQNS, YLEFT, YRIGHT, P, F)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NEQNS
    REAL P, YLEFT (NEQNS), YRIGHT (NEQNS), F(NEQNS)
    COMMON /PARAM/ A0, A1, E
    REAL A0, A1, E
    Define boundary conditions
F(1) = YLEFT(1)
F(2) = YLEFT(3)
F(3) = YLEFT(4)
```

```
    F(4) = YRIGHT(1)
    F(5) = YRIGHT (3)
    F(6) = YRIGHT (4)
    RETURN
    END
    SUBROUTINE FCNJAC (NEQNS, X, Y, P, DYPDY)
!
    INTEGER NEQNS
    REAL X, P, Y(NEQNS), DYPDY (NEQNS,NEQNS)
    COMMON /PARAM/ A0, A1, E
    REAL A0, A1, E
! SPECIFICATIONS FOR SUBROUTINES
    DYPDY = 0.0EO
    DYPDY (1,2) = 1.0
    DYPDY(1,4) = -P*Y(4)
    DYPDY (3,4) = 1.0
    DYPDY (4,5) = -1.0
    DYPDY (5,6) = 1.0
    DYPDY (6,2) = P*Y(5)*A0/A1
    DYPDY (6,5) = P*Y(2)*A0/A1
    RETURN
    END
```

        Output
    |  | Displacement |  |
| ---: | ---: | ---: |
| X | Axial | Transverse |
| 0.0 | $1.631 \mathrm{E}-11$ | $-8.677 \mathrm{E}-10$ |
| 5.0 | $1.914 \mathrm{E}-05$ | $-1.273 \mathrm{E}-03$ |
| 10.0 | $2.839 \mathrm{E}-05$ | $-4.697 \mathrm{E}-03$ |
| 15.0 | $2.461 \mathrm{E}-05$ | $-9.688 \mathrm{E}-03$ |
| 20.0 | $1.008 \mathrm{E}-05$ | $-1.567 \mathrm{E}-02$ |
| 25.0 | $-9.550 \mathrm{E}-06$ | $-2.206 \mathrm{E}-02$ |
| 30.0 | $-2.721 \mathrm{E}-05$ | $-2.830 \mathrm{E}-02$ |
| 35.0 | $-3.644 \mathrm{E}-05$ | $-3.382 \mathrm{E}-02$ |
| 40.0 | $-3.379 \mathrm{E}-05$ | $-3.811 \mathrm{E}-02$ |
| 45.0 | $-2.016 \mathrm{E}-05$ | $-4.083 \mathrm{E}-02$ |
| 50.0 | $-4.414 \mathrm{E}-08$ | $-4.176 \mathrm{E}-02$ |
| 55.0 | $2.006 \mathrm{E}-05$ | $-4.082 \mathrm{E}-02$ |
| 60.0 | $3.366 \mathrm{E}-05$ | $-3.810 \mathrm{E}-02$ |
| 65.0 | $3.627 \mathrm{E}-05$ | $-3.380 \mathrm{E}-02$ |
| 70.0 | $2.702 \mathrm{E}-05$ | $-2.828 \mathrm{E}-02$ |
| 75.0 | $9.378 \mathrm{E}-06$ | $-2.205 \mathrm{E}-02$ |
| 80.0 | $-1.021 \mathrm{E}-05$ | $-1.565 \mathrm{E}-02$ |
| 85.0 | $-2.468 \mathrm{E}-05$ | $-9.679 \mathrm{E}-03$ |
| 90.0 | $-2.842 \mathrm{E}-05$ | $-4.692 \mathrm{E}-03$ |
| 95.0 | $-1.914 \mathrm{E}-05$ | $-1.271 \mathrm{E}-03$ |
| 100.0 | $0.000 \mathrm{E}+00$ | $0.000 \mathrm{E}+00$ |

## DASPG

Solves a first order differential-algebraic system of equations, $g\left(t, y, y^{\prime}\right)=0$, using the PetzoldGear BDF method.

## Required Arguments

$\boldsymbol{T}$ - Independent variable, $t$. (Input/Output)
Set T to the starting value $t_{0}$ at the first step.
TOUT - Final value of the independent variable. (Input)
Update this value when re-entering after output, $I D O=2$.
IDO - Flag indicating the state of the computation. (Input/Output)

## IDO State

1 Initial entry
2 Normal re-entry after obtaining output
3 Release workspace
4 Return because of an error condition
The user sets IDO $=1$ or IDO $=3$. All other values of IDO are defined as output. The initial call is made with IDO $=1$ and $\mathrm{T}=t_{0}$. The routine then sets IDO $=2$, and this value is used for all but the last entry that is made with IDO $=3$. This call is used to release workspace and other final tasks. Values of IDO larger than 4 occur only when calling the second-level routine D2SPG and using the options associated with reverse communication.
$\boldsymbol{Y}$ - Array of size NEQ containing the dependent variable values, $y$. This array must contain initial values. (Input/Output)

YPR - Array of size NEQ containing derivative values, $y^{\prime}$. This array must contain initial values. (Input/Output)
The routine will solve for consistent values of $y^{\prime}$ to satisfy the equations at the starting point.
$\boldsymbol{G C N}$ - User-supplied SUBROUTINE to evaluate $g\left(t, y, y^{\prime}\right)$. The usage is CALL GCN (NEQ, T, Y, YPR, GVAL), where GCN must be declared EXTERNAL in the calling program. The routine will solve for values of $y^{\prime}\left(t_{0}\right)$ so that $g\left(t_{0}, y, y^{\prime}\right)=0$. The user can signal that $g$ is not defined at requested values of $\left(t, y, y^{\prime}\right)$ using an option. This causes the routine to reduce the step size or else quit.

NEQ - Number of differential equations. (Input)
$T$ - Independent variable. (Input)
Y - Array of size NEQ containing the dependent variable values $y(t)$. (Input)
YPR - Array of size NEQ containing the derivative values $y^{\prime}(t)$. (Input)
GVAL - Array of size NEQ containing the function values, $g\left(t, y, y^{\prime}\right)$. (Output)

## Optional Arguments

NEQ - Number of differential equations. (Input)
Default: NEQ $=\operatorname{size}(\mathrm{y}, 1)$

## FORTRAN 90 Interface

Generic: CALL DASPG (T, TOUT, IDO, Y, YPR, GCN [,...])
Specific: The specific interface names are S_DASPG and D_DASPG.

## FORTRAN 77 Interface

Single: CALL DASPG (NEQ, T, TOUT, IDO, Y, YPR, GCN)
Double: The double precision name is DDASPG.

## Description

Routine DASPG finds an approximation to the solution of a system of differential-algebraic equations $g\left(t, y, y^{\prime}\right)=0$, with given initial data for $y$ and $y^{\prime}$. The routine uses BDF formulas, appropriate for systems of stiff ODEs, and attempts to keep the global error proportional to a userspecified tolerance. See Brenan et al. (1989). This routine is efficient for stiff systems of index 1 or index 0. See Brenan et al. (1989) for a definition of index. Users are encouraged to use DOUBLE PRECISION accuracy on machines with a short REAL precision accuracy. The examples given below are in REAL accuracy because of the desire for consistency with the rest of IMSL MATH/LIBRARY examples. The routine DASPG is based on the code DASSL designed by L. Petzold (1982-1990).

## Comments

Users can often get started using the routine DASPG/DDASPG without reading beyond this point in the documentation. There is often no reason to use options when getting started. Those readers who do not want to use options can turn directly to the first two examples. The following tables give numbers and key phrases for the options. A detailed guide to the options is given below in Comment 2.

| Value | Brief or Key Phrase for INTEGER Option |
| :--- | :--- |
| $\mathbf{6}$ | INTEGER option numbers |
| 7 | Floating-point option numbers |


| Value | Brief or Key Phrase for INTEGER Option |
| :---: | :---: |
| IN(1) | First call to DASPG, D2SPG |
| IN(2) | Scalar or vector tolerances |
| IN(3) | Return for output at intermediate steps |
| IN(4) | Creep up on special point, TSTOP |
| IN(5) | Provide (analytic) partial derivative formulas |
| IN(6) | Maximum number of steps |
| IN(7) | Control maximum step size |
| IN(8) | Control initial step size |
| IN(9) | Not Used |
| IN(10) | Constrain dependent variables |
| IN(11) | Consistent initial data |
| IN(12-15) | Not Used |
| IN(16) | Number of equations |
| IN(17) | What routine did, if any errors |
| IN(18) | Maximum BDF order |
| IN(19) | Order of BDF on next move |
| IN(20) | Order of BDF on previous move |
| IN(21) | Number of steps |
| IN(22) | Number of $g$ evaluations |
| IN(23) | Number of derivative matrix evaluations |
| IN(24) | Number of error test failures |
| IN(25) | Number of convergence test failures |
| IN(26) | Reverse communiction for $g$ |
| IN(27) | Where is $g$ stored? |
| IN(28) | Panic flag |
| IN(29) | Reverse communication, for partials |
| IN(30) | Where are partials stored? |
| IN(31) | Reverse communication, for solving |
| IN(32) | Not Used |
| IN(33) | Where are vector tolerances stored? |
| IN(34) | Is partial derivative array allocated? |
| IN(35) | User's work arrays sizes are checked |
| IN(36-50) | Not used |

Table 1. Key Phrases for Floating-Point Options

| Value | Brief or Key Phrase for Floating-Point Option |
| :--- | :--- |
| INR(1) | Value of $t$ |
| INR(2) | Farthest internal $t$ vaue of integration |
| INR(3) | Value of TOUT |
| INR(4) | A stopping point of integration before TOUT |
| INR(5) | Values of two scalars ATOL, RTOL |
| INR(6) | Initial step size to use |
| INR(7) | Maximum step allowed |
| INR(8) | Condition number reciprocal |
| INR(9) | Value of $c_{j}$ for partials |
| INR(10) | Step size on the next move |
| INR(11) | Step size on the previous move |
| INR(12-20) | Not Used |

Table 2. Number and Key Phrases for Floating-Point Options

1. Workspace may be explicitly provided, and many of the options utilized by directly calling D2SPG/DD2SPG. The reference is:

CALL D2SPG (N, T, TOUT, IDO, Y, YPR, GCN, JGCN, IWK, WK)
The additional arguments are as follows:

## IDO State

5 Return for evaluation of $g\left(t, y, y^{\prime}\right)$
6 Return for evaluation of matrix $A=\left[\partial g / \partial y+c_{j} \partial g / \partial y^{\prime}\right]$
7 Return for factorization of the matrix $A=\left[\partial g / \partial y+c_{j} \partial g / \partial y^{\prime}\right]$
8 Return for solution of $A \Delta y=\Delta g$

These values of IDO occur only when calling the second-level routine D2SPG and using options associated with reverse communication. The routine D2SPG/DD2SPG is reentered.

GCN - A Fortran SUBROUTINE to compute $g\left(t, y, y^{\prime}\right)$. This routine is normally provided by the user. That is the default case. The dummy IMSL routine DGSPG/DDGSPG may be used as this argument when $g\left(t, y, y^{\prime}\right)$ is evaluated by reverse communication. In either case, a name must be declared in a Fortran EXTERNAL statement. If usage of the dummy IMSL routine is intended, then the name DGSPG/DDGSPG should be specified. The dummy IMSL routine will never
be called under this optional usage of reverse communication. An example of reverse communication for evaluation of $g$ is given in Example 4.
$J G C N$ - A Fortran SUBROUTINE to compute partial derivatives of $g\left(t, y, y^{\prime}\right)$. This routine may be provided by the user. The dummy IMSL routine DJSPG/DDJSPG may be used as this argument when partial derivatives are computed using divided differences. This is the default. The dummy routine is not called under default conditions. If partial derivatives are to be explicitly provided, the routine JGCN must be written by the user or reverse communication can be used. An example of reverse communication for evaluation of the partials is given in Example 4.

If the user writes a routine with the fixed name DJSPG/DDJSPG, then partial derivatives can be provided while calling DASPG. An option is used to signal that formulas for partial derivatives are being supplied. This is illustrated in Example 3. The name of the partial derivative routine must be declared in a Fortran EXTERNAL statement when calling D2SPG. If usage of the dummy IMSL routine is intended, then the name DJSPG/DDJSPG should be specified for this EXTERNAL name. Whenever the user provides partial derivative evaluation formulas, by whatever means, that must be noted with an option. Usage of the derivative evaluation routine is
CALL JGCN (N, T, Y, YPR, CJ, PDG, LDPDG) where

## Arg Definition

N Number of equations. (Input)
$\mathrm{T} \quad$ Independent variable, $t$. (Input)
$\mathrm{Y} \quad$ Array of size N containing the values of the dependent variables, $y$. (Input)
YPR Array of size N containing the values of the derivatives, $y^{\prime}$. (Input)

CJ The value $c_{j}$ used in computing the partial derivatives returned in PDG. (Input)

PDG Array of size LDPDG * N containing the partial derivatives $A=\left[\partial g / \partial y+c_{j} \partial g /\right.$ $\left.\partial y^{\prime}\right]$. Each nonzero derivative entry $a_{i j}$ is returned in the array location $\operatorname{PDG}(i, j)$. The array contents are zero when the routine is called. Thus, only the nonzero derivatives have to be defined in the routine JGCN. (Output)

LDPDG The leading dimension of PDG. Normally, this value is $N$. It is a value larger than N under the conditions explained in option $\mathbf{1 6}$ of LSLRG (Chapter 1, Linear Systems).

JGCN must be declared EXTERNAL in the calling program.

IWK - Work array of integer values. The size of this array is $35+\mathrm{N}$. The contents of IWK must not be changed from the first call with IDO $=1$ until after the final call with $I D O=3$.

WK - Work ahrray of floating-point values in the working precision. The size of this array is $41+($ MAXORD +6$) N+(N+K) N(1-L)$ where $K$ is determined from the values IVAL (3) and IVAL(4) of option 16 of LSLRG (Chapter 1, Linear Systems). The value of $L$ is 0 unless option $\mathbf{I N}(34)$ is used to avoid allocation of the array containing the partial derivatives. With the use of this option, L can be set to 1 . The contents of array WK must not be changed from the first call with IDO $=1$ until after the final call.
2. Integer and Floating-Point Options with Chapter 11 Options Manager

The routine DASPG allows the user access to many interface parameters and internal working variables by the use of options. The options manager subprograms IUMAG and SUMAg/DUMAG (Chapter 11, Utilities), are used to change options from their default values or obtain the current values of required parameters.

Options of type INTEGER:
6 This is the list of numbers used for INTEGER options. Users will typically call this option first to get the numbers, $\operatorname{IN}(\mathrm{I}), I=1,50$. This option has 50 entries. The default values are $I N(I)=I+50, I=1,50$.

7 This is the list of numbers used for REAL and DOUBLE PRECISION options. Users will typically call this option first to get the numbers, $\operatorname{INR}(\mathrm{I}), \mathrm{I}=1,20$. This option has 20 entries. The default values are $\operatorname{INR}(I)=I+50, I=1,20$.

IN(1) This is the first call to the routine DASPG or D2SPG. Value is 0 for the first call, 1 for further calls. Setting IDO $=1$ resets this option to its default. Default value is 0 .
$\mathbf{I N}(2)$ This flag controls the kind of tolerances to be used for the solution. Value is 0 for scalar values of absolute and relative tolerances applied to all components. Value is 1 when arrays for both these quantities are specified. In this case, use D2SPG. Increase the size of $W \mathrm{~K}$ by $2 * \mathrm{~N}$, and supply the tolerance arrays at the end of wk. Use option $\mathbf{I N}(33)$ to specify the offset into wk where the 2 N array values are to be placed: all ATOL values are followed by all RTOL values. Also see $\mathbf{I N}(33)$. Default value is 0 .

IN(3) This flag controls when the code returns to the user with output values of $y$ and $y^{\prime}$. If the value is 0 , it returns to the user at $\mathrm{T}=$ TOUT only. If the value is 1 , it returns to the user at an internal working step. Default value is 0 .

IN(4) This flag controls whether the code should integrate past a special point, TSTOP, and then interpolate to get $y$ and $y^{\prime}$ at TOUT. If the value is 0 , this is permitted. If the value is 1 , the code assumes the equations either change on the alternate side
of TSTOP or they are undefined there. In this case, the code creeps up to TSTOP in the direction of integration. The value of TSTOP is set with option INR(4). Default value is 0 .

IN(5) This flag controls whether partial derivatives are computed using divided onesided differences, or they are to be computed using user-supplied evaluation formulas. If the value is 0 , use divided differences. If the value is 1 , use formulas for the partial derivatives. See Example 3 for an illustration of one way to do this. Default value is 0 .

IN(6) The maximum number of steps. Default value is 500 .
IN(7) This flag controls a maximum magnitude constraint for the step size. If the value is 0 , the routine picks its own maximum. If the value is 1 , a maximum is specified by the user. That value is set with option number INR(7). Default value is 0 .

IN(8) This flag controls an initial value for the step size. If the value is 0 , the routine picks its own initial step size. If the value is 1 , a starting step size is specified by the user. That value is set with option number INR(6). Default value is 0 .
$\mathbf{I N}(9)$ Not used. Default value is 0 .

IN(10) This flag controls attempts to constrain all components to be nonnegative. If the value is 0 , no constraints are enforced. If value is 1 , constraint is enforced. Default value is 0 .

IN(11) This flag controls whether the initial values $\left(t, y, y^{\prime}\right)$ are consistent. If the value is $0, g\left(t, y, y^{\prime}\right)=0$ at the initial point. If the value is 1 , the routine will try to solve for $y^{\prime}$ to make this equation satisfied. Default value is 0 .
$\mathbf{I N}(12-15)$ Not used. Default value is 0 for each option.
$\mathbf{I N}(16) \quad$ The number of equations in the system, $n$. Default value is 0 .
$\mathbf{I N}(17)$ This value reports what the routine did. Default value is 0 .

| Value | Explanation |
| :--- | :--- |
| 1 | A step was taken in the intermediate output mode. The value <br> TOUT has not been reached. |
| 2 | The integration to exactly TSTOP was completed. |
| 3 | The integration to TSTOP was completed by stepping past TSTOP <br> and interpolating to evaluate $y$ and $y^{\prime}$. |
| -1 | Too many steps taken. |
| -2 | Error tolerances are too small. |
| -3 | A pure relative error tolerance can't be satisfied. |


| Value | Explanation |
| :--- | :--- |
| -6 | There were repeated error test failures on the last step. |
| -7 | The BDF corrector equation solver did not converge. |
| -8 | The matrix of partial derivatives is singular. |
| -10 | The BDF corrector equation solver did not converge because the <br> evaluation failure flag was raised. |
| -11 | The evaluation failure flag was raised to quit. |
| -12 | The iteration for the initial vaule of $y^{\prime}$ did not converge. |
| -33 | There is a fatal error, perhaps caused by invalid input. |

Table 3. What the Routine DASPG or D2SPG Did
IN(18) The maximum order of BDF formula the routine should use. Default value is 5 .

IN(19) The order of the BDF method the routine will use on the next step. Default value is $\operatorname{IMACH}(5)$.

IN(20) The order of the BDF method used on the last step. Default value is IMACH(5).
$\mathbf{I N}(21)$ The number of steps taken so far. Default value is 0 .
$\mathbf{I N}(22) \quad$ The number of times that $g$ has been evaluated. Default value is 0 .
IN(23) The number of times that the partial derivative matrix has been evaluated. Default value is 0 .

IN(24) The total number of error test failures so far. Default value is 0 .
IN(25) The total number of convergence test failures so far. This includes singular iteration matrices. Default value is 0 .

IN(26) Use reverse communication to evaluate $g$ when this value is 0 . If the value is 1 , forward communication is used. Use the routine D2SPG for reverse communication. With reverse communication, a return will be made with IDO $=5$. Compute the value of $g$, place it into the array $W K$ at the offset obtained with option $\mathbf{I N}(27)$, and re-enter the routine. Default value is 1 .
$\mathbf{I N}(27)$ The user is to store the evaluated function $g$ during reverse communication in the work array wK using this value as an offset. Default value is IMACH(5).
$\mathbf{I N}(28) \quad$ This value is a "panic flag." After an evaluation of $g$, this value is checked. The value of $g$ is used if the flag is 0 . If it has the value -1 , the routine reduces the step size and possibly the order of the BDF. If the value is -2 , the routine returns control to the user immediately. This option is also used to signal a singular or poorly conditioned partial derivative matrix encountered during the
factor phase in reverse communication. Use a nonzero value when the matrix is singular. Default value is 0 .

IN(29) Use reverse communication to evaluate the partial derivative matrix when this value is 0 . If the value is 1 , forward communication is used. Use the routine D2SPG for reverse communication. With reverse communication, a return will be made with IDO $=6$. Compute the partial derivative matrix $A$ and re-enter the routine. If forward communication is used for the linear solver, return the partials using the offset into the array WK. This offset value is obtained with option $\mathbf{I N}(30)$. Default value is 1 .

IN(30) The user is to store the values of the partial derivative matrix $A$ by columns in the work array wK using this value as an offset. The option 16 for LSLRG is used here to compute the row dimension of the internal working array that contains $A$. Users can also choose to store this matrix in some convenient form in their calling program if they are providing linear system solving using reverse communication. See options IN(31) and IN(34). Default value is IMACH(5).

IN(31) Use reverse communication to solve the linear system $A \Delta y=\Delta g$ if this value is 0 . If the value is 1 , use forward communication into the routines L2CRG and LFSRG (Chapter 1, Linear Systems) for the linear system solving. Return the solution using the offset into the array wK where $g$ is stored. This offset value is obtained with option $\mathbf{I N}(27)$. With reverse communication, a return will be made with IDO $=7$ for factorization of $A$ and with $I D O=8$ for solving the system. Re-enter the routine in both cases. If the matrix A is singular or poorly conditioned, raise the "panic flag," option IN(28), during the factorization. Default value is 1 .

IN(32) Not used. Default value is 0 .
$\mathbf{I N}(33)$ This value is used when $\mathbf{I N}(2)$ is set to 1 , indicating that arrays of absolute and relative tolerances are input in the WK array of D2SPG. Set this parameter to the offset, ioff, into $W K$ where the tolerances are stored. Increase the size of WK by $2 *_{\mathrm{N}}$, and store tolerance values beginning at ioff $=$ size(WK) $-2 *_{\mathrm{N}}+1$.
Absolute tolerances will be stored in $W K(i \circ f f+i-1)$ for $i=1, N$ and relative tolerances will be stored in $W K(i \circ f f+N+i-1)$ for $i=1, N$. Also, use $\operatorname{IN}(35)$ to specify the size of the work arrays.

IN(34) This flag is used if the user has not allocated storage for the matrix $A$ in the array $W K$. If the value is 0 , storage is allocated. If the value is 1 , storage was not allocated. In this case, the user must be using reverse communication to evaluate the partial derivative matrix and to solve the linear systems $A \Delta y=\Delta g$. Default value is 0 .

IN(35) These two values are the sizes of the arrays IWK and wK allocated in the users program. The values are checked against the program requirements. These checks are made only if the values are positive. Users will normally set this option when directly calling D2SPG. Default values are $(0,0)$.

## Options of type REAL or DOUBLE PRECISION:

$\operatorname{INR}(1) \quad$ The value of the independent variable, $t$. Default value is AMACH(6).

INR(2) The farthest working $t$ point the integration has reached. Default value is AMACH(6) .

INR(3) The current value of TOUT. Default value is AMACH(6).
INR(4) The next special point, TSTOP, before reaching TOUT. Default value is AMACH(6). Used with option IN(4).

INR(5) The pair of scalar values ATOL and RTOL that apply to the error estimates of all components of $y$. Default values for both are $\operatorname{SQRT}(\operatorname{AMACH}(4))$.

INR(6) The initial step size if DASPG is not to compute it internally. Default value is AMACH(6).

INR(7) The maximum step size allowed. Default value is AMACH(2).
$\operatorname{INR}(8)$ This value is the reciprocal of the condition number of the matrix $A$. It is defined when forward communication is used to solve for the linear updates to the BDF corrector equation. No further program action, such as declaring a singular system, based on the condition number. Users can declare the system to be singular by raising the "panic flag" using option IN(28). Default value is АМАСН(6).

INR(9) The value of $c_{j}$ used in the partial derivative matrix for reverse communication evaluation. Default value is АМАСН(6).

INR(10) The step size to be attempted on the next move. Default value is AMACH(6).

INR(11) The step size taken on the previous move. Default value is AMACH(6).
4. Norm Function Subprogram

The routine DASPG uses a weighted Euclidean-RMS norm to measure the size of the estimated error in each step. This is done using a FUNCTION subprogram: REAL FUNCTION D10PG ( $\mathrm{N}, \mathrm{V}, \mathrm{WT}$ ). This routine returns the value of the RMS weighted norm given by:

$$
\mathrm{D} 10 \mathrm{PG}=\sqrt{N^{-1} \sum_{i=1}^{N}\left(v_{i} / w t_{i}\right)^{2}}
$$

Users can replace this function with one of their own choice. This should be done only for problem-related reasons.

## Example 1

The Van der Pol equation $u^{\prime \prime}+\mu\left(u^{2}-1\right) u^{\prime}+u=0, \mu>0$, is a single ordinary differential equation with a periodic limit cycle. See Hartman (1964, page 181). For the value $\mu=5$, the equations are integrated from $t=0$ until the limit has clearly developed at $t=26$. The (arbitrary) initial conditions used here are $u(0)=2$ and $u^{\prime}(0)=-2 / 3$. Except for these initial conditions and the final $t$ value, this is problem (E2) of the Enright and Pryce (1987) test package. This equation is solved as a differential-algebraic system by defining the first-order system:

$$
\begin{aligned}
& \varepsilon=1 / \mu \\
& y_{1}=u \\
& g_{1}=y_{2}-y_{1}^{\prime}=0 \\
& g_{2}=\left(1-y_{1}^{2}\right) y_{2}-\varepsilon\left(y_{1}+y_{2}^{\prime}\right)=0
\end{aligned}
$$

Note that the initial condition for

$$
y_{2}^{\prime}
$$

in the sample program is not consistent, $g_{2} \neq 0$ at $t=0$. The routine DASPG solves for this starting value. No options need to be changed for this usage. The set of pairs $\left(u\left(t_{j}\right), u^{\prime}\left(t_{j}\right)\right)$ are accumulated for the 260 values $t_{j}=0.1,26,(0.1)$.

```
    USE UMACH_INT
    USE DASPG_INT
    IMPLICIT NONE
    INTEGER N, NP, IDO
    PARAMETER (N=2, NP=260)
    INTEGER ISTEP, NOUT, NSTEP
    REAL DELT, T, TEND, U(NP), UPR(NP), Y(N), YPR(N)
! SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL GCN
    IDO = 1
    T = 0.0
    TEND = 26.0
    DELT = 0.1
    NSTEP = TEND/DELT
    Y(1) = 2.0
    Y(2) = -2.0/3.0
    YPR(1) = Y(2)
    YPR(2) = 0.
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
    ISTEP = 0
    10 CONTINUE
```

!

```
    ISTEP = ISTEP + 1
    CALL DASPG (T, T+DELT, IDO, Y, YPR, GCN)
                                    Save solution for plotting
    IF (ISTEP .LE. NSTEP) THEN
        U(ISTEP) = Y(1)
        UPR(ISTEP) = YPR(1)
            IF (ISTEP .EQ. NSTEP) IDO = 3
            GO TO 10
    END IF
    WRITE (NOUT,99999) TEND, Y, YPR
99998 FORMAT (11X, 'T', 14X, 'Y(1)', 11X, 'Y(2)', 10X, 'Y''(1)', 10X, &
            'Y''(2)')
99999 FORMAT (5F15.5)
! Start plotting
CALL SCATR (NSTEP, U, UPR)
    CALL EFSPLT (0, '')
        END
!
    SUBROUTINE GCN (N, T, Y, YPR, GVAL)
    INTEGER N
        REAL T, Y(N), YPR(N), GVAL(N)
        REAL EPS
!
        EPS = 0.2
!
    GVAL(1) = Y(2) - YPR(1)
    GVAL(2) = (1.0-Y(1)**2)*Y(2) - EPS*(Y(1) +YPR(2))
        RETURN
        END
```

    Output
    | $T$ | $Y(1)$ | $Y(2)$ | $Y^{\prime}(1)$ | $Y^{\prime}(2)$ |
| :---: | ---: | ---: | ---: | ---: |
| 26.00000 | 1.45330 | -0.24486 | -0.24713 | -0.09399 |



Figure 5-1 Van der Pol Cycle, $\left(\mathrm{u}(\mathrm{t}), \mathrm{u}^{\prime}(\mathrm{t})\right), \mu=5$.

## Additional Examples

## Example 2

The first-order equations of motion of a point-mass $m$ suspended on a massless wire of length $\ell$ under the influence of gravity force, $m g$ and tension value $\lambda$, in Cartesian coordinates, $(p, q)$, are

$$
\begin{aligned}
p^{\prime} & =u \\
q^{\prime} & =v \\
m u^{\prime} & =-p \lambda \\
m v^{\prime} & =-q \lambda-m g \\
p^{2}+q^{2}-\ell^{2} & =0
\end{aligned}
$$

This is a genuine differential-algebraic system. The problem, as stated, has an index number equal to the value 3. Thus, it cannot be solved with DASPG directly. Unfortunately, the fact that the index is greater than 1 must be deduced indirectly. Typically there will be an error processed which states that the (BDF) corrector equation did not converge. The user then differentiates and replaces the constraint equation. This example is transformed to a problem of index number of value 1 by differentiating the last equation twice. This resulting equation, which replaces the given equation, is the total energy balance:

$$
m\left(u^{2}+v^{2}\right)-m g q-\ell^{2} \lambda=0
$$

With initial conditions and systematic definitions of the dependent variables, the system becomes:

$$
\begin{aligned}
p(0) & =\ell, q(0)=u(0)=v(0)=\lambda(0)=0 \\
y_{1} & =p \\
y_{2} & =q \\
y_{3} & =u \\
y_{4} & =v \\
y_{5} & =\lambda \\
g_{1} & =y_{3}-y_{1}^{\prime}=0 \\
g_{2} & =y_{4}-y_{2}^{\prime}=0 \\
g_{3} & =-y_{1} y_{5}-m y_{3}^{\prime}=0 \\
g_{4} & =-y_{2} y_{5}-m g-m y_{4}^{\prime}=0 \\
g_{5} & =m\left(y_{3}^{2}+y_{4}^{2}\right)-m g y_{2}-\ell^{2} y_{5}=0
\end{aligned}
$$

The problem is given in English measurement units of feet, pounds, and seconds. The wire has length 6.5 ft , and the mass at the end is 98 lb . Usage of the software does not require it, but standard or "SI" units are used in the numerical model. This conversion of units is done as a first step in the user-supplied evaluation routine, GCN. A set of initial conditions, corresponding to the pendulum starting in a horizontal position, are provided as output for the input signal of $n=0$. The maximum magnitude of the tension parameter, $\lambda(t)=y_{5}(t)$, is computed at the output points, $t=0.1, \pi,(0.1)$. This extreme value is converted to English units and printed.

```
USE DASPG INT
USE CUNIT INT
USE DASPG_INT
USE CUNIT INT
USE UMACH_INT
USE CONST_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=5)
INTEGER IDO, ISTEP, NOUT, NSTEP
REAL DELT, GVAL (N), MAXLB, MAXTEN, T, TEND, TMAX, Y(N), &
            YPR(N)
                                    SPECIFICATIONS FOR INTRINSICS
INTRINSIC ABS
REAL ABS
SPECIFICATIONS FOR SUBROUTINES
SPECIFICATIONS FOR FUNCTIONS
Define initial data
IDO = 1
T = 0.0
```

```
    TEND = CONST('pi')
    DELT = 0.1
    NSTEP = TEND/DELT
    CALL UMACH (2, NOUT)
!
    CALL GCN (0, T, Y, YPR, GVAL)
    ISTEP = 0
    MAXTEN = 0.
    10 CONTINUE
    ISTEP = ISTEP + 1
    CALL DASPG (T, T+DELT, IDO, Y, YPR, GCN)
    IF (ISTEP .LE. NSTEP) THEN
                Note max tension value
        IF (ABS (Y(5)) .GT. ABS (MAXTEN)) THEN
                TMAX = T
                MAXTEN = Y(5)
        END IF
        IF (ISTEP .EQ. NSTEP) IDO = 3
        GO TO 10
    END IF
    CALL CUNIT (MAXTEN, 'kg/s**2', MAXLB, 'lb/s**2')
                                    Print maximum tension
    WRITE (NOUT,99999) MAXLB, TMAX
99999 FORMAT (' Extreme string tension of', F10.2, ' (lb/s**2)', &
            ' occurred at ', 'time ', F10.2)
    END
!
    SUBROUTINE GCN (N, T, Y, YPR, GVAL)
    USE CUNIT INT
    USE CONST_INT
! - SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
    REAL T, Y(*), YPR(*), GVAL(*)
                            SPECIFICATIONS FOR LOCAL VARIABLES
    REAL FEETL, GRAV, LENSQ, MASSKG, MASSLB, METERL, MG
                                    SPECIFICATIONS FOR SAVE VARIABLES
    LOGICAL FIRST
    SAVE FIRST
    SPECIFICATIONS FOR SUBROUTINES
    SPECIFICATIONS FOR FUNCTIONS
    DATA FIRST/.TRUE./
    IF (FIRST) GO TO 20
    10 CONTINUE
!
    IF (N .EQ. O) THEN
                                    The pendulum is horizontal
                                    with these initial y values
    Y(1) = METERL
    Y(2) = 0.
    Y(3) = 0.
    Y(4) = 0.
    Y(5) = 0.
```

```
        YPR(1) = 0.
        YPR(2) = 0.
        YPR(3) = 0.
        YPR(4)=0.
        YPR(5) = 0.
        RETURN
        END IF
!
    GVAL(1) = Y(3) - YPR(1)
    GVAL(2) = Y(4) - YPR(2)
    GVAL(3) = -Y(1)*Y(5) - MASSKG*YPR(3)
    GVAL (4) = -Y(2)*Y(5) - MASSKG*YPR(4) - MG
    GVAL(5) = MASSKG*(Y(3)**2+Y(4)**2) - MG*Y(2) - LENSQ*Y(5)
    RETURN
! Convert from English to
20 CONTINUE
    FEETL = 6.5
    MASSLB = 98.0
! Change to meters
    CALL CUNIT (FEETL, 'ft', METERL, 'meter')
        Change to kilograms
    CALL CUNIT (MASSLB, 'lb', MASSKG, 'kg')
    Get standard gravity
    GRAV = CONST('StandardGravity')
    MG = MASSKG*GRAV
    LENSQ = METERL**2
    FIRST = .FALSE.
    GO TO 10
    END
```


## Output

Extreme string tension of 1457.24 (lb/s**2) occurred at time 2.50

## Example 3

In this example, we solve a stiff ordinary differential equation (E5) from the test package of Enright and Pryce (1987). The problem is nonlinear with nonreal eigenvalues. It is included as an example because it is a stiff problem, and its partial derivatives are provided in the usersupplied routine with the fixed name DJSPG. Users who require a variable routine name for partial derivatives can use the routine D2SPG. Providing explicit formulas for partial derivatives is an important consideration for problems where evaluations of the function $g\left(t, y, y^{\prime}\right)$ are expensive. Signaling that a derivative matrix is provided requires a call to the Chapter 11 options manager utility, IUMAG. In addition, an initial integration step-size is given for this test problem. A signal for this is passed using the options manager routine IUMAG. The error tolerance is changed from the defaults to a pure absolute tolerance of $0.1 * \operatorname{SQRT}(\operatorname{AMACH}(4))$. Also see IUMAG, and SUMAG/DUMAG in Chapter 11, Utilities, for further details about the options manager routines.

USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER N

```
    PARAMETER (N=4)
    ICHAP, IGET, INUM, IPUT, IRNUM
    PARAMETER (ICHAP=5, IGET=1, INUM=6, IPUT=2, IRNUM=7)
                            SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER IDO, IN(50), INR(20), IOPT(2), IVAL(2), NOUT
    REAL C0, PREC, SVAL(3), T, TEND, Y(N), YPR(N)
        SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL GCN
    IDO = 1
    T = 0.0
    TEND = 1000.0
    CO = 1.76E-3
    Y(1) = C0
    Y(2) = 0.
    Y(3) = 0.
    Y(4) = 0.
! Initial derivatives
    YPR(1) = 0.
    YPR(2) = 0.
    YPR(3) = 0.
    YPR(4) = 0.
    IOPT(1) = INUM
    CALL IUMAG ('math', ICHAP, IGET, 1, IOPT, IN)
    IOPT(1) = IRNUM
    CALL IUMAG ('math', ICHAP, IGET, 1, IOPT, INR)
    Provide initial step
    IOPT(1) = INR(6)
    SVAL(1) = 5.0E-5
    IOPT(2) = INR(5)
    PREC = AMACH(4)
    SVAL(2) = 0.1*SQRT(PREC)
    SVAL(3) = 0.0
    CALL UMAG ('math', ICHAP, IPUT, IOPT, SVAL)
    IOPT(1) = IN(5)
    IVAL(1) = 1
    IOPT(2) = IN(8)
    IVAL(2) = 1
    CALL IUMAG ('math', ICHAP, IPUT, 2, IOPT, IVAL)
    Write title
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
    CALL DASPG (T, TEND, IDO, Y, YPR, GCN)
    WRITE (NOUT,99999) T, Y, YPR
    Reset floating options
    to defaults
```

```
IOPT(1) = -INR(5)
IOPT(2) = -INR(6)
CALL UMAG ('math', ICHAP, IPUT, IOPT, SVAL)
                                    Reset integer options
                                    to defaults
IOPT(1) = -IN(5)
IOPT(2) = -IN(8)
CALL IUMAG ('math', ICHAP, IPUT, 2, IOPT, IVAL)
99998 FORMAT (11X, 'T', 14X, 'Y followed by Y''')
99999 FORMAT (F15.5/(4F15.5))
END
SUBROUTINE GCN (N, T, Y, YPR, GVAL)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER N
REAL T, Y(N), YPR(N), GVAL (N)
                        SPECIFICATIONS FOR LOCAL VARIABLES
REAL C1, C2, C3, C4
C1 = 7.89E-10
C2 = 1.1E7
C3 = 1.13E9
C4 = 1.13E3
GVAL(1) = -C1*Y(1) - C2*Y(1)*Y(3) - YPR(1)
GVAL(2) = C1*Y(1) - C3*Y(2)*Y(3) - YPR(2)
GVAL(3) = C1*Y(1) - C2*Y(1)*Y(3) + C4*Y(4) - C3*Y(2)*Y(3) - &
YPR(3)
GVAL(4) = C2*Y(1)*Y(3) - C4*Y(4) - YPR(4)
RETURN
END
SUBROUTINE DJSPG (N, T, Y, YPR, CJ, PDG, LDPDG)
                        SPECIFICATIONS FOR ARGUMENTS
INTEGER N, LDPDG
REAL T, CJ, Y(N), YPR(N), PDG(LDPDG,N)
    SPECIFICATIONS FOR LOCAL VARIABLES
REAL C1, C2, C3, C4
C1 = 7.89E-10
C2 = 1.1E7
C3 = 1.13E9
C4 = 1.13E3
PDG(1,1) = -C1 - C2*Y(3) - CJ
PDG(1,3) = -C2*Y(1)
PDG (2,1) = C1
PDG (2,2) = -C3*Y(3) - CJ
PDG (2,3) = - C3*Y(2)
PDG(3,1) = C1 - C2*Y(3)
PDG (3,2) = -C3*Y(3)
PDG(3,3) = -C2*Y(1) - C3*Y(2) - CJ
PDG (3,4) = C4
```

```
PDG(4,1) = C2*Y(3)
PDG(4,3) = C2*Y(1)
PDG(4,4) = -C4 - CJ
RETURN
END
```


## Output

## T

1000.00000

| 0.00162 | 0.00000 | 0.00000 | 0.00000 |
| :--- | :--- | :--- | :--- |

$\begin{array}{llll}0.00000 & 0.00000 & 0.00000 & 0.00000\end{array}$

## Example 4

In this final example, we compute the solution of $n=10$ ordinary differential equations, $g=H y-y^{\prime}$, where $y(0)=y_{0}=(1,1, \ldots, 1)^{T}$. The value

$$
\sum_{i=1}^{n} y_{i}(t)
$$

is evaluated at $t=1$. The constant matrix $H$ has entries $h_{i, j}=\min (j-i, 0)$ so it is lower Hessenberg. We use reverse communication for the evaluation of the following intermediate quantities:

1. The function $g$,
2. The partial derivative matrix $A=\partial g / \partial y+c_{j} \partial g / \partial y^{\prime}=H-c_{j} I$,
3. The solution of the linear system $A \Delta y=\Delta g$.

In addition to the use of reverse communication, we evaluate the partial derivatives using formulas. No storage is allocated in the floating-point work array for the matrix. Instead, the matrix $A$ is stored in an array A within the main program unit. Signals for this organization are passed using the routine IUMAG (Chapter 11, Utilities).
An algorithm appropriate for this matrix, Givens transformations applied from the right side, is used to factor the matrix $A$. The rotations are reconstructed during the solve step. See SROTG (Chapter 9, Basic Matrix/Vector Operations) for the formulas.

The routine $\operatorname{D} 2 \mathrm{SPG}$ stores the value of $c_{j}$. We get it with a call to the options manager routine SUMAG (Chapter 11, Utilities). A pointer, or offset into the work array, is obtained as an integer option. This gives the location of $g$ and $\Delta g$. The solution vector $\Delta y$ replaces $\Delta g$ at that location. Caution: If a user writes code wherein $g$ is computed with reverse communication and partials are evaluated with divided differences, then there will be two distinct places where $g$ is to be stored. This example shows a correct place to get this offset.

This example also serves as a prototype for large, structured (possibly nonlinear) DAE problems where the user must use special methods to store and factor the matrix $A$ and solve the linear system $A \Delta y=\Delta g$. The word "factor" is used literally here. A user could, for instance, solve the system using an iterative method. Generally, the factor step can be any preparatory phase required for a later solve step.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=10)
    SPECIFICATIONS FOR PARAMETERS
    INTEGER ICHAP, IGET, INUM, IPUT, IRNUM
    PARAMETER (ICHAP=5, IGET=1, INUM=6, IPUT=2, IRNUM=7)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER I, IDO, IN(50), INR(20), IOPT(6), IVAL(7), IWK(35+N), &
        J, NOUT
    REAL A(N,N), GVAL(N), H(N,N), SC, SS, SUMY, SVAL(1), T, &
        TEND, WK(41+11*N), Y(N), YPR(N), Z
    SPECIFICATIONS FOR INTRINSICS
    INTRINSIC ABS, SQRT
    REAL ABS, SQRT
    SPECIFICATIONS FOR SUBROUTINES
    SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL DGSPG, DJSPG
    Define initial data
    IDO = 1
    T = 0.0E0
    TEND = 1.OEO
! Initial values
    CALL SSET (N, 1.0E0, Y, 1)
    CALL SSET (N, 0.0, YPR, 1)
    Initial lower Hessenberg matrix
    CALL SSET (N*N, O.OEO, H, 1)
    DO 20 I=1, N - 1
        DO 10 J=1, I + 1
            H(I,J) = J - I
        CONTINUE
    CONTINUE
    DO 30 J=1, N
        H(N,J) = J - N
    CONTINUE
    IOPT(1) = INUM
    CALL IUMAG ('math', ICHAP, IGET, 1, IOPT, IN)
    Get floating point option numbers
    IOPT(1) = IRNUM
    CALL IUMAG ('math', ICHAP, IGET, 1, IOPT, INR)
    Set for reverse communication
    evaluation of g.
    IOPT(1) = IN(26)
    IVAL(1) = 0
    Set for evaluation of partial
    derivatives.
    IOPT(2) = IN(5)
    IVAL(2) = 1
    IOPT(3) = IN(29)
    IVAL(3) = 0
    Set for reverse communication
```

```
! solution of linear equations.
    IOPT(4) = IN(31)
    IVAL(4) = 0
    IOPT(5) = IN(34)
    IVAL(5) = 1
! Set the sizes of IWK, WK
    for internal checking.
    IOPT(6) = IN(35)
    IVAL(6) = 35 + N
    IVAL(7) = 41 + 11*N
! 'Put' integer options.
    CALL IUMAG ('math', ICHAP, IPUT, 6, IOPT, IVAL)
    Write problem title.
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
! Integrate ODE/DAE. Use
! dummy IMSL external names.
    4 0 ~ C O N T I N U E ~
    CALL D2SPG (N, T, TEND, IDO, Y, YPR, DGSPG, DJSPG, IWK, WK)
    Find where g goes.
    (It only goes in one place
    here, but can vary if
    divided differences are used
    for partial derivatives.)
    IOPT(1) = IN(27)
    CALL IUMAG ('math', ICHAP, IGET, 1, IOPT, IVAL)
    Direct user response.
    GO TO (50, 180, 60, 50, 90, 100, 130, 150), IDO
    5 0 ~ C O N T I N U E
    WRITE (NOUT,*) ' Unexpected return with IDO = ', IDO
    6 0 ~ C O N T I N U E ~
        DO 70 I=1, 50
        IN(I) = -IN(I)
    7 0 ~ C O N T I N U E ~
        CALL IUMAG ('math', ICHAP, IPUT, 50, IN, IVAL)
        DO 80 I=1, 20
            INR(I) = -INR(I)
    8 0 ~ C O N T I N U E ~
        CALL UMAG ('math', ICHAP, IPUT, INR, SVAL, numopts=1)
        STOP
    90 CONTINUE
        CALL SCOPY (N, YPR, 1, GVAL, 1)
        CALL SGEMV ('NO', N, N, 1.0E0, H, N, Y, 1, -1.0E0, GVAL, 1)
                            Put g into place.
        CALL SCOPY (N, GVAL, 1, WK(IVAL(1:)), 1)
        GO TO 40
    100 CONTINUE
! Return came for partial
    1 1 0 ~ C A L L ~ S C O P Y ~ ( N * N , ~ H , ~ 1 , ~ A , ~ 1 )
```

```
! Get value of c_j for partials.
    IOPT(1) = INR(9)
    CALL UMAG ('math', ICHAP, IGET, IOPT, SVAL, numopts=1)
    Subtract c_j from diagonals
    to compute (partials for y')*c_j.
    DO 120 I=1, N
        A(I,I) = A(I,I) - SVAL(1)
    120 CONTINUE
    GO TO 40
    130 CONTINUE
    DO 140 J=1, N - 1
! Construct and apply Givens
        CALL SROTG (A (J,J), A (J,J+1), SC, SS)
        CALL SROT (N-J, A((J+1):,1), 1, A((J+1):,J+1), 1, SC, SS)
    140 CONTINUE
    GO TO 40
    150 CONTINUE
! Return came to solve the system
    CALL SCOPY (N, WK(IVAL(1)), 1, GVAL, 1)
    DO 160 J=1, N - 1
        GVAL (J) = GVAL (J)/A (J,J)
        CALL SAXPY (N-J, -GVAL(J), A((J+1):,J), 1, GVAL((J+1):, 1))
    160 CONTINUE
    GVAL (N) = GVAL (N)/A (N,N)
!
    DO 170 J=N - 1, 1, -1
        Z = A(J,J+1)
        IF (ABS(Z) .LT. 1.OE0) THEN
            SC = SQRT (1.0E0-Z**2)
                SS = Z
        ELSE IF (ABS(Z) .GT. 1.OE0) THEN
                SC = 1.0E0/Z
                SS = SQRT(1.0E0-SC**2)
        ELSE
                SC = 0.0E0
                SS = 1.0E0
        END IF
        CALL SROT (1, GVAL(J:), 1, GVAL((J+1):), 1, SC, SS)
    170 CONTINUE
    CALL SCOPY (N, GVAL, 1, WK(IVAL(1)), 1)
    GO TO 40
!
    180 CONTINUE
    SUMY = 0.EO
    DO 190 I=1, N
        SUMY = SUMY + Y(I)
    190 CONTINUE
    WRITE (NOUT,99999) TEND, SUMY
! Finish up internally
    IDO = 3
    GO TO 40
99998 FORMAT (11X, 'T', 6X, 'Sum of Y(i), i=1,n')
99999 FORMAT (2F15.5)
```


## Output

T Sum of $Y(i), i=1, n$
$1.00000 \quad 65.17058$

## Introduction to Subroutine PDE_1D_MG

The section describes an algorithm and a corresponding integrator subroutine PDE_1D_MG for solving a system of partial differential equations

$$
u_{t} \equiv \frac{\partial u}{\partial t}=f(u, x, t), \quad x_{L}<x<x_{R}, t>t_{0}
$$

## Equation 1

This software is a one-dimensional solver. It requires initial and boundary conditions in addition to values of $u_{t}$. The integration method is noteworthy due to the maintenance of grid lines in the space variable, $x$. Details for choosing new grid lines are given in Blom and Zegeling, (1994). The class of problems solved with $\operatorname{PDE} \_1 D_{\text {_ }} M G$ is expressed by equations:

$$
\begin{aligned}
& \sum_{k=1}^{N P D E} C_{j, k}\left(x, t, u, u_{x}\right) \frac{\partial u^{k}}{\partial t}=x^{-m} \frac{\partial}{\partial x}\left(x^{m} R_{j}\left(x, t, u, u_{x}\right)\right)-Q_{j}\left(x, t, u, u_{x}\right), \\
& j=1, \ldots, N P D E, \quad x_{L}<x<x_{R,} \quad t>t_{0}, \quad m \in\{0,1,2\}
\end{aligned}
$$

## Equation 2

The vector

$$
u \equiv\left[u^{\prime} \ldots, u^{N P D E}\right]^{T}
$$

is the solution. The integer value NPDE $\geq 1$ is the number of differential equations. The functions $R_{j}$ and $Q_{j}$ can be regarded, in special cases, as flux and source terms. The functions

$$
u, C_{j, k}, R_{j} \text { and } Q_{j}
$$

are expected to be continuous. Allowed values

$$
m=0, m-1, \text { and } m=2
$$

are for problems in Cartesian, cylindrical or polar, and spherical coordinates. In the two cases $m>0$, the interval

$$
\left[x_{L}, x_{R}\right]
$$

must not contain $x=0$ as an interior point.
The boundary conditions have the master equation form

$$
\begin{aligned}
& \beta_{j}(x, t) R_{j}\left(x, t, u, u_{x}\right)=\gamma_{j}\left(x, t, u, u_{x}\right), \\
& \text { at } x=x_{L} \text { and } x=x_{R}, j=1, \ldots, N P D E
\end{aligned}
$$

## Equation 3

In the boundary conditions the

$$
\beta_{j} \text { and } \gamma_{j}
$$

are continuous functions of their arguments. In the two cases $m>0$ and an endpoint occurs at 0 , the finite value of the solution at $x=0$ must be ensured. This requires the specification of the solution at $x=0$, or implies that

$$
\left.R_{j}\right|_{x=x_{L}}=0
$$

or

$$
\left.R_{j}\right|_{x=x_{R}}=0
$$

The initial values satisfy

$$
u\left(x, t_{0}\right)=u_{0}(x), \quad x \in\left[x_{L}, x_{R}\right]
$$

where $u_{0}$ is a piece-wise continuous vector function of $x$ with $N P D E$ components.
The user must pose the problem so that mathematical definitions are known for the functions

$$
C_{k, j}, R_{j}, Q_{j}, \beta_{j}, \gamma_{j} \text { and } u_{0}
$$

These functions are provided to the routine PDE_1D_MG in the form of three subroutines. Optionally, this information can be provided by reverse communication. These forms of the interface are explained below and illustrated with examples. Users may turn directly to the examples if they are comfortable with the description of the algorithm.

## PDE_1D_MG

Invokes a module, with the statement USE PDE_1D_MG, near the second line of the program unit. The integrator is provided with single or double precision arithmetic, and a generic named interface is provided. We do not recommend using 32-bit floating point arithmetic here. The routine is called within the following loop, and is entered with each value of IDO. The loop continues until a value of IDO results in an exit.

```
IDO=1
DO
    CASE(IDO == 1) {Do required initialization steps}
    CASE(IDO == 2) {Save solution, update TO and TOUT }
            IF{Finished with integration} IDO=3
    CASE(IDO == 3) EXIT {Normal}
    CASE(IDO == 4) EXIT {Due to errors}
    CASE(IDO == 5) {Evaluate initial data}
    CASE(IDO == 6) {Evaluate differential equations}
    CASE(IDO == 7) {Evaluate boundary conditions}
    CASE (IDO == 8) {Prepare to solve banded system}
    CASE(IDO == 9) {Solve banded system}
```

```
CALL PDE 1D MG (TO,TOUT,IDO,U, &
initial_\overline{conditions,&}
pde_system_definition,&
boundary_conditions, IOPT)
END DO
```

The arguments to PDE_1D_MG are required or optional.

## Required Arguments

TO-(Input/Output)
This is the value of the independent variable $t$ where the integration of $u_{t}$ begins. It is set to the value tout on return.

TOUT-(Input)
This is the value of the independent variable $t$ where the integration of $u_{t}$ ends. Note: Values of TO < TOUT imply integration in the forward direction, while values of TO > TOUT imply integration in the backward direction. Either direction is permitted.

IDO-(Input/Output)
This in an integer flag that directs program control and user action. Its value is used for initialization, termination, and for directing user response during reverse communication:

IDO =1 This value is assigned by the user for the start of a new problem. Internally it causes allocated storage to be reallocated, conforming to the problem size. Various initialization steps are performed.

IDO $=\mathbf{2}$ This value is assigned by the routine when the integrator has successfully reached the end point, тоUT.

IDO $=\mathbf{3}$ This value is assigned by the user at the end of a problem. The routine is called by the user with this value. Internally it causes termination steps to be performed.

IDO=4 This value is assigned by the integrator when a type FATAL or TERMINAL error condition has occurred, and error processing is set NOT to STOP for these types of errors. It is not necessary to make a final call to the integrator with IDO $=\mathbf{3}$ in this case.

Values of IDO $=\mathbf{5 , 6}, \mathbf{7 , 8}, 9$ are reserved for applications that provide problem information or linear algebra computations using reverse communication. When problem information is provided using reverse communication, the differential equations, boundary conditions and initial data must all be given. The absence of optional subroutine names in the calling sequence directs the routine to use reverse communication. In the module PDE_1D_MG_INT, scalars and arrays for evaluating results are named below. The names are preceded by the prefix "s_pde_1d_mg_" or "d_pde_1d_mg_", depending on the precision. We use the prefix "?_pde_1d_mg_", for the appropriate choice.

IDO=5 This value is assigned by the integrator, requesting data for the initial conditions. Following this evaluation the integrator is re-entered.
(Optional) Update the grid of values in array locations $U(N P D E+1, j) j=2, \ldots, N$. This grid is returned to the user equally spaced, but can be updated as desired, provided the values are increasing.
(Required) Provide initial values for all components of the system at the grid of values $U(N P D E+1, j) j=1, \ldots, N$. If the optional step of updating the initial grid is performed, then the initial values are evaluated at the updated grid.

IDO $=6$ This value is assigned by the integrator, requesting data for the differential equations. Following this evaluation the integrator is re-entered. Evaluate the terms of the system of Equation 2. A default value of $m=0$ is assumed, but this can be changed to one of the other choices $m=1$ or 2 . Use the optional argument IOPT (: ) for that purpose. Put the values in the arrays as indicated ${ }^{1}$.

$$
\begin{aligned}
& x \equiv ? p d e_{-} 1 d_{-} m g_{-} x \\
& t \equiv ? p d e_{-} 1 d \_m g_{-} t \\
& u^{j} \equiv ? p d e_{-} 1 d \_m g_{-} u(j) \\
& \frac{\partial u^{j}}{\partial x}=u_{x}^{j} \equiv ? p d e_{-} 1 d \_m g_{-} d u d x(j) \\
& ? \quad p d e_{-} 1 d \_m g \_c(j, k):=C_{j, k}\left(x, t, u, u_{x}\right) \\
& ? . p d e_{-} 1 d \_m g_{-} r(j):=r_{j}\left(x, t, u, u_{x}\right) \\
& ? p d e \_1 d \_m g_{-} q(j):=q_{j}\left(x, t, u, u_{x}\right) \\
& j, k=1, \ldots, N P D E
\end{aligned}
$$

If any of the functions cannot be evaluated, set pde_1d_mg_ires=3. Otherwise do not change its value.

IDO=7 This value is assigned by the integrator, requesting data for the boundary conditions, as expressed in Equation 3. Following the evaluation the integrator is reentered.

[^0]\[

$$
\begin{aligned}
& x \equiv ? p d e_{-} 1 d_{-} m g_{-} x \\
& t \equiv ? p d \__{-} 1 d \_m g_{-} t \\
& u^{j} \equiv ? p d e_{-} 1 d \_m g_{-} u(j) \\
& \frac{\partial u^{j}}{\partial x}=u_{x}^{j} \equiv ? p d e_{-} 1 d \_m g_{-} d u d x(j) \\
& ? \quad p d e \_1 d \_m g_{-} b e t a(j):=\beta_{j}\left(x, t, u, u_{x}\right) \\
& ? \quad p d e \_1 d \_m g \_g a m m a(j):=\gamma_{j}\left(x, t, u, u_{x}\right) \\
& j=1, \ldots, N P D E
\end{aligned}
$$
\]

The value $x \in\left\{x_{L}, x_{R}\right\}$, and the logical flag pde_1d_mg_LEFT=. TRUE. for $x=x_{L}$. It has the value pde_1d_mg_LEFT=. FALSE. for $x=x_{R}$. If any of the functions cannot be evaluated, set pde_1d_mg_ires=3. Otherwise do not change its value.

IDO=8 This value is assigned by the integrator, requesting the calling program to prepare for solving a banded linear system of algebraic equations. This value will occur only when the option for "reverse communication solving" is set in the array IOPT (: ), with option PDE_1D_MG_REV_COMM_FACTOR_SOLVE. The matrix data for this system is in Band Storage $\bar{M}$ ode, described $\overline{\text { in }}$ the section: Reference Material for the IMSL Fortran Numerical Libraries.

| PDE_1D_MG_IBAND | Half band-width of linear system |
| :--- | :--- |
| PDE_1D_MG_LDA | The value $3 * P D E \_1 D_{-}$MG_IBAND+1, with <br> $N E Q=(N P D E+1) N$ |
| ?_PDE_1D_MG_A | Array of size PDE_1D_MG_LDA by NEQ <br> holding the problem matrix in Band Storage <br> Mode |
| PDE_1D_MG_PANIC_FLAG | Integer set to a non-zero value only if the linear <br> system is detected as singular |

IDO $=\mathbf{9}$ This value is assigned by the integrator, requesting the calling program to solve a linear system with the matrix defined as noted with IDO $=\mathbf{8}$.

| $?_{-}$PDE_1D_MG_RHS | Array of size NEQ holding the linear <br> system problem right-hand side |
| :--- | :--- |
| PDE_1D_MG_PANIC_FLAG | Integer set to a non-zero value only if the <br> linear system is singular |
| ?_PDE_1D_MG_SOL | Array of size NEQ to receive the solution, <br> after the solving step |

U (1:NPDE $+1,1: N$ - (Input/Output)
This assumed-shape array specifies Input information about the problem size and boundaries. The dimension of the problem is obtained from $N P D E+1=\operatorname{size}(U, 1)$. The number of grid points is obtained by $N=\operatorname{size}(U, 2)$. Limits for the variable $x$ are assigned as input in array locations, $U(N P D E+1,1)=x_{L}, U(N P D E+1, N)=x_{R}$. It is not required to define $U(N P D E+1, j), j=2, \ldots, N-1$. At completion, the array $\mathrm{U}(1: \mathrm{NPDE}, 1: \mathrm{N})$ contains the approximate solution value $U_{i}\left(x_{j}(T O U T), T O U T\right)$ in location $\mathrm{U}(\mathrm{I}, \mathrm{J})$. The grid value $x_{j}(T O U T)$ is in location $\mathrm{U}(\mathrm{NPDE}+1, \mathrm{~J})$. Normally the grid values are equally spaced as the integration starts. Variable spaced grid values can be provided by defining them as Output from the subroutine initial_conditions or during reverse communication, IDO $=5$.

## Optional Arguments

```
initial_conditions-(Input)
```

The name of an external subroutine, written by the user, when using forward communication. If this argument is not used, then reverse communication is used to provide the problem information. The routine gives the initial values for the system at the starting independent variable value T 0 . This routine can also provide a nonuniform grid at the initial value.

```
SUBROUTINE initial_conditions (NPDE,N,U)
    Integer NPDE,N
    REAL(kind(TO)) U(:,)
END SUBROUTINE
```

(Optional) Update the grid of values in array locations $U(N P D E+1, j), j=2, \ldots, N-1$. This grid is input equally spaced, but can be updated as desired, provided the values are increasing.
(Required) Provide initial values $U(:, j), j=1, \ldots, N$ for all components of the system at the grid of values $U(N P D E+1, j), j=1, \ldots, N$. If the optional step of updating the initial grid is performed, then the initial values are evaluated at the updated grid.

```
pde_system_definition-(Input)
```

The name of an external subroutine, written by the user, when using forward communication. It gives the differential equation, as expressed in Equation 2.

```
SUBROUTINE pde_system_definition&
    (t, x, NPDE, u, dudx, c, q, r, IRES)
    Integer NPDE, IRES
    REAL(kind(TO)) t, x, u(:), dudx(:)
    REAL(kind(TO)) c(:,:), q(:), r(:)
END SUBROUTINE
```

Evaluate the terms of the system of. A default value of $m=0$ is assumed, but this can be changed to one of the other choices $m=1$ or 2 . Use the optional argument IOPT (: ) for that purpose. Put the values in the arrays as indicated.

$$
\begin{aligned}
& u^{j} \equiv u(j) \\
& \frac{\partial u^{j}}{\partial x}=u_{x}^{j} \equiv d u d x(j) \\
& c(j, k):=C_{j, k}\left(x, t, u, u_{x}\right) \\
& r(j):=r_{j}\left(x, t, u, u_{x}\right) \\
& q(j):=q_{j}\left(x, t, u, u_{x}\right) \\
& j, k=1, \ldots, N P D E
\end{aligned}
$$

If any of the functions cannot be evaluated, set $\operatorname{IRES}=3$. Otherwise do not change its value.
boundary_conditions-(Input)
The name of an external subroutine, written by the user when using forward communication. It gives the boundary conditions, as expressed in Equation 2.

```
SUBROUTINE BOUNDARY_CONDITIONS(T,BETA,GAMMA,U,DUDX,NPDE,LEFT,IRES)
                    real(kind(1d0)),intent(in)::t
                    real(kind(1d0)),intent(out),dimension(:)::BETA, GAMMA
                    real(kind(1d0)),intent(in),dimension(:)::U, DUDX
                    integer,intent(in)::NPDE
                    logical,intent(in)::LEFT
            integer,intent(out)::IRES
        END SUBROUTINE
```

$$
\begin{aligned}
& u^{j} \equiv u(j) \\
& \frac{\partial u^{j}}{\partial x}=u_{x}^{j} \equiv \operatorname{dudx}(j) \\
& \operatorname{beta}(j):=\beta_{j}\left(x, t, u, u_{x}\right) \\
& \operatorname{gamma}(j):=\gamma_{j}\left(x, t, u, u_{x}\right) \\
& j=1, \ldots, N P D E
\end{aligned}
$$

The value $x \in\left\{x_{L}, x_{R}\right\}$, and the logical flag LEFT=.TRUE. for $x=x_{L}$. The flag has the value LEFT=. FALSE. for $x=x_{R}$.

IOPT-(Input)
Derived type array s_options or d_options, used for passing optional data to PDE_1D_MG. See the section Optional Data in the Introduction for an explanation of the derived type and its use. It is necessary to invoke a module, with the statement USE ERROR_OPTION_PACKET, near the second line of the program unit. Examples 2-8 use this optional argument. The choices are as follows:

| Packaged Options for PDE_1D_MG |  |  |  |
| :---: | :--- | :---: | :---: |
| Option Prefix =? |  | Option Name |  |
| S_, d_ | PDE_1D_MG_CART_COORDINATES | Option Value |  |
| S_, d_ | PDE_1D_MG_CYL_COORDINATES | 1 |  |
| S_, d_ | PDE_1D_MG_SPH_COORDINATES | 2 |  |
| S_, d_ | PDE_1D_MG_TIME_SMOOTHING | 3 |  |
| S_, d_ | PDE_1D_MG_SPATIAL_SMOOTHING | 4 |  |
| S_, d_ | PDE_1D_MG_MONITOR_REGULARIZING | 5 |  |
| S_, d_ | PDE_1D_MG_RELATIVE_TOLERANCE | 6 |  |
| S_, d_ | PDE_1D_MG_ABSOLUTE_TOLERANCE | 7 |  |
| S_, d_ | PDE_1D_MG_MAX_BDF_ORDER | 8 |  |
| S_, d_ | PDE_1D_MG_REV_COMM_FACTOR_SOLVE | 9 |  |
| S_, d_ | PDE_1D_MG_NO_NULLIFY_STACK | 10 |  |

IOPT(IO) = PDE_1D_MG_CART_COORDINATES
Use the value $m=0$ in Equation 2. This is the default.
IOPT(IO) = PDE_1D_MG_CYL_COORDINATES
Use the value $m=1$ in Equation 2. The default value is $m=0$.
IOPT(IO) = PDE_1D_MG_SPH_COORDINATES
Use the value $\bar{m}=2 \overline{\text { in }}$ Equation 2. The default value is $m=0$.
IOPT(IO) =
?_OPTIONS (PDE_1D_MG_TIME_SMOOTHING,TAU)
This option resets the value of the parameter $\tau \geq 0$, described above.
The default value is $\tau=0$.
IOPT(IO) =
?_OPTIONS (PDE_1D_MG_SPATIAL_SMOOTHING, KAP)
This option resets the value of the parameter $\kappa \geq 0$, described above.
The default value is $\kappa=2$.

IOPT(IO) =
?_OPTIONS (PDE_1D_MG_MONITOR_REGULARIZING,ALPH)
This option resets the value of the parameter $\alpha \geq 0$, described above.
The default value is $\alpha=0.01$.

IOPT(IO) = ?_OPTIONS
(PDE_1D_MG_RELATIVE_TOLERANCE,RTOL)
This option resets the value of the relative accuracy parameter used in DASPG. The default value is $\mathrm{RTOL}=1 \mathrm{E}-2$ for single precision and RTOL=1D-4 for double precision.

IOPT(IO) = ?_OPTIONS
(PDE_1D_MG_ABSOLUTE_TOLERANCE,ATOL)
This $\overline{\text { option }}$ resets the value of the absolute accuracy parameter used in DASPG. The default value is $\mathrm{ATOL}=1 \mathrm{E}-2$ for single precision and
$A T O L=1 D-4$ for double precision.
$\operatorname{IOPT}(I O)=$ PDE_1D_MG_MAX_BDF_ORDER
$\operatorname{IOPT}(I O+1)=$ MAXBDF
Reset the maximum order for the BDF formulas used in DASPG. The default value is $M A X B D F=2$. The new value can be any integer between 1 and 5 . Some problems will benefit by making this change. We used the default value due to the fact that DASPG may cycle on its selection of order and step-size with orders higher than value 2 .

IOPT (IO) = PDE_1D_MG_REV_COMM_FACTOR_SOLVE
The calling program unit will solve the banded linear systems required in the stiff differential-algebraic equation integrator. Values of $\mathbf{I D O}=\mathbf{8}, \mathbf{9}$ will occur only when this optional value is used.

```
IOPT(IO) = PDE_1D_MG_NO_NULLIFY_STACK
```

To maintain an efficient interface, the routine PDE_1D_MG collapses the subroutine call stack with CALL_E1PSH ("NULLIFY_STACK"). This implies that the overhead of maintaining the stack will be eliminated, which may be important with reverse communication. It does not eliminate error processing. However, precise information of which routines have errors will not be displayed. To see the full call chain, this option should be used. Following completion of the integration, stacking is turned back on with CALL_E1POP ("NULLIFY_STACK").

## FORTRAN 90 Interface

Generic: CALL PDE_1D_MG (TO, TOUT, IDO, [,...])
Specific: The specific interface names are S_PDE_1D_MG and D_PDE_1D_MG.

## Description

The equation

$$
u_{t}=f(u, x, t), x_{L}<x<x_{R}, t>t_{0}
$$

is approximated at $N$ time-dependent grid values

$$
x_{L}=x_{0}<\ldots<x_{i}(t)<x_{i+1}(t)<\ldots<x_{N}=x_{R} .
$$

Using the total differential

$$
\frac{d u}{d t}=u_{t}+u_{x} \frac{d x}{d t}
$$

transforms the differential equation to

$$
\frac{d u}{d t}-u_{x} \frac{d x}{d t}=u_{t}=f(u, x, t) .
$$

Using central divided differences for the factor $u_{x}$ leads to the system of ordinary differential equations in implicit form

$$
\frac{d U_{i}}{d t}-\frac{\left(U_{i+1}-U_{i-1}\right)}{\left(x_{i+1}-x_{i-1}\right)} \frac{d x_{i}}{d t}=F_{i}, \quad t>t_{0}, \quad i=1, \ldots, N .
$$

The terms $U_{i}, F_{i}$ respectively represent the approximate solution to the partial differential equation and the value of $f(u, x, t)$ at the point $(x, t)=\left(x_{i},(t), t\right)$. The truncation error is second-order in the space variable, $x$. The above ordinary differential equations are underdetermined, so additional equations are added for the variation of the time-dependent grid points. It is necessary to discuss these equations, since they contain parameters that can be adjusted by the user. Often it will be necessary to modify these parameters to solve a difficult problem. For this purpose the following quantities are defined ${ }^{2}$ :

$$
\begin{aligned}
& \Delta x_{i}=x_{i+1}-x_{i}, n_{i}=\left(\Delta x_{i}\right)^{-1} \\
& \mu_{i}=n_{i}-\kappa(\kappa+1)\left(n_{i+1}-2 n_{i}+n_{i-1}\right), 0 \leq i \leq N \\
& n_{-1} \equiv n_{0}, n_{N+1} \equiv n_{N}
\end{aligned}
$$

The values $n_{i}$ are the so-called point concentration of the grid, and $\kappa \geq 0$ denotes a spatial smoothing parameter. Now the grid points are defined implicitly so that

$$
\frac{\mu_{i-1}+\tau \frac{d \mu_{i-1}}{d t}}{M_{i-1}}=\frac{\mu_{i}+\tau \frac{d \mu_{1}}{d t}}{M_{i}}, 1 \leq i \leq N
$$

where $\tau \geq 1$ is a time-smoothing parameter. Choosing $\tau$ very large results in a fixed grid.
Increasing the value of $\tau$ from its default avoids the error condition where grid lines cross. The divisors are

$$
M_{i}^{2}=\alpha+N P D E^{-1} \sum_{j=1}^{\text {NDDE }} \frac{\left(U_{i+1}^{j}-U_{i}^{j}\right)^{2}}{\left(\Delta x_{i}\right)^{2}} .
$$

The value $\kappa$ determines the level of clustering or spatial smoothing of the grid points. Decreasing $\kappa$ from its default decrease the amount of spatial smoothing. The parameters $M_{i}$ approximate arc length and help determine the shape of the grid or $x_{i}$-distribution. The parameter $\tau$ prevents the grid movement from adjusting immediately to new values of the $M_{i}$, thereby avoiding oscillations in the grid that cause large relative errors. This is important when applied to solutions with steep gradients.
The discrete form of the differential equation and the smoothing equations are combined to yield the implicit system of differential equations.

[^1]\[

$$
\begin{aligned}
& A(Y) \frac{d Y}{d t}=L(Y), \\
& Y=\left[U_{1}^{1}, \ldots, U_{1}^{\text {NPDE }}, x_{1}, \ldots, U_{j}^{1}, \ldots, U_{j}^{\text {NPDE }} x_{j}, \ldots\right]^{T}
\end{aligned}
$$
\]

This is frequently a stiff differential-algebraic system. It is solved using the integrator DASPG and its subroutines, including D2SPG. These are documented in this chapter. Note that DASPG is restricted to use within PDE_1D_MG until the routine exits with the flag IDO = 3. If DASPG is needed during the evaluations of the differential equations or boundary conditions, use of a second processor and inter-process communication is required. The only options for DASPG set by PDE_1D_MG are the Maximum BDF Order, and the absolute and relative error values, ATOL and RTOL. Users may set other options using the Options Manager. This is described in routine DASPG and generally in Chapter 11 of this manual.

## Remarks on the Examples

Due to its importance and the complexity of its interface, this subroutine is presented with several examples. Many of the program features are exercised. The problems complete without any change to the optional arguments, except where these changes are required to describe or to solve the problem.

In many applications the solution to a PDE is used as an auxiliary variable, perhaps as part of a larger design or simulation process. The truncation error of the approximate solution is commensurate with piece-wise linear interpolation on the grid of values, at each output point. To show that the solution is reasonable, a graphical display is revealing and helpful. We have not provided graphical output as part of our documentation, but users may already have the Visual Numerics, Inc. product, PV-WAVE, not included with Fortran Numerical Library. Examples 1-8 write results in files "PDE_ex0? out" that can be visualized with PV-WAVE. We provide a script of commands, "pde_1d_mg_plot.pro", for viewing the solutions (see example below). The grid of values and each consecutive solution component is displayed in separate plotting windows. The script and data files written by examples 1-8 on a SUN-SPARC system are in the directory for Fortran Numerical Library examples. When inside PV_WAVE, execute the command line "pde_1d_mg_plot,filename=' PDE_ex0?.out" to view the output of a particular example.

## Code for PV-WAVE Plotting (Examples Directory)

```
PRO PDE_1d_mg_plot, FILENAME = filename, PAUSE = pause
;
    if keyword_set(FILENAME) then file = filename else file = "res.dat"
    if keyword_set(PAUSE) then twait = pause else twait \(=.1\)
        Define floating point variables that will be read
        from the first line of the data file.
    \(\mathrm{xl}=0 \mathrm{DO}\)
    \(\mathrm{xr}=0 \mathrm{DO}\)
    t0 \(=0 D 0\)
    tlast \(=0 \mathrm{DO}\)
;
; Open the data file and read in the problem parameters.
    openr, lun, filename, /get_lun
    readf, lun, npde, np, nt, xl, xr, t0, tlast
```

```
; Define the arrays for the solutions and grid.
    u = dblarr(nt, npde, np)
    g = dblarr(nt, np)
    times = dblarr(nt)
    Define a temporary array for reading in the data.
    tmp = dblarr(np)
    t_tmp = 0D0
;
; Read in the data.
    for i = 0, nt-1 do begin ; For each step in time
        readf, lun, t_tmp
        times(i) = t_tmp
        for k = 0, npde-1 do begin ; For each PDE:
            rmf, lun, tmp
            u(i,k,*) = tmp ; Read in the components.
        end
        rmf, lun, tmp
        g(i,*) = tmp ; Read in the grid.
    end
;
; Close the data file and free the unit.
    close, lun
    free_lun, lun
            We now have all of the solutions and grids.
            Delete any window that is currently open.
    while (!d.window NE -1) do WDELETE
            Open two windows for plotting the solutions
            and grid.
    window, 0, xsize = 550, ysize = 420
    window, 1, xsize = 550, ysize = 420
            Plot the grid.
    wset, 0
    plot, [xl, xr], [t0, tlast], /nodata, ystyle = 1, $
        title = "Grid Points", xtitle = "X", ytitle = "Time"
    for i = 0, np-1 do begin
        oplot, g(*, i), times, psym = -1
    end
;
Plot the solution(s):
    wset, 1
    for k = 0, npde-1 do begin
        umin = min(u(*,k,*))
        umax = max(u(*,k,*))
        for i = 0, nt-1 do begin
        title = strcompress("U_"+string(k+1), /remove_all)+ $
            " at time "+string(times(i))
        plot, g(i, *), u(i,k,*), ystyle = 1, $
```

```
                title = title, xtitle = "X", $
                ytitle = strcompress("U_"+string(k+1), /remove_all), $
                xr = [xl, xr], yr = [umīn, umax], $
                psym = -4
                wait, twait
            end
    end
end
```


## Example 1 - Electrodynamics Model

This example is from Blom and Zegeling (1994). The system is

$$
\begin{aligned}
& u_{t}=\varepsilon p u_{x x}-g(u-v) \\
& v_{t}=p v_{x x}+g(u-v),
\end{aligned}
$$

$$
\text { where } g(z)=\exp (\eta z / 3)-\exp (-2 \eta z / 3)
$$

$$
0 \leq x \leq 1,0 \leq t \leq 4
$$

$$
u_{x}=0 \text { and } v=0 \text { at } x=0
$$

$$
u=1 \text { and } v_{x}=0 \text { at } x=1
$$

$$
\varepsilon=0.143, p=0.1743, \eta=17.19
$$

We make the connection between the model problem statement and the example:

$$
\begin{aligned}
& C=I_{2} \\
& m=0, R_{1}=\varepsilon p u_{x}, R_{2}=p v_{x} \\
& Q_{1}=g(u-v), Q_{2}=-Q_{1}
\end{aligned}
$$

The boundary conditions are

$$
\begin{aligned}
& \beta_{1}=1, \beta_{2}=0, \gamma_{1}=0, \gamma_{2}=v, \text { at } x=x_{L}=0 \\
& \beta_{1}=0, \beta_{2}=1, \gamma_{1}=u-1, \gamma_{2}=0, \text { at } x=x_{R}=1
\end{aligned}
$$

## Rationale: Example 1

This is a non-linear problem with sharply changing conditions near $t=0$. The default settings of integration parameters allow the problem to be solved. The use of PDE_1D_MG with forward communication requires three subroutines provided by the user to describe the initial conditions, differential equations, and boundary conditions.

```
    program PDE_EX1
! Electrodynamics}\mathrm{ Model:
    USE PDE_1d_mg_int
        IMPLICI\overline{T}}\mathrm{ NO
        INTEGER, PARAMETER :: NPDE=2, N=51, NFRAMES=5
        INTEGER I, IDO
! Define array space for the solution.
```

```
    real(kind(1d0)) U(NPDE+1,N), T0, TOUT
    real(kind(1d0)) :: ZERO=0D0, ONE=1D0, &
    DELTA T=10D0, TEND=4D0
    EXTERNAL IC_01, PDE_01, BC_01
! Start loop to integrate and write solution values.
    IDO=1
    DO
        SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                    T0=ZERO
                    TOUT=1D-3
                        U(NPDE+1,1)=ZERO;U(NPDE+1,N)=ONE
                    OPEN(FILE='PDE ex01.out',UNIT=7)
                    WRITE(7, "(3I5, 4F10.5)") NPDE, N, NFRAMES,&
                        U(NPDE+1,1), U(NPDE+1,N), T0, TEND
! Update to the next output point.
! Write solution and check for final point.
        CASE (2)
            WRITE(7,"(F10.5)") TOUT
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")U(I, :)
            END DO
            TO=TOUT;TOUT=TOUT*DELTA T
            IF(TO >= TEND) IDO=3
            TOUT=MIN(TOUT, TEND)
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
        END SELECT
! Forward communication is used for the problem data.
        CALL PDE_1D_MG (T0, TOUT, IDO, U, &
            initial_conditions= IC_01,&
            PDE_system_definition=-}\mathrm{ PDE 01,&
            boundary_conditions= BC_01)
        END DO
    END
    SUBROUTINE IC 01 (NPDE, NPTS, U)
! This is the initial data for Example 1.
        IMPLICIT NONE
        INTEGER NPDE, NPTS
        REAL (KIND (1D0)) U (NPDE+1,NPTS)
        U(1, :)=1D0;U(2, :) =0D0
    END SUBROUTINE
    SUBROUTINE PDE_01(T, X, NPDE, U, DUDX, C, Q, R, IRES)
```

```
! This is the differential equation for Example 1.
    IMPLICIT NONE
    INTEGER NPDE, IRES
    REAL(KIND(1DO)) T, X, U(NPDE), DUDX (NPDE) , &
        C(NPDE,NPDE), Q(NPDE), R(NPDE)
    REAL (KIND (1D0)) : : EPS=0.143D0, P=0.1743D0, &
            ETA=17.19D0, Z, TWO=2D0, THREE=3D0
        C=0D0;C (1,1)=1D0;C (2, 2)=1D0
        R=P*DUDX;R(1)=R(1)*EPS
        Z=ETA* (U (1) -U (2))/THREE
        Q(1)=EXP(Z)-EXP (-TWO*Z)
        Q(2)=-Q(1)
    END SUBROUTINE
    SUBROUTINE BC_01(T, BTA, GAMA, U, DUDX, NPDE, LEFT, IRES)
! These are the boundary conditions for Example 1.
    IMPLICIT NONE
    INTEGER NPDE, IRES
    LOGICAL LEFT
    REAL(KIND(1D0)) T, BTA(NPDE), GAMA(NPDE),&
        U(NPDE), DUDX(NPDE)
    IF(LEFT) THEN
        BTA (1) =1D0;BTA (2) =0D0
        GAMA (1) =0D0; GAMA (2)=U (2)
    ELSE
        BTA (1) =0D0;BTA (2)=1D0
        GAMA (1)=U (1) -1D0;GAMA (2)=0D0
    END IF
    END SUBROUTINE
```


## Additional Examples

## Example 2 - Inviscid Flow on a Plate

This example is a first order system from Pennington and Berzins, (1994). The equations are

$$
\begin{aligned}
& u_{t}=-v_{x} \\
& u u_{t}=-v u_{x}+w_{x x} \\
& w=u_{x}, \text { implying that } u u_{t}=-v u_{x}+u_{x x} \\
& u(0, t)=v(0, t)=0, u(\infty, t)=u\left(x_{R}, t\right)=1, t \geq 0 \\
& u(x, 0)=1, v(x, 0)=0, x \geq 0
\end{aligned}
$$

Following elimination of $w$, there remain $N P D E=2$ differential equations. The variable $t$ is not time, but a second space variable. The integration goes from $t=0$ to $t=5$. It is necessary to truncate the variable $x$ at a finite value, say $x_{\max }=x_{R}=25$. In terms of the integrator, the system is defined by letting $m=0$ and

$$
C=\left\{C_{j k}\right\}=\left[\begin{array}{ll}
1 & 0 \\
u & 0
\end{array}\right], R=\left[\begin{array}{l}
-v \\
u_{x}
\end{array}\right], Q=\left[\begin{array}{c}
0 \\
v u_{x}
\end{array}\right]
$$

The boundary conditions are satisfied by

$$
\begin{aligned}
& \beta=0, \gamma=\left[\begin{array}{c}
u-\exp (-20 t) \\
v
\end{array}\right], \text { at } x=x_{L} \\
& \beta=0, \gamma=\left[\begin{array}{c}
u-1 \\
v_{x}
\end{array}\right], \text { at } x=x_{R}
\end{aligned}
$$

We use $N=10+51=61$ grid points and output the solution at steps of $\Delta t=0.1$.

## Rationale: Example 2

This is a non-linear boundary layer problem with sharply changing conditions near $t=0$. The problem statement was modified so that boundary conditions are continuous near $t=0$. Without this change the underlying integration software, DASPG, cannot solve the problem. The continuous blending function $u-\exp (-20 t)$ is arbitrary and artfully chosen. This is a mathematical change to the problem, required because of the stated discontinuity at $t=0$. Reverse communication is used for the problem data. No additional user-written subroutines are required when using reverse communication. We also have chosen 10 of the initial grid points to be concentrated near $x_{L}=0$, anticipating rapid change in the solution near that point. Optional changes are made to use a pure absolute error tolerance and non-zero time-smoothing.

```
    program PDE_1D_MG_EX02
! Inviscid Flow Ove\overline{r a-}Pl\overline{a}te
    USE PDE_1d_mg_int
    USE ERRO
    IMPLICIT NONE
    INTEGER, PARAMETER : : NPDE=2, N1=10, N2=51, N=N1+N2
    INTEGER I, IDO, NFRAMES
! Define array space for the solution.
    real(kind(1d0)) U(NPDE+1,N), T0, TOUT, DX1, DX2, DIF
    real(kind(1d0)) :: ZERO=0D0, ONE=1D0, DELTA_T=1D-1,&
        TEND=5D0, XMAX=25D0
    real(kind(1d0)) :: U0=1D0, U1=0D0, TDELTA=1D-1, TOL=1D-2
    TYPE(D_OPTIONS) IOPT(3)
! Start loop to integrate and record solution values.
    IDO=1
    DO
        SELECT CASE (IDO)
! Define values that determine limits and options.
        CASE (1)
            T0=ZERO
            TOUT=DELTA T
            U(NPDE+1,1)=ZERO; U (NPDE+1,N) = XMAX
            OPEN(FILE='PDE_ex02.out',UNIT=7)
            NFRAMES=NINT((吕END+DELTA_T)/DELTA_T)
            WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES, &
```

```
    U(NPDE+1,1), U(NPDE+1,N), T0, TEND
    DX1=XMAX/N2;DX2=DX1/N1
    IOPT (1)=D OPTIONS (PDE 1D MG RELATIVE TOLERANCE, ZERO)
    IOPT (2)=D_OPTIONS (PDE_1D_MG_ABSOLUTE_TOLERANCE,TOL)
    IOPT (3)=D_OPTIONS(PDE_1D_MG_TIME_SMOOTHING,1D-3)
! Update to the next output point.
! Write solution and check for final point.
    CASE (2)
            T0=TOUT
            IF(TO <= TEND) THEN
                    WRITE(7,"(F10.5)") TOUT
                    DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I,:)
            END DO
            TOUT=MIN(TOUT+DELTA_T,TEND)
            IF(TO == TEND) IDO=3
            END IF
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
        CASE (5)
            U(:NPDE,: )=ZERO;U(1,:)=ONE
            DO I=1,N1
                    U (NPDE+1,I)=(I-1)*DX2
            END DO
            DO I=N1+1,N
                    U (NPDE+1,I)=(I-N1)*DX1
            END DO
            WRITE(7,"(F10.5)")T0
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")U(I,:)
            END DO
! Define differential equations.
        CASE (6)
            D_PDE_1D_MG_C=ZERO
            D_PDE_1D_MG_C (1,1)=ONE
            D_PDE_1D_MG_C (2,1)=D_PDE_1D_MG_U (1)
            D_PDE_1D_MG_R(1)=-D_PDE_1D_MG_U(2)
            D_PDE_1D_MG_R(2)= D_PDE_1D_MG_DUDX(1)
            D PDE 1D MG Q (1)= ZERO
            D_PDE_1D_MG_Q (2)= &
            D_PDE_1D_MG_U (2)*D_PDE_1D_MG_DUDX(1)
! Define boundary con\overline{dit}\overline{i}on\overline{s}.
        CASE (7)
            D_PDE 1D_MG_BETA=ZERO
            I\overline{F}(PD\overline{E}_1\overline{D}_M\overline{G}_LEFT) THEN
```

```
    DIF=EXP(-20D0*D_PDE_1D_MG_T)
! Blend the left boundary value down to zero.
            D_PDE_1D_MG_GAMMA=(/D_PDE_1D_MG_U(1)-DIF,D_PDE_1D_MG_U (2)/)
        ELSE
            D_PDE_1D_MG_GAMMA=(/D_PDE_1D_MG_U(1) -
ONE,D_PDE_1D_MG_DUDDX(2)/)
    END SELECT
! Reverse communication is used for the problem data.
            CALL PDE 1D MG (T0, TOUT, IDO, U, IOPT=IOPT)
        END DO
        end program
```


## Example 3 - Population Dynamics

This example is from Pennington and Berzins (1994). The system is

$$
\begin{aligned}
& u_{t}=-u_{x}-I(t) u, x_{L}=0 \leq x \leq a=x_{R}, t \geq 0 \\
& I(t)=\int_{0}^{a} u(x, t) d x \\
& u(x, 0)=\frac{\exp (-x)}{2-\exp (-a)} \\
& u(0, t)=g\left(\int_{0}^{a} b(x, I(t)) u(x, t) d x, t\right), \text { where } \\
& b(x, y)=\frac{x y \exp (-x)}{(y+1)^{2}}, \text { and } \\
& \mathrm{g}(z, t)= \\
& \frac{4 z(2-2 \exp (-a)+\exp (-t))^{2}}{(1-\exp (-a))(1-(1+2 a) \exp (-2 a)(1-\exp )(-a)+\exp (-t))}
\end{aligned}
$$

This is a notable problem because it involves the unknown $u(x, t)=\frac{\exp (-x)}{1-\exp (-a)+\exp (-t)}$ across the entire domain. The software can solve the problem by introducing two dependent algebraic equations:

$$
\begin{aligned}
& v_{1}(t)=\int_{0}^{a} u(x, t) d x, \\
& v_{2}(t)=\int_{0}^{a} x \exp (-x) u(x, t) d x
\end{aligned}
$$

This leads to the modified system

$$
\begin{aligned}
& u_{t}=-u_{x}-v_{1} u, \quad 0 \leq x \leq a, t \geq 0 \\
& u(0, t)=\frac{g(1, t) v_{1} v_{2}}{\left(v_{1}+1\right)^{2}}
\end{aligned}
$$

In the interface to the evaluation of the differential equation and boundary conditions, it is necessary to evaluate the integrals, which are computed with the values of $u(x, t)$ on the grid. The integrals are approximated using the trapezoid rule, commensurate with the truncation error in the integrator.

## Rationale: Example 3

This is a non-linear integro-differential problem involving non-local conditions for the differential equation and boundary conditions. Access to evaluation of these conditions is provided using reverse communication. It is not possible to solve this problem with forward communication, given the current subroutine interface. Optional changes are made to use an absolute error tolerance and non-zero time-smoothing. The time-smoothing value $\tau=1$ prevents grid lines from crossing.

```
    program PDE_1D_MG_EX03
! Population Dynāmi\overline{cs M}
        USE PDE_1d_mg_int
        USE ERROR_OPTION_PACKET
        IMPLICIT N}ON
        INTEGER, PARAMETER : : NPDE=1, N=101
        INTEGER IDO, I, NFRAMES
! Define array space for the solution.
        real(kind(1d0)) U(NPDE+1,N), MID(N-1), T0, TOUT, V_1, V_2
        real(kind(1d0)) :: ZERO=0D0, HALF=5D-1, ONE=1D0,&
            TWO=2D0, FOUR=4D0, DELTA_T=1D-1,TEND=5D0, A=5D0
        TYPE(D_OPTIONS) IOPT(3)
! Start loop tō integrate and record solution values.
        IDO=1
        DO
            SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                T0=ZERO
                    TOUT=DELTA_T
                    U(NPDE+1,1)=ZERO; U (NPDE+1,N ) = A
                    OPEN(FILE='PDE_ex03.out',UNIT=7)
                    NFRAMES=NINT ((TEND+DELTA_T)/DELTA_T)
                    WRITE(7, "(3I5, 4D14.5)")}\mathrm{ NPDE, N, NFRAMES, &
                    U(NPDE+1,1), U(NPDE+1,N), T0, TEND
                    IOPT (1)=D_OPTIONS(PDE_1D_MG_RELATIVE_TOLERANCE, ZERO)
                    IOPT (2) =D_OPTIONS (PDE_1D__MG_ABSOLUTE__TOLERANCE,1D-2)
                    IOPT (3)=D_OPTIONS (PDE_1D_MG_TIME_SMOOTHING,1D0)
! Update to the next output point.
! Write solution and check for final point.
            CASE (2)
            T0=TOUT
            IF(TO <= TEND) THEN
                WRITE(7,"(F10.5)") TOUT
```

```
            DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I,:)
            END DO
            TOUT=MIN(TOUT+DELTA_T,TEND)
            IF(TO == TEND)IDO=3
            END IF
! All completed. Solver is shut down.
            CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
    CASE (5)
            U(1,:)=EXP(-U (2,:)) / (TWO-EXP (-A))
            WRITE(7,"(F10.5)")T0
            DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I,:)
            END DO
! Define differential equations.
    CASE (6)
                            D_PDE_1D_MG_C (1,1)=ONE
                            D_PDE_1D_MG_-R(1)=-D_PDE_1D_MG_U(1)
! Evaluate the a
            V_1=HALF*SUM((U (1,1:N-1) +U (1,2:N) ) *&
                    (U(2,2:N) - U(2,1:N-1)))
                            D_PDE_1D_MG_Q(1)=V_1*D_PDE_1D_MG_U(1)
! Define boundary conditions.
                        CASE (7)
                            IF(PDE_1D_MG_LEFT) THEN
! Evaluate the approximate integral, for this t.
! A second integral is needed at the edge.
                            V_1=HALF*SUM((U (1,1:N-1) +U (1,2:N) ) *&
                        (U(2,2:N) - U(2,1:N-1)))
    MID=HALF* (U (2, 2:N) +U(2,1:N-1))
    V_2=HALF*SUM(MID*EXP(-MID) *&
    (\overline{U}(1,1:N-1) +U(1,2:N))* (U (2,2:N) -U (2,1:N-1)))
                    D_PDE_1D_MG_BETA=ZERO
D_PDE_1D_MG_GAMMA=G (ONE,D_PDE_1D_MG_T)*V_1*V_2/(V_1+ONE)**2-&
                    D_PDE_1D_MG_U
            ELSE
                    D_PDE_1D_MG_BETA=ZERO
                    D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_DUDX (1)
            END \overline{IF}
        END SELECT
! Reverse communication is used for the problem data.
            CALL PDE_1D_MG (TO, TOUT, IDO, U, IOPT=IOPT)
    END DO
CONTAINS
    FUNCTION G(z,t)
    IMPLICIT NONE
        REAL(KIND(1d0)) Z, T, G
        G=FOUR*Z* (TWO-TWO*EXP (-A) +EXP (-T) ) **2
        G=G/((ONE-EXP (-A)) * (ONE-(ONE+TWO*A) *&
            EXP(-TWO*A))* (1-EXP (-A) +EXP (-T)))
    END FUNCTION
```


## Example 4 - A Model in Cylindrical Coordinates

This example is from Blom and Zegeling (1994). The system models a reactor-diffusion problem:

$$
\begin{aligned}
& T_{z}=r^{-1} \frac{\partial\left(\beta r T_{r}\right)}{\partial r}+\gamma \exp \left(\frac{T}{1+\varepsilon T}\right) \\
& T_{r}(0, z)=0, T(1, z)=0, z>0 \\
& T(r, 0)=0,0 \leq r<1 \\
& \beta=10^{-4}, \gamma=1, \varepsilon=0.1
\end{aligned}
$$

The axial direction $Z$ is treated as a time coordinate. The radius $r$ is treated as the single space variable.

## Rationale: Example 4

This is a non-linear problem in cylindrical coordinates. Our example illustrates assigning $m=1$ in Equation 2. We provide an optional argument that resets this value from its default, $m=0$. Reverse communication is used to interface with the problem data.

```
    program PDE 1D MG EX04
! Reactor-Diffusion prōblem in cylindrical coordinates.
        USE pde_1d_mg_int
        USE error option packet
        IMPLICIT NONE
        INTEGER, PARAMETER :: NPDE=1, N=41
        INTEGER IDO, I, NFRAMES
! Define array space for the solution.
        real(kind(1d0)) T(NPDE+1,N), Z0, ZOUT
        real(kind(1d0)) :: ZERO=0D0, ONE=1D0, DELTA Z=1D-1,&
            ZEND=1D0, ZMAX=1D0, BTA=1D-4, GAMA=1D0, EPS=1D-1
        TYPE(D_OPTIONS) IOPT(1)
! Start loop tō integrate and record solution values.
        IDO=1
        DO
            SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                Z0=ZERO
                ZOUT=DELTA Z
                    T(NPDE+1,1)=ZERO;T(NPDE+1,N ) = ZMAX
                    OPEN(FILE='PDE_ex04.out',UNIT=7)
                    NFRAMES=NINT ((\overline{Z}END+DELTA_Z)/DELTA_Z)
                    WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES,&
                    T(NPDE+1,1), T(NPDE+1,N), Z0, ZEND
                    IOPT (1)=PDE_1D_MG_CYL_COORDINATES
! Update to the next output point.
! Write solution and check for final point.
            CASE (2)
            IF(ZO <= ZEND) THEN
                WRITE(7,"(F10.5)") ZOUT
```

```
            DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")T(I,:)
            END DO
            ZOUT=MIN(ZOUT+DELTA_Z,ZEND)
            IF(ZO == ZEND)IDO=3
            END IF
! All completed. Solver is shut down.
    CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
    CASE (5)
            T(1,:)=ZERO
            WRITE(7,"(F10.5)") Z0
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")T(I,:)
            END DO
! Define differential equations.
    CASE (6)
            D_PDE_1D_MG_C (1,1)=ONE
            D_PDE_1D_MG_R(1)=BTA*D_PDE_1D_MG_DUDX (1)
            D_PDE_1D_-MG_Q (1) = -GAM\overline{A}*EX\overline{P}(D_PDE__1D_MG_U (1)/&
            (ONE}+E\overline{P}S*\overline{D}_PDE_1D_MG_U(1))
! Define boundary conditions.
        CASE (7)
            IF(PDE_1D_MG_LEFT) THEN
                    D_PDE_1D_MG_BETA=ONE; D_PDE_1D_MG_GAMMA=ZERO
            ELSE
                    D_PDE_1D_MG_BETA=ZERO; D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_U (1)
            END \overline{IF}
    END SELECT
! Reverse communication is used for the problem data.
! The optional derived type changes the internal model
! to use cylindrical coordinates.
            CALL PDE_1D_MG (Z0, ZOUT, IDO, T, IOPT=IOPT)
        END DO
    end program
```


## Example 5 - A Flame Propagation Model

This example is presented more fully in Verwer, et al., (1989). The system is a normalized problem relating mass density $u(x, t)$ and temperature $v(x, t)$ :

$$
\begin{aligned}
& u_{t}=u_{x x}-u f(v) \\
& v_{t}=v_{x x}+u f(v) \\
& \text { where } f(z)=\gamma \exp (-\beta / z), \beta=4, \gamma=3.52 \times 10^{6} \\
& 0 \leq x \leq 1,0 \leq \mathrm{t} \leq 0.006 \\
& u(x, 0)=1, v(x, 0)=0.2 \\
& \begin{aligned}
u_{x} & =v_{x}=0, x=0 \\
u_{x} & =0, v=b(t), x=1, \text { where } \\
b(t) & =1.2, \text { for } t \geq 2 \times 10^{-4}, \text { and } \\
& =0.2+5 \times 10^{3} t, \text { for } 0 \leq t \leq 2 \times 10^{-4}
\end{aligned}
\end{aligned}
$$

## Rationale: Example 5

This is a non-linear problem. The example shows the model steps for replacing the banded solver in the software with one of the user's choice. Reverse communication is used for the interface to the problem data and the linear solver. Following the computation of the matrix factorization in DL2CRB, we declare the system to be singular when the reciprocal of the condition number is smaller than the working precision. This choice is not suitable for all problems. Attention must be given to detecting a singularity when this option is used.

```
    program PDE 1D MG EX05
! Flame propagation model
        USE pde_1d_mg_int
        USE ERRO\overline{O}_\overline{OPTİON_PACKET}
        USE Numerical_Libraries, ONLY :&
        dl2crb, dlfsrb
        IMPLICIT NONE
        INTEGER, PARAMETER : : NPDE=2, N=40, NEQ=(NPDE+1)*N
        INTEGER I, IDO, NFRAMES, IPVT(NEQ)
! Define array space for the solution.
        real(kind(1d0)) U(NPDE+1,N), T0, TOUT
! Define work space for the banded solver.
        real(kind(1d0)) WORK(NEQ), RCOND
        real(kind(1d0)) : : ZERO=0D0, ONE=1D0, DELTA_T=1D-4, &
            TEND=6D-3, XMAX=1D0, BTA=4D0, GAMA=3.52D6
        TYPE(D_OPTIONS) IOPT(1)
! Start loop to integrate and record solution values.
        IDO=1
        DO
            SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                TO=ZERO
                    TOUT=DELTA T
                    U (NPDE+1,1)=ZERO; U (NPDE+1,N) = XMAX
                    OPEN(FILE='PDE_ex05.out',UNIT=7)
                    NFRAMES=NINT ((TEND+DELTA_T)/DELTA_T)
```

```
            WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES, &
                U(NPDE+1,1), U(NPDE+1,N), T0, TEND
            IOPT(1)=PDE_1D_MG_REV_COMM_FACTOR_SOLVE
! Update to the next outpu
! Write solution and check for final point.
        CASE (2)
            T0=TOUT
                        IF(TO <= TEND) THEN
                        WRITE(7," (F10.5)") TOUT
                        DO I=1,NPDE+1
                        WRITE(7,"(4E15.5)")U(I,:)
                        END DO
                        TOUT=MIN(TOUT+DELTA T,TEND)
                        IF(TO == TEND)IDO=3
            END IF
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
        CASE (5)
            U(1,:)=ONE; U (2,:)=2D-1
            WRITE(7," (F10.5)")T0
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")U(I,:)
            END DO
! Define differential equations.
        CASE (6)
            D_PDE_1D_MG_C=ZERO
            D_PDE_1D_MG_C (1,1)=ONE; D_PDE_1D_MG_C (2, 2)=ONE
            D_PDE_1D_MG_R=D_PDE_1D_MG_DUDX
            D_PDE_1D_MG_Q(1)= D_PDE_1D_MG_U(1)*F(D_PDE_1D_MG_U (2))
            D_PDE_1D_-MG_Q (2)= -D_PDEE_1D_MG_Q (1)
! Define bounda\overline{ry cōndītiōons.}
        CASE (7)
            IF(PDE_1D_MG_LEFT) THEN
                    D P\overline{DE \}1D M
            ELSE
                D PDE 1D MG BETA (1)=ONE
                D_PDE_1D_MG_GAMMA (1)=ZERO
                D_PDE_1D_MG_BETA (2)=ZERO
                    IF(D PDE 1D MG T >= 2D-4) THEN
                    D_\overline{PDE_1'D_M}\mp@subsup{M}{_}{\prime}\overline{G}AMMA (2)=12D-1
                    ELSE
                    D_PDE_1D_MG_GAMMA (2)=2D-1+5D3*D_PDE_1D_MG_T
                END IF
                D_PDE_1D_MG_GAMMA (2)=D_PDE_1D_MG_GAMMA (2) - &
                    \overline{D}_PDE
            END IF
        CASE (8)
! Factor the banded matrix. This is the same solver used
```

```
! internally but that is not required. A user can substitute
! one of their own.
    call dl2crb (neq, d pde 1d mg a, pde 1d mg lda, &
    pde_1d_mg_iband, pde_1d_mg_iband, d_pde_1d_mg_a, &
    pde_1d_mg_lda, ipvt, rcond, work)
            IF(\overline{rcoñd}<== EPSILON(ONE)) pde_1d_mg_panic_flag = 1
        CASE (9)
! Solve using the factored banded matrix.
            call dlfsrb(neq, d_pde 1d_mg_a, pde_1d_mg_lda, &
            pde_1d_mg_iband, pde_1d_mg_iband, ipvt, &
            d_p\overline{de_\overline{1}d_mg_rhs, 1, \overline{d_p\overline{de_\}}\mathbf{1}d_mg_sol)}
        END SELECT
! Reverse communication is used for the problem data.
            CALL PDE_1D_MG (TO, TOUT, IDO, U, IOPT=IOPT)
        END DO
CONTAINS
        FUNCTION F(Z)
        IMPLICIT NONE
        REAL(KIND(1D0)) Z, F
            F=GAMA*EXP(-BTA/Z)
            END FUNCTION
        end program
```


## Example 6 - A 'Hot Spot’ Model

This example is presented more fully in Verwer, et al., (1989). The system is a normalized problem relating the temperature $u(x, t)$, of a reactant in a chemical system. The formula for $h(z)$ is equivalent to their example.

$$
\begin{aligned}
& u_{t}=u_{x x}+h(u) \\
& \text { where } h(z)=\frac{R}{a \delta}(1+a-z) \exp (-\delta(1 / z-1)) \\
& a=1, \delta=20, R=5 \\
& 0 \leq x \leq 1,0 \leq t \leq 0.29 \\
& u(x, 0)=1 \\
& u_{x}=0, x=0 \\
& u=1, x=1
\end{aligned}
$$

## Rationale: Example 6

This is a non-linear problem. The output shows a case where a rapidly changing front, or hot-spot, develops after a considerable way into the integration. This causes rapid change to the grid. An option sets the maximum order BDF formula from its default value of 2 to the theoretical stable maximum value of 5 .

```
USE pde_1d_mg_int
USE errōr_\overline{option_packet}
IMPLICIT NONE
```

```
    INTEGER, PARAMETER :: NPDE=1, N=80
    INTEGER I, IDO, NFRAMES
! Define array space for the solution.
    real(kind(1dO)) U(NPDE+1,N), TO, TOUT
    real(kind(1d0)) :: ZERO=0D0, ONE=1D0, DELTA_T=1D-2,&
        TEND=29D-2, XMAX=1D0, A=1D0, DELTA=2D1, R=5D0
    TYPE(D OPTIONS) IOPT(2)
! Start loop to integrate and record solution values.
    IDO=1
    DO
        SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                T0=ZERO
                        TOUT=DELTA T
                            U(NPDE+1,1)=ZERO; U (NPDE+1,N) =XMAX
                    OPEN(FILE='PDE ex06.out',UNIT=7)
            NFRAMES=(TEND+\overline{DELTA T)/DELTA T}
            WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES,&
                U(NPDE+1,1), U(NPDE+1,N), TO, TEND
! Illustrate allowing the BDF order to increase
! to its maximum allowed value.
                            IOPT (1)=PDE_1D_MG_MAX_BDF_ORDER
                IOPT (2)=5
! Update to the next output point.
! Write solution and check for final point.
        CASE (2)
            T0=TOUT
            IF(TO <= TEND) THEN
                WRITE(7,"(F10.5)")TOUT
                DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I,:)
                END DO
                TOUT=MIN(TOUT+DELTA_T,TEND)
                IF(TO == TEND)IDO=3
            END IF
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
        CASE (5)
            U(1,:)=ONE
            WRITE(7,"(F10.5)")T0
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")U(I,:)
            END DO
! Define differential equations.
        CASE (6)
            D_PDE_1D_MG_C=ONE
            D_PDE_1D_MG_R=D_PDE_1D_MG_DUDX
```



```
! Define boundary conditions.
        CASE (7)
            IF(PDE_1D_MG_LEFT) THEN
                    D PDE 1D MG BETA=ZERO
                            D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_DUDX
        ELSE
                        D_PDE_1D_MG_BETA=ZERO
                            D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_U(1) -ONE
        END \overline{IF}
        END SELECT
! Reverse communication is used for the problem data.
            CALL PDE_1D_MG (T0, TOUT, IDO, U, IOPT=IOPT)
        END DO
CONTAINS
        FUNCTION H(Z)
        real(kind(1dO)) Z, H
            H=(R/(A*DELTA))* (ONE+A-Z)*EXP (-DELTA* (ONE/Z-ONE))
        END FUNCTION
        end program
```


## Example 7 - Traveling Waves

This example is presented more fully in Verwer, et al., (1989). The system is a normalized problem relating the interaction of two waves, $u(x, t)$ and $v(x, t)$ moving in opposite directions. The waves meet and reduce in amplitude, due to the non-linear terms in the equation. Then they separate and travel onward, with reduced amplitude.

$$
\begin{aligned}
& u_{t}=-u_{x}-100 u v \\
& \left.\begin{array}{l}
v_{t}= \\
-0.5
\end{array}\right)=100 u v \\
& \begin{aligned}
u(x, 0) & =0.5,0 \leq t \leq 0.5 \\
\quad & =0, \text { otherwise }
\end{aligned} \\
& \begin{aligned}
v(x, 0) & =0.5(1+\cos (10 \pi x)) x \in[0.1,0.3], \text { and } \\
\quad & 0, \text { otherwise }
\end{aligned} \\
& \begin{aligned}
u & =v=0 \text { at both ends, } \mathrm{t} \geq 0
\end{aligned}
\end{aligned}
$$

## Rationale: Example 7

This is a non-linear system of first order equations.

```
    program PDE 1D MG EX07
! Traveling Waves
    USE pde_1d_mg_int
    USE errōr ōption packet
    IMPLICIT N
    INTEGER, PARAMETER :: NPDE=2, N=50
```

```
    INTEGER I, IDO, NFRAMES
! Define array space for the solution.
    real(kind(1d0)) U(NPDE+1,N), TEMP(N), TO, TOUT
    real(kind(1d0)) :: ZERO=0D0, HALF=5D-1, &
        ONE=1D0, DELTA_T=5D-2,TEND=5D-1, PI
    TYPE(D_OPTIONS) IOPT(5)
! Start loop to integrate and record solution values.
    IDO=1
    DO
        SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                        TO=ZERO
                        TOUT=DELTA T
                        U(NPDE+1,1)=-HALF; U (NPDE+1,N)=HALF
                        OPEN(FILE='PDE_ex07.out',UNIT=7)
            NFRAMES = (TEND+DELTA T)/DELTA T
                        WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES,&
                U(NPDE+1,1), U(NPDE+1,N), T0, TEND
            IOPT(1)=D OPTIONS (PDE 1D MG TIME SMOOTHING,1D-3)
            IOPT (2)=D_OPTIONS(PDE_1D_MG_RELATIVE_TOLERANCE, ZERO)
            IOPT (3) =D_OPTIONS(PDE_1D_MG_ABSOLUTE_TOLERANCE,1D-3)
            IOPT (4) =P\overline{DE_1D_MG_MAX_BD\overline{F}}0\overline{O}DER
                IOPT (5)=3
! Update to the next output point.
! Write solution and check for final point.
        CASE (2)
            T0=TOUT
            IF(TO <= TEND) THEN
                WRITE(7," (F10.5)") TOUT
                DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I, :)
                END DO
                TOUT=MIN(TOUT+DELTA_T,TEND)
                IF(TO == TEND)IDO=3
            END IF
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
! Define initial data values.
        CASE (5)
            TEMP=U (3,:)
            U(1,:)=PULSE (TEMP) ; U (2,:)=U(1, :)
            WHERE (TEMP < -3D-1 .or. TEMP > -1D-1) U(1,:)=ZERO
            WHERE (TEMP < 1D-1 .or. TEMP > 3D-1) U(2,:)=ZERO
            WRITE(7,"(F10.5)")T0
            DO I=1,NPDE+1
                WRITE (7, "(4E15.5)")U(I,:)
            END DO
```

```
! Define differential equations.
    CASE (6)
            D PDE 1D MG C=ZERO
            D_PDE_1D_MG_C (1,1)=ONE; D_PDE_1D_MG_C (2,2)=ONE
            D_PDE_1D_MG_R=D_PDE_1D_MG_U
            D_PDE_1D_MG_R(1)}=-\mp@subsup{D}{~}{-
            D_PDE_1D_MG_Q(1)=100D0*D_PDE_1D_MG_U(1)*D_PDE_1D_MG_U(2)
            D_PDE_1D_MG_Q (2) = D_PDE_1D_MG_Q(\overline{1})
! Define boundary conditions.
    CASE (7)
            D_PDE_1D_MG_BETA=ZERO;D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_U
            END SELECT
! Reverse communication is used for the problem data.
            CALL PDE 1D MG (T0, TOUT, IDO, U, IOPT=IOPT)
            END DO
CONTAINS
    FUNCTION PULSE(Z)
    real(kind(1d0)) Z(:), PULSE(SIZE(Z))
            PI=ACOS (-ONE)
            PULSE=HALF* (ONE+COS (10D0*PI*Z))
        END FUNCTION
    end program
```


## Example 8 - Black-Scholes

The value of a European "call option," $c(s, t)$, with exercise price $e$ and expiration date $T$, satisfies the "asset-or-nothing payoff" $c(s, T)=s, s \geq e ;=0, s<e$. Prior to expiration $c(s, t)$ is estimated by the Black-Scholes differential equation

$$
c_{t}+\frac{\sigma^{2}}{2} s^{2} c_{s s}+r s c_{s}-r c \equiv c_{t}+\frac{\sigma^{2}}{2}\left(s^{2} c_{s}\right)_{s}+\left(r-\sigma^{2}\right) s c_{s}-r c=0 .
$$

The parameters in the model are the risk-free interest rate, $r$, and the stock volatility, $\sigma$. The boundary conditions are $c(0, t)=0$ and $c_{s}(s, t) \approx 1, s \rightarrow \infty$. This development is described in Wilmott, et al. (1995), pages 41-57. There are explicit solutions for this equation based on the Normal Curve of Probability. The normal curve, and the solution itself, can be efficiently computed with the IMSL function ANORDF, IMSL (1994), page 186. With numerical integrationthe equation itself or the payoff can be readily changed to include other formulas, $c(s, T)$, and corresponding boundary conditions. We use

$$
e=100, r=0.08, T-t=0.25, \sigma^{2}=0.04, s_{L}=0 \text {, and } s_{R}=150 .
$$

## Rationale: Example 8

This is a linear problem but with initial conditions that are discontinuous. It is necessary to use a positive time-smoothing value to prevent grid lines from crossing. We have used an absolute tolerance of $10^{-3}$. In \$US, this is one-tenth of a cent.

```
    program PDE_1D_MG_EX08
! Black-Scholes \overline{cal\overline{l}}\mathrm{ price}
    USE pde_1d_mg_int
    USE error_option_packet
    IMPLICIT N
    INTEGER, PARAMETER : : NPDE=1, N=100
    INTEGER I, IDO, NFRAMES
! Define array space for the solution.
    real(kind(1d0)) U(NPDE+1,N), TO, TOUT, SIGSQ, XVAL
    real(kind(1d0)) :: ZERO=0D0, HALF=5D-1, ONE=1D0, &
        DELTA T=25D-3, TEND=25D-2, XMAX=150, SIGMA=2D-1, &
        R=8D-\overline{2}, E=100D0
    TYPE(D_OPTIONS) IOPT(5)
! Start loop tō integrate and record solution values.
    IDO=1
    DO
        SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                        T0=ZERO
                        TOUT=DELTA T
                        U(NPDE+1,1)=ZERO; U (NPDE+1,N) = XMAX
                OPEN(FILE='PDE ex08.Out',UNIT=7)
                NFRAMES=NINT ((\overline{TEND+DELTA T)/DELTA T)}
                WRITE(7, "(3I5, 4D14.5)") NPDE, N, NFRAMES, &
                    U(NPDE+1,1), U(NPDE+1,N), T0, TEND
                    SIGSQ=SIGMA**2
! Illustrate allowing the BDF order to increase
! to its maximum allowed value.
                IOPT(1)=PDE_1D_MG_MAX_BDF_ORDER
                IOPT (2)=5
                IOPT (3)=D_OPTIONS(PDE_1D_MG_TIME_SMOOTHING, 5D-3)
                IOPT (4)=D_OPTIONS(PDE_1D__MG_RELATIIVE_TOLERANCE, ZERO)
                IOPT (5) =D_OPTIONS (PDE__1D_MG_ABSOLUTE__TOLERANCE,1D-2)
! Update to the next output point.
! Write solution and check for final point.
        CASE (2)
            T0=TOUT
            IF(TO <= TEND) THEN
                WRITE(7,"(F10.5)") TOUT
                DO I=1,NPDE+1
                    WRITE(7,"(4E15.5)")U(I,:)
                END DO
                TOUT=MIN(TOUT+DELTA_T,TEND)
                    IF(TO == TEND)IDO=3
                END IF
! All completed. Solver is shut down.
        CASE (3)
            CLOSE (UNIT=7)
            EXIT
```

```
! Define initial data values.
    CASE (5)
            U(1,:)=MAX(U(NPDE+1,:) -E,ZERO) ! Vanilla European Call
            U(1,:)=U(NPDE+1,:) ! Asset-or-nothing Call
            WHERE (U(1,:) <= E) U(1,:)=ZERO ! on these two lines
            WRITE (7,"(F10.5)")T0
            DO I=1,NPDE+1
                WRITE(7,"(4E15.5)")U(I,:)
            END DO
! Define differential equations.
        CASE (6)
            XVAL=D_PDE_1D_MG_X
            D PDE 1D MG C=ONE
            D PDE 1D MG R=D PDE 1D MG DUDX*XVAL**2*SIGSQ*HALF
            D_PDE_1D_MG_Q=-(R-SIGSQ)*XVAL*D_PDE_1D_MG_DUDX+R*D_PDE_1D_MG_U
! Define boundary conditions.
        CASE (7)
            IF(PDE_1D_MG_LEFT) THEN
                D_P\overline{DE_}\=1D_M
                    D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_U
            ELSE
                    D_PDE_1D_MG_BETA=ZERO
                    D_PDE_1D_MG_GAMMA=D_PDE_1D_MG_DUDX (1) -ONE
            END IF
        END SELECT
! Reverse communication is used for the problem data.
        CALL PDE_1D_MG (TO, TOUT, IDO, U, IOPT=IOPT)
        END DO
    end program
```


## Example 9 - Electrodynamics, Parameters Studied with MPI



REQUIRED
For a detailed description of MPI Requirements see "Dense Matrix Parallelism Using MPI" in Chapter 10 of this manual.

This example, described above in Example 1, is from Blom and Zegeling (1994). The system parameters $\varepsilon, p$, and $\eta$, are varied, using uniform random numbers. The intervals studied are $0.1 \leq \varepsilon \leq 0.2,0.1 \leq p \leq 0.2$, and $10 \leq \eta \leq 20$. Using $N=21$ grid values and other program options, the elapsed time, parameter values, and the value $\left.v(x, t)\right|_{x=1, t=4}$ are sent to the root node. This information is written on a file. The final summary includes the minimum value of

$$
\left.v(x, t)\right|_{x=1, t=4},
$$

and the maximum and average time per integration, per node.

## Rationale: Example 9

This is a non-linear simulation problem. Using at least two integrating processors and MPI allows more values of the parameters to be studied in a given time than with a single processor. This code is valuable as a study guide when an application needs to estimate timing and other output parameters. The simulation time is controlled at the root node. An integration is started, after receiving results, within the first SIM_TIME seconds. The elapsed time will be longer than SIM_TIME by the slowest processor's time for its last integration.

```
        program PDE_1D_MG_EX09
! Electrodynamics Model, parameter study.
        USE PDE_1d_mg_int
        USE MPI_SETUP_INT
        USE RAN\overline{D INT}
        USE SHOW INT
        IMPLICIT NONE
        INCLUDE "mpif.h"
        INTEGER, PARAMETER :: NPDE=2, N=21
        INTEGER I, IDO, IERROR, CONTINUE, STATUS(MPI_STATUS_SIZE)
        INTEGER, ALLOCATABLE :: COUNTS(:)
! Define array space for the solution.
        real(kind(1d0)) :: U(NPDE+1,N), T0, TOUT
        real(kind(1d0)) :: ZERO=0D0, ONE=1D0,DELTA_T=10D0, TEND=4D0
! SIM_TIME is the number of seconds to run the simulation.
        real(kind(1d0)) :: EPS, P, ETA, Z, TWO=2D0, THREE=3D0, SIM_TIME=60D0
        real(kind(1d0)) :: TIMES, TIMEE, TIMEL, TIME, TIME_SIM, V_MIN, &
        DATA(5)
        real(kind(1d0)), ALLOCATABLE :: AV_TIME(:), MAX_TIME(:)
        TYPE(D_OPTIONS) IOPT(4), SHOW_IOPT(2)
        TYPE (S_OPTIONS) SHOW_INTOPT(2)
        MP_NPRŌCS=MP_SETUP(1)
        MPI_NODE_PRIORITY=(/ (I-1,I=1,MP_NPROCS) / )
! If NP_NPRO}CS=1\overline{,}\mathrm{ the program stops. Change
! MPI_ROOOT_WORKS=.TRUE. if MP_NPROCS=1.
        MPI__ROOT_WORKS=.FALSE.
        IF(.NOT. MPI_ROOT_WORKS .and. MP_NPROCS == 1) STOP
        ALLOCATE (AV_\overline{TIME (MP_NPROCS), MAX_TIME (MP_NPROCS), COUNTS (MP_NPROCS))}
! Get time start for simulation timing.
        TIME=MPI WTIME()
        IF(MP_RANNK == 0) OPEN(FILE='PDE_ex09.out',UNIT=7)
    SIMULATE: DO-
! Pick random parameter values.
    EPS=1D-1*(ONE+rand(EPS))
    P=1D-1* (ONE+rand (P))
    ETA=10D0* (ONE+rand (ETA))
! Start loop to integrate and communicate solution times.
            IDO=1
! Get time start for each new problem.
        DO
            IF(.NOT. MPI_ROOT_WORKS .and. MP_RANK == 0) EXIT
            SELECT CASE (IDO)
! Define values that determine limits.
            CASE (1)
                                    T0=ZERO
                                TOUT=1D-3
```

```
    U(NPDE+1,1)=ZERO;U(NPDE+1,N) =ONE
    IOPT (1)=PDE_1D_MG_MAX_BDF_ORDER
    IOPT (2)=5
    IOPT (3)=D_OPTIONS (PDE_1D_MG_RELATIVE_TOLERANCE,1D-2)
    IOPT (4)=D_OPTIONS (PDE_1D_MG_ABSOLUTE_TOLERANCE,1D-2)
    TIMES=MPI_WTIME()
! Update to the next output point.
! Write solution and check for final point.
    CASE (2)
    TO=TOUT;TOUT=TOUT*DELTA_T
    IF(TO >= TEND) IDO=3
    TOUT=MIN(TOUT, TEND)
! All completed. Solver is shut down.
            CASE (3)
                            TIMEE=MPI_WTIME()
                            EXIT
! Define initial data values.
        CASE (5)
                            U(1,:)=1D0;U(2,:)=0D0
! Define differential equations.
    CASE (6)
D_PDE_1D_MG_C=0D0;D_PDE_1D_MG_C (1,1)=1D0;D_PDE_1D_MG_C (2,2)=1D0
    D_\overline{PDE_}\mp@subsup{\overline{1}}{D}{D}\\overline{M}G_\overline{R}=P*D_PDE_1D_\overline{MG_DUUDX}
D_PDE_1D_MG_R(1)=D_PDE_1D_MG_R(1)*EPS
    Z=ETA}\overline{*}(D_PD\overline{E}_1D_MG_U(1)-D_PDE_1D_MG_U(2))/THREE
    D_PDE_1D_MG_\overline{Q}(1)=E\overline{X}P(Z)-E\overline{X}P(-\overline{TWO}\overline{*}Z)
    D_PDE_1D_MG_Q (2)=-D_PDE_1D_MG_Q(1)
! Define boundary conditions.
        CASE (7)
            IF(PDE_1D_MG_LEFT) THEN
                                    D_PDE_\overline{1D_MG_BETA(1)=1D0;D_PDE_1D_MG_BETA (2)=0D0}
D_PDE_1D_MG_GAMMA (1)=0D0;D_PDE_1D_MG_GAMMA (2)=D_PDE_1D_MG_U (2)
        ELSE
                D_PDE_1D_MG_BETA (1)=0D0;D_PDE_1D_MG_BETA (2)=1D0
                D - PDE - 1D - MG - GAMMA (1)=D PDE \ 1D - MG - U(\overline{1})- &
                    1\overline{D}0;D_PD\overline{E}_1\overline{D}_MG_GAMMA (\overline{2})=0\overline{D}0
            END IF
        END SELECT
! Reverse communication is used for the problem data.
        CALL PDE_1D_MG (T0, TOUT, IDO, U)
        END DO
        TIMEL=TIMEE-TIMES
        DATA=(/EPS, P, ETA, U(2,N), TIMEL/)
        IF(MP_RANK > 0) THEN
! Send parameters and time to the root.
        CALL MPI SEND(DATA, 5, MPI DOUBLE PRECISION,0, MP RANK, &
        MP_LIBRA\overline{RY_WORLD, IERROR)}
! Receive back a "go/stop" flag.
        CALL MPI RECV(CONTINUE, 1, MPI INTEGER, 0, MPI ANY TAG, &
        MP_LIBRA\overline{R}Y_WORLD, STATUS, IERRO\overline{R})
! If root notes that time is up, it sends node a quit flag.
        IF(CONTINUE == 0) EXIT SIMULATE
```

ELSE
! If root is working, record its result and then stand ready
! for other nodes to send.
IF (MPI_ROOT_WORKS) WRITE (7,*) MP_RANK, DATA
! If all nodes have repor̄̄ed, then quit.
IF (COUNT (MPI_NODE_PRIORITY >= 0) == 0) EXIT SIMULATE
! See if time is up. Some nodes still must report.
IF (MPI WTIME () -TIME >= SIM_TIME) THEN
CONTINUE=0
ELSE
CONTINUE=1
END IF
! Root receives simulation data and finds which node sent it. IF (MP_NPROCS > 1) THEN

CALL MPI_RECV (DATA, 5, MPI_DOUBLE_PRECISION, \&
MPI_ANY_SOURCE, MPI_ANY_TAG, MP_LIBRARY_WORLD, \& STAT̄US, ${ }^{-}$IERROR)
WRITE (7,*) STATUS (MPI_SOURCE), DATA
! If time at the root has elapsed, nodes receive signal to stop.
! Send the reporting node the "go/stop" flag.
! Mark if a node has been stopped.
CALL MPI SEND (CONTINUE, 1, MPI INTEGER, \& STATUS (MPI_SOURCE), \&0, MP_LIBRARY_WORLD, IERROR)
IF (CONTINUE $==0$ ) MPI_NODE_PRIORITY $\left(S T A T U S ~\left(M P I \_S O U R C E\right) ~+1\right) ~ \& ~$ =- MPI_NODE_PRIORITY(STATUS $\left.\overline{\left(M P I \_S O U R C E\right)}+1\right)-1$
END IF
IF (CONTINUE == 0) MPI_NODE_PRIORITY(1)=-1
END IF
END DO SIMULATE
IF (MP_RANK == 0) THEN
END $\mathrm{DILE}(\mathrm{UNIT}=7$ ) ; REWIND (UNIT=7)
! Read the data. Find extremes and averages.
MAX_TIME=ZERO;AV_TIME=ZERO; COUNTS=0;V_MIN=HUGE (ONE)
DO
$\operatorname{READ}(7, *, \operatorname{END}=10) \mathrm{I}, \mathrm{DATA}$
$\operatorname{COUNTS}(I+1)=\operatorname{COUNTS}(I+1)+1$
AV_TIME $(I+1)=A V \_T I M E(I+1)+D A T A(5)$
$\operatorname{IF} \overline{( } \operatorname{MAX}$ TIME $(I+1 \overline{)}<$ DATA (5)) MAX_TIME $(I+1)=$ DATA (5)
V_MIN= $\bar{M} I N\left(V \_M I N, ~ D A T A(4)\right)$
END $\bar{D} O$
10 CONTINUE
CLOSE (UNIT=7)
! Set printing Index to match node numbering.
SHOW_IOPT (1) = SHOW_STARTING_INDEX_IS
SHOW_IOPT (2) $=0$
SHOW_INTOPT (1)=SHOW_STARTING_INDEX_IS
SHOW_INTOPT (2) $=0$
CALL SHOW (MAX_TIME,"Maximum Integration Time, per
process:",IOPT=SHOW_IOPT)
AV_TIME= $\bar{A} V$ _TIME/MAX $(1$, COUNTS $)$
CALL SHOW (AV_TIME, "Average Integration Time, per
process:",IOPT=SHOW_IOPT)
CALL SHOW̄ (COUNTS, "Number of Integrations", IOPT=SHOW_INTOPT)
WRITE(*," (1x, A, F6.3)") "Minimum value for $v(x, t)$, at $x=1, t=4$ :
", V_MIN

```
    END IF
    MP_NPROCS=MP_SETUP("Final")
end program
```


## MOLCH

Solves a system of partial differential equations of the form $u_{t}=f\left(x, t, u, u_{x}, u_{x x}\right)$ using the method of lines. The solution is represented with cubic Hermite polynomials.

## Required Arguments

IDO - Flag indicating the state of the computation. (Input/Output)

## IDO State

1 Initial entry
2 Normal reentry

3 Final call, release workspace

Normally, the initial call is made with $I D O=1$. The routine then sets $I D O=2$, and this value is then used for all but the last call that is made with $I D O=3$.

FCNUT - User-supplied SUBROUTINE to evaluate the function $u_{t}$. The usage is
CALL FCNUT (NPDES, X, T, U, UX, UXX, UT), where
NPDES - Number of equations. (Input)
X - Space variable, $x$. (Input)
T - Time variable, $t$. (Input)
U - Array of length NPDES containing the dependent variable values,
u. (Input)

UX - Array of length NPDES containing the first derivatives $u_{x}$.
(Input)
UXX - Array of length NPDES containing the second derivative $u_{x x}$. (Input)
UT - Array of length NPDES containing the computed derivatives, $u_{t}$. (Output)

The name FCNUT must be declared EXTERNAL in the calling program.
FCNBC - User-supplied SUBROUTINE to evaluate the boundary conditions. The boundary conditions accepted by MOLCH are $\alpha_{k} u_{k}+\beta_{k} u_{x} \equiv \gamma_{k}$. Note: Users must supply the values $\alpha_{k}$ and $\beta_{k}$, which determine the values $\gamma_{k}$. Since the $\gamma_{k}$ can depend on $t$, values of $\gamma^{\prime}{ }_{k}$ are also required. Users must supply these values. The usage is
CALL FCNBC (NPDES, X, T, ALPHA, BTA, GAMMAP), where

NPDES - Number of equations. (Input)
X - Space variable, $x$. This value directs which boundary condition to compute. (Input)
T - Time variable, $t$. (Input)
ALPHA - Array of length NPDES containing the $\alpha_{k}$ values. (Output)

BTA - Array of length NPDES containing the $\beta_{k}$ values. (Output)
GAMMAP - Array of length NPDES containing the values of the derivatives, $\frac{d \gamma_{k}}{d t}=\gamma_{k}^{\prime}$ (Output)

The name FCNBC must be declared EXTERNAL in the calling program.
$\boldsymbol{T}$ - Independent variable, $t$. (Input/Output)
On input, $T$ supplies the initial time, $t_{0}$. On output, $T$ is set to the value to which the integration has been updated. Normally, this new value is TEND.

TEND - Value of $t=$ tend at which the solution is desired. (Input)
XBREAK - Array of length NX containing the break points for the cubic Hermite splines used in the $x$ discretization. (Input)
The points in the array XBREAK must be strictly increasing. The values XBREAK (1) and $\operatorname{XBREAK}(N X)$ are the endpoints of the interval.
$\boldsymbol{Y}$ - Array of size NPDES by NX containing the solution. (Input/Output)
The array $Y$ contains the solution as $\mathrm{Y}(k, i)=u_{k}(x$, tend $)$ at $x=\operatorname{XBREAK}(i)$. On input, Y contains the initial values. It MUST satisfy the boundary conditions. On output, y contains the computed solution.
There is an optional application of MOLCH that uses derivative values, $u_{x}\left(x, t_{0}\right)$. The user allocates twice the space for $Y$ to pass this information. The optional derivative information is input as

$$
\mathrm{Y}(\mathrm{k}, \mathrm{i}+\mathrm{NX})=\frac{\partial u_{k}}{\partial x}\left(x, t_{0}\right)
$$

at $x=\mathrm{X}(i)$. The array Y contains the optional derivative values as output:

$$
\mathrm{Y}(\mathrm{k}, \mathrm{i}+\mathrm{NX})=\frac{\partial u_{k}}{\partial x}(x, \text { tend })
$$

at $x=\mathrm{x}(i)$. To signal that this information is provided, use an options manager call as outlined in Comment 3 and illustrated in Examples 3 and 4.

## Optional Arguments

NPDES - Number of differential equations. (Input)
Default: NPDES $=\operatorname{size}(\mathrm{Y}, 1)$.
$N X$ - Number of mesh points or lines. (Input)
Default: $N X=\operatorname{size}(Y, 2)$.

TOL - Differential equation error tolerance. (Input)
An attempt is made to control the local error in such a way that the global relative error is proportional to TOL.
Default: TOL $=100 . *$ machine precision.
HINIT - Initial step size in the $t$ integration. (Input)
This value must be nonnegative. If HINIT is zero, an initial step size of $0.001 \mid$ tend $-t_{0} \mid$ will be arbitrarily used. The step will be applied in the direction of integration.
Default: $\mathrm{HINIT}=0.0$.
$\boldsymbol{L D Y}$ - Leading dimension of $Y$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDY $=\operatorname{size}(Y, 1)$.

## FORTRAN 90 Interface

Generic: CALL MOLCH (IDO, FCNUT, FCNBC, T, TEND, XBREAK, Y [,...])

Specific: The specific interface names are S_MOLCH and D_MOLCH.

## FORTRAN 77 Interface

Single: CALL MOLCH (IDO, FCNUT, FCNBC, NPDES, T, TEND, NX, XBREAK, TOL, HINIT, Y, LDY)

Double: The double precision name is DMOLCH.

## Description

Let $M=$ NPDES, $N=$ NX and $x_{i}=\operatorname{XBREAK}(\mathrm{I})$. The routine MOLCH uses the method of lines to solve the partial differential equation system

$$
\frac{\partial u_{k}}{\partial t}=f_{k}\left(x, t, u_{1}, \ldots u_{M}, \frac{\partial u_{1}}{\partial x}, \ldots \frac{\partial u_{M}}{\partial x}, \frac{\partial^{2} u_{1}}{\partial x^{2}}, \ldots \frac{\partial^{2} u_{M}}{\partial x^{2}}\right)
$$

with the initial conditions

$$
u_{k}=u_{k}(x, t) \quad \text { at } t=t_{0}
$$

and the boundary conditions

$$
\alpha_{k} u_{k}+\beta_{k} \frac{\partial u_{k}}{\partial x}=\gamma_{k}(t) \quad \text { at } x=x_{1} \text { and at } x=x_{N}
$$

for $k=1, \ldots, M$.
Cubic Hermite polynomials are used in the $x$ variable approximation so that the trial solution is expanded in the series

$$
\hat{u}_{k}(x, t)=\sum_{i=1}^{N}\left(a_{i, k}(t) \phi_{i}(x)+b_{i, k}(t) \psi_{i}(x)\right)
$$

where $\phi_{i}(x)$ and $\psi_{i}(x)$ are the standard basis functions for the cubic Hermite polynomials with the knots $x_{1}<x_{2}<\ldots<x_{N}$. These are piecewise cubic polynomials with continuous first derivatives. At the breakpoints, they satisfy

$$
\begin{aligned}
& \phi_{i}\left(x_{l}\right)=\delta_{i l} \psi_{i}\left(x_{l}\right)=0 \\
& \frac{d \phi_{i}}{d x}\left(x_{l}\right)=0 \quad \frac{d \psi_{i}}{d x}\left(x_{l}\right)=\delta_{i l}
\end{aligned}
$$

According to the collocation method, the coefficients of the approximation are obtained so that the trial solution satisfies the differential equation at the two Gaussian points in each subinterval,

$$
\begin{aligned}
& p_{2 j-1}=x_{j}+\frac{3-\sqrt{3}}{6}\left(x_{j+1}-x_{j}\right) \\
& p_{2 j}=x_{j}+\frac{3+\sqrt{3}}{6}\left(x_{j+1}+x_{j}\right)
\end{aligned}
$$

for $j=1, \ldots, N$. The collocation approximation to the differential equation is

$$
\begin{aligned}
& \sum_{i=1}^{N} \frac{d a_{i, k}}{d t} \phi_{i}\left(p_{j}\right)+\frac{d b_{i, k}}{d t} \psi_{i}\left(p_{j}\right)= \\
& \quad f_{k}\left(p_{j}, t, \hat{u}_{1}\left(p_{j}\right), \ldots, \hat{u}_{M}\left(p_{j}\right), \ldots,\left(\hat{u}_{1}\right)_{x x}\left(p_{j}\right), \ldots,\left(\hat{u}_{M}\right)_{x x}\left(p_{j}\right)\right)
\end{aligned}
$$

for $k=1, \ldots, M$ and $j=1, \ldots, 2(N-1)$.
This is a system of $2 M(N-1)$ ordinary differential equations in $2 M N$ unknown coefficient functions, $a_{i, k}$ and $b_{i, k}$. This system can be written in the matrix-vector form as $A d c / d t=F(t, y)$ with $c\left(t_{0}\right)=c_{0}$ where $c$ is a vector of coefficients of length $2 M N$ and $c_{0}$ holds the initial values of the coefficients. The last $2 M$ equations are obtained by differentiating the boundary conditions

$$
\alpha_{k} \frac{d a_{k}}{d t}+\beta_{k} \frac{d b_{k}}{d t}=\frac{d \gamma_{k}}{d t}
$$

for $k=1, \ldots, M$.
The initial conditions $u_{k}\left(x, t_{0}\right)$ must satisfy the boundary conditions. Also, the $\gamma_{k}(t)$ must be continuous and have a smooth derivative, or the boundary conditions will not be properly imposed for $t>t_{0}$.

If $\alpha_{k}=\beta_{k}=0$, it is assumed that no boundary condition is desired for the $k$-th unknown at the left endpoint. A similar comment holds for the right endpoint. Thus, collocation is done at the endpoint. This is generally a useful feature for systems of first-order partial differential equations.
If the number of partial differential equations is $M=1$ and the number of breakpoints is $N=4$, then

$$
A=\left[\begin{array}{ccccccc}
\alpha_{1} & \beta_{1} & & & & & \\
\phi_{1}\left(p_{1}\right) & \psi_{1}\left(p_{1}\right) & \phi_{2}\left(p_{1}\right) & \psi_{2}\left(p_{1}\right) & & & \\
\phi_{1}\left(p_{2}\right) & \psi_{1}\left(p_{2}\right) & \phi_{2}\left(p_{2}\right) & \psi_{2}\left(p_{2}\right) & & & \\
& & \phi_{3}\left(p_{3}\right) & \psi_{3}\left(p_{3}\right) & \phi_{4}\left(p_{3}\right) & \psi_{4}\left(p_{3}\right) & \\
& & \phi_{3}\left(p_{4}\right) & \psi_{3}\left(p_{4}\right) & \phi_{4}\left(p_{4}\right) & \psi_{4}\left(p_{4}\right) & \\
& & & & \phi_{5}\left(p_{5}\right) & \psi_{5}\left(p_{5}\right) & \phi_{6}\left(p_{5}\right) \\
& & & \psi_{5}\left(p_{5}\right) \\
& & & & \left.\phi_{6}\right) & \psi_{5}\left(p_{6}\right) & \phi_{6}\left(p_{6}\right) \\
\psi_{6}\left(p_{6}\right) \\
& & & & & & \alpha_{4} \\
& & & & & \beta_{4}
\end{array}\right]
$$

The vector $c$ is

$$
c=\left[a_{1}, b_{1}, a_{2}, b_{2}, a_{3}, b_{3}, a_{4}, b_{4}\right]^{T}
$$

and the right-side $F$ is

$$
F=\left[\gamma^{\prime}\left(x_{1}\right), f\left(p_{1}\right), f\left(p_{2}\right), f\left(p_{3}\right), f\left(p_{4}\right), f\left(p_{5}\right), f\left(p_{6}\right), \gamma^{\prime}\left(x_{4}\right)\right]^{T}
$$

If $M>1$, then each entry in the above matrix is replaced by an $M \times M$ diagonal matrix. The element $\alpha_{1}$ is replaced by $\operatorname{diag}\left(\alpha_{1,1}, \ldots, \alpha_{1, \mathrm{M}}\right)$. The elements $\alpha_{N}, \beta_{1}$ and $\beta_{N}$ are handled in the same manner. The $\phi_{i}\left(p_{j}\right)$ and $\psi_{i}\left(p_{j}\right)$ elements are replaced by $\phi_{i}\left(p_{j}\right) I_{M}$ and $\psi_{i}\left(p_{j}\right) I_{M}$ where $I_{M}$ is the identity matrix of order $M$. See Madsen and Sincovec (1979) for further details about discretization errors and Jacobian matrix structure.

The input/output array Y contains the values of the $a_{k, i}$. The initial values of the $b_{k, i}$ are obtained by using the IMSL cubic spline routine CSINT (see Chapter 3, Interpolation and Approximation) to construct functions

$$
\hat{u}_{k}\left(x, t_{0}\right)
$$

such that

$$
\hat{u}_{k}\left(x_{i}, t_{0}\right)=a_{k i}
$$

The IMSL routine CSDER, see Chapter 3, Interpolation and Approximation, is used to approximate the values

$$
\frac{d \hat{U}_{k}}{d x}\left(x_{i}, t_{0}\right) \equiv b_{k, i}
$$

There is an optional usage of MOLCH that allows the user to provide the initial values of $b_{k, i}$.
The order of matrix $A$ is $2 M N$ and its maximum bandwidth is $6 M-1$. The band structure of the Jacobian of $F$ with respect to $c$ is the same as the band structure of $A$. This system is solved using a modified version of IVPAG. Some of the linear solvers were removed. Numerical Jacobians are used exclusively. The algorithm is unchanged. Gear's BDF method is used as the default because the system is typically stiff.

We now present four examples of PDEs that illustrate how users can interface their problems with IMSL PDE solving software. The examples are small and not indicative of the complexities that most practitioners will face in their applications. A set of seven sample application problems, some of them with more than one equation, is given in Sincovec and Madsen (1975). Two further examples are given in Madsen and Sincovec (1979).

## Comments

1. Workspace may be explicitly provided, if desired, by use of M2LCH/DM2LCH. The reference is:

CALL M2LCH (IDO, FCNUT, FCNBC, NPDES, T, TEND, NX, XBREAK, TOL, HINIT, Y, LDY, WK, IWK)

The additional arguments are as follows:
$\boldsymbol{W K}$ - Work array of length 2 NX * $\operatorname{NPDES}\left(12\right.$ * NPDES $^{2}+21$ * NPDES + 9). WK should not be changed between calls to M2LCH.
$\boldsymbol{I W K}$ - Work array of length 2NX * NPDES. IWK should not be changed between calls to M2LCH.
2. Informational errors

Type Code
41 After some initial success, the integration was halted by repeated error test failures.
42 On the next step, $\mathrm{X}+\mathrm{H}$ will equal X . Either TOL is too small or the problem is stiff.
43 After some initial success, the integration was halted by a test on TOL.
44 Integration was halted after failing to pass the error test even after reducing the step size by a factor of $1.0 \mathrm{E}+10$. TOL may be too small.
45 Integration was halted after failing to achieve corrector convergence even after reducing the step size by a factor of $1.0 \mathrm{E}+10$. TOL may be too small.
3. Optional usage with Chapter 11 Option Manager

11 This option consists of the parameter PARAM, an array with 50 components. See IVPAG for a more complete documentation of the contents of this array. To reset this option, use the subprogram SUMAG for single precision, and DUMAG (see Chapter 11, Utilities) for double precision. The entry PARAM(1) is assigned the initial step, hinit. The entries PARAM(15) and PARAM(16) are assigned the values equal to the number of lower and upper diagonals that will occur in the Newton method for solving the BDF corrector equations. The value $\operatorname{PARAM}(17)=1$ is used to signal that the $x$ derivatives of the initial data are provided in the the array Y. The output values PARAM(31)-PARAM(36), showing technical data about the ODE integration, are available with another option
manager subroutine call. This call is made after the storage for MOLCH is released. The default values for the first 20 entries of PARAM are ( $0,0, \operatorname{amach}(2)$, 500., 0., 5., 0, 0, 1., 3., 1., 2., 2., 1., amach(6), amach(6), $0, \operatorname{sqrt}(\operatorname{amach}(4)), 1 .$, 0.$)$. Entries 21-50 are defaulted to amach(6).

## Example 1

The normalized linear diffusion PDE, $u_{t}=u_{x x}, 0 \leq x \leq 1, t>t_{0}$, is solved. The initial values are $t_{0}=0, u\left(x, t_{0}\right)=u_{0}=1$. There is a "zero-flux" boundary condition at $x=1$, namely $u_{x}(1, t)=0$, ( $t>t_{0}$ ). The boundary value of $u(0, t)$ is abruptly changed from $u_{0}$ to the value $u_{1}=0.1$. This transition is completed by $t=t_{\delta}=0.09$.

Due to restrictions in the type of boundary conditions sucessfully processed by MOLCH, it is necessary to provide the derivative boundary value function $\gamma^{\prime}$ at $x=0$ and at $x=1$. The function $\gamma$ at $x=0$ makes a smooth transition from the value $u_{0}$ at $t=t_{0}$ to the value $u_{1}$ at $t=t_{\delta}$. We compute the transition phase for $\gamma^{\prime}$ by evaluating a cubic interpolating polynomial. For this purpose, the function subprogram CSDER, see Chapter 3, Interpolation and Approximation, is used. The interpolation is performed as a first step in the user-supplied routine FCNBC. The function and derivative values $\gamma\left(t_{0}\right)=u_{0}, \gamma^{\prime}\left(t_{0}\right)=0, \gamma\left(t_{\delta}\right)=u_{1}$, and $\gamma^{\prime}\left(t_{\delta}\right)=0$, are used as input to routine C2HER, to obtain the coefficients evaluated by CSDER. Notice that $\gamma^{\prime}(t)=0, t>t_{\delta}$. The evaluation routine CSDER will not yield this value so logic in the routine $\operatorname{FCNBC}$ assigns $\gamma^{\prime}(t)=0, t>t_{\delta}$.

```
    USE MOLCH INT
    USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
    USE AMACH_INT
    USE WRRRN INT
    IMPLICIT NONE
    PARAMETER (NPDES=1, NX=8, LDY=NPDES)
                                SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER I, IDO, J, NOUT, NSTEP
    REAL HINIT, PREC, T, TEND, TOL, XBREAK(NX), Y(LDY,NX), U0
    CHARACTER TITLE*19
    INTRINSIC FLOAT
    REAL FLOAT
    EXTERNAL FCNBC, FCNUT
    U0 = 1.0
    DO 10 I=1, NX
        XBREAK(I) = FLOAT(I-1)/(NX-1)
        Y(1,I) = U0
    10 CONTINUE
PREC = AMACH(4)
TOL = SQRT(PREC)
```

```
HINIT = 0.01*TOL
T = 0.0
IDO = 1
NSTEP = 10
CALL UMACH (2, NOUT)
J = 0
    20 CONTINUE
    J = J + 1
    TEND = FLOAT (J)/FLOAT (NSTEP)
                                    This puts more output for small
                                    t values where action is fastest.
    TEND = TEND**2
        Solve the problem
    CALL MOLCH (IDO, FCNUT, FCNBC, T, TEND, XBREAK, Y, TOL=TOL, &
        HINIT=HINIT)
    IF (J .LE. NSTEP) THEN
        Print results
        WRITE (TITLE,'(A,F4.2)') 'Solution at T =', T
        CALL WRRRN (TITLE, Y)
        Final call to release workspace
        IF (J .EQ. NSTEP) IDO = 3
        GO TO 20
    END IF
    END
    SUBROUTINE FCNUT (NPDES, X, T, U, UX, UXX, UT)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NPDES
    REAL X, T, U(*), UX(*), UXX(*), UT(*)
                                    Define the PDE
    UT(1) = UXX(1)
    RETURN
    END
SUBROUTINE FCNBC (NPDES, X, T, ALPHA, BTA, GAMP)
USE CSDER_INT
USE C2HER_INT
USE WRRRN INT
    INTEGER NPDES
REAL X, T, ALPHA(*), BTA(*), GAMP(*)
                                    SPECIFICATIONS FOR PARAMETERS
REAL TDELTA, U0, U1
PARAMETER (TDELTA=0.09, U0=1.0, U1=0.1)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER IWK(2), NDATA
    REAL DFDATA(2), FDATA(2), XDATA(2)
                                    SPECIFICATIONS FOR SAVE VARIABLES
    REAL BREAK(2), CSCOEF (4,2)
    LOGICAL FIRST
    SAVE BREAK, CSCOEF, FIRST
        SPECIFICATIONS FOR SUBROUTINES
    DATA FIRST/.TRUE./
    IF (FIRST) GO TO 20
```

```
! Define the boundary conditions
    IF (X .EQ. O.0) THEN
        ALPHA(1) = 1.0
        BTA(1) = 0.0
        GAMP(1) = 0.
                            If in the boundary layer,
                            compute nonzero gamma prime.
        IF (T .LE. TDELTA) GAMP(1) = CSDER(1,T,BREAK,CSCOEF)
        ELSE
    These are for x=1.
        ALPHA(1) = 0.0
        BTA(1) = 1.0
        GAMP(1) = 0.0
        END IF
        RETURN
        2 0 ~ C O N T I N U E ~
! = 2 Compute the boundary layer data.
        NDATA = 2
        XDATA(1) = 0.0
        XDATA(2) = TDELTA
        FDATA(1) = U0
        FDATA(2) = U1
        DFDATA(1) = 0.0
        DFDATA(2) = 0.0
! Do Hermite cubic interpolation.
        CALL C2HER (NDATA, XDATA, FDATA, DFDATA, BREAK, CSCOEF, IWK)
        FIRST = .FALSE.
        GO TO 10
        END
```


## Output

| 1 |  | 2 | Solution at $T=0.01$ |  |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.9 | 4 | 5 | 6 | 7 | 8 |  |  |
| 0.969 | 0.997 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |


| Solution at $\mathrm{T}=0.04$ |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 0.625 | 0.871 | 0.963 | 0.991 | 0.998 | 1.000 | 1.000 | 1.000 |


| 1 | 2 | Solu | $\text { on at } T$ | $\text { . } 09$ | 6 | 7 | 8 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0.0998 | 0.4603 | 0.7171 | 0.8673 | 0.9437 | 0.9781 | 0.9917 | 0.9951 |


| 1 | 2 | $\underset{3}{\text { Solut }}$ | $\begin{array}{r} \mathrm{n} \text { at } \begin{array}{r} \mathrm{T} \\ 4 \end{array} \mathbf{r} \end{array}$ | . 16 | 6 | 7 | 8 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0.0994 | 0.3127 | 0.5069 | 0.6680 | 0.7893 | 0.8708 | 0.9168 | 0.9316 |
| Solution at $T=0.25$ |  |  |  |  |  |  |  |
| 0.0994 | 0.2564 | 0.4043 | 0.5352 | 0.6428 | 0.7223 | 0.7709 | 0.7873 |



## Additonal Examples

## Example 2

In this example, using MOLCH, we solve the linear normalized diffusion PDE $u_{t}=u_{x x}$ but with an optional usage that provides values of the derivatives, $u_{x}$, of the initial data. Due to errors in the numerical derivatives computed by spline interpolation, more precise derivative values are required when the initial data is $u(x, 0)=1+\cos [(2 n-1) \pi x], n>1$. The boundary conditions are "zero flux" conditions $u_{x}(0, t)=u_{x}(1, t)=0$ for $t>0$. Note that the initial data is compatible with these end conditions since the derivative function

$$
u_{x}(x, 0)=\frac{d u(x, 0)}{d x}=-(2 n-1) \pi \sin [(2 n-1) \pi x]
$$

vanishes at $x=0$ and $x=1$.
The example illustrates the use of the IMSL options manager subprograms SUMAG or, for double precision, DUMAG, see Chapter 11, Utilities, to reset the array PARAM used for control of the specialized version of IVPAG that integrates the system of ODEs. This optional usage signals that the derivative of the initial data is passed by the user. The values $u(x$, tend $)$ and $u_{x}(x$, tend $)$ are output at the breakpoints with the optional usage.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDY, NPDES, NX, IAC
PARAMETER (NPDES=1, NX=10, LDY=NPDES)
SPECIFICATIONS FOR PARAMETERS
INTEGER ICHAP, IGET, IPUT, KPARAM
PARAMETER (ICHAP=5, IGET=1, IPUT=2, KPARAM=11)
```

$!$

```
    INTEGER I, IACT, IDO, IOPT(1), J, JGO, N, NOUT, NSTEP
    REAL ARG1, HINIT, PREC, PARAM(50), PI, T, TEND, TOL, &
            XBREAK (NX), Y(LDY,2*NX)
    CHARACTER TITLE*36
                                    SPECIFICATIONS FOR INTRINSICS
    INTRINSIC COS, FLOAT, SIN, SQRT
    REAL COS, FLOAT, SIN, SQRT
                                    SPECIFICATIONS FOR FUNCTIONS
                                    Set breakpoints and initial
                                    conditions.
    N = 5
    PI = CONST('pi')
    IOPT(1) = KPARAM
    DO 10 I=1, NX
        XBREAK(I) = FLOAT(I-1)/(NX-1)
        ARG1 = (2.*N-1)*PI
                                Set function values.
        Y(1,I) = 1. + COS(ARG1*XBREAK(I))
                                    Set first derivative values.
        Y(1,I+NX) = -ARG1*SIN(ARG1*XBREAK(I))
    1 0 \text { CONTINUE}
    PREC = AMACH (4)
    TOL = SQRT(PREC)
    HINIT = 0.01*TOL
    T = 0.0
    IDO = 1
    NSTEP = 10
    CALL UMACH (2, NOUT)
    J = 0
        Get and reset the PARAM array
        so that user-provided derivatives
                        of the initial data are used.
        JGO = 1
        IACT = IGET
        GO TO 70
    20 CONTINUE
        This flag signals that
        derivatives are passed.
        PARAM(17) = 1.
        JGO = 2
        IACT = IPUT
        GO TO 70
    3 0 ~ C O N T I N U E
        Look at output at steps
        of 0.001.
    TEND = 0.
    4 0 ~ C O N T I N U E
    J = J + 1
    TEND = TEND + 0.001
!
    CALL MOLCH (IDO, FCNUT, FCNBC, T, TEND, XBREAK, Y, NPDES=NPDES, &
        NX=NX, HINIT=HINIT, TOL=TOL)
    IF (J .LE. NSTEP) THEN
```

```
! Print results
        WRITE (TITLE,'(A,F5.3)') 'Solution and derivatives at T =', T
        CALL WRRRN (TITLE, Y)
        Final call to release workspace
        IF (J .EQ. NSTEP) IDO = 3
        GO TO 40
        END IF
                            Show, for example, the maximum
                                    step size used.
    JGO = 3
    IACT = IGET
    GO TO 70
    5 0 ~ C O N T I N U E ~
    WRITE (NOUT,*) ' Maximum step size used is: ', PARAM(33)
! Reset option to defaults
    JGO = 4
    IAC = IPUT
    IOPT(1) = -IOPT(1)
    GO TO 70
    6 0 ~ C O N T I N U E ~
        RETURN
Internal routine to work options
    7 0 ~ C O N T I N U E ~
    CALL SUMAG ('math', ICHAP, IACT, IOPT, PARAM, numopt=1)
    GO TO (20, 30, 50, 60), JGO
    END
    SUBROUTINE FCNUT (NPDES, X, T, U, UX, UXX, UT)
                                    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NPDES
    REAL X, T, U(*), UX(*), UXX(*), UT(*)
                            Define the PDE
    UT(1) = UXX(1)
    RETURN
    END
    SUBROUTINE FCNBC (NPDES, X, T, ALPHA, BTA, GAMP)
                                    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NPDES
    REAL X, T, ALPHA(*), BTA(*), GAMP (*)
    ALPHA(1) = 0.0
    BTA(1) = 1.0
    GAMP(1) = 0.0
        RETURN
    END
```


## Output

| Solution and derivatives at $\mathrm{T}=0.001$ |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 1.483 | 0.517 | 1.483 | 0.517 | 1.483 | 0.517 | 1.483 | 0.517 | 1.483 | 0.517 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |


| Solution and derivatives at $T=0.002$ |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 1.233 | 0.767 | 1.233 | 0.767 | 1.233 | 0.767 | 1.233 | 0.767 | 1.233 | 0.767 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $T=0.003$ |  |  |  |  |  |  |  |
| 1.113 | 0.887 | 1.113 | 0.887 | 1.113 | 0.887 | 1.113 | 0.887 | 1.113 | 0.887 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $\mathrm{T}=0.004$ |  |  |  |  |  |  |  |
| 1.054 | 0.946 | 1.054 | 0.946 | 1.054 | 0.946 | 1.054 | 0.946 | 1.054 | 0.946 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $T=0.005$ |  |  |  |  |  |  |  |
| 1.026 | 0.974 | 1.026 | 0.974 | 1.026 | 0.974 | 1.026 | 0.974 | 1.026 | 0.974 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $T=0.006$ |  |  |  |  |  |  |  |
| 1.012 | 0.988 | 1.012 | 0.988 | 1.012 | 0.988 | 1.012 | 0.988 | 1.012 | 0.988 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $\mathrm{T}=0.007$ |  |  |  |  |  |  |  |
| 1.006 | 0.994 | 1.006 | 0.994 | 1.006 | 0.994 | 1.006 | 0.994 | 1.006 | 0.994 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $\mathrm{T}=0.008$ |  |  |  |  |  |  |  |
| 1.003 | 0.997 | 1.003 | 0.997 | 1.003 | 0.997 | 1.003 | 0.997 | 1.003 | 0.997 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| 1 | 2 | Solution and derivatives at $\mathrm{T}=0.009$ |  |  |  |  |  |  |  |
| 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |


| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 2 | Solution and derivatives at $T=0.010$ |  |  |  |  |  |  |  |
| 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 | 1.001 | 0.999 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Maximum | step si | ze used | is: | 1.000 | 0E-02 |  |  |  |  |

## Example 3

In this example, we consider the linear normalized hyperbolic PDE, $u_{t t}=u_{x x}$, the "vibrating string" equation. This naturally leads to a system of first order PDEs. Define a new dependent variable $u_{t}=v$. Then, $v_{t}=u_{x x}$ is the second equation in the system. We take as initial data $u(x, 0)=\sin (\pi x)$ and $u_{t}(x, 0)=v(x, 0)=0$. The ends of the string are fixed so $u(0, t)=u(1, t)=v(0, t)=v(1, t)=0$. The exact solution to this problem is $u(x, t)=\sin (\pi x) \cos (\pi t)$. Residuals are computed at the output values of $t$ for $0<t \leq 2$. Output is obtained at 200 steps in increments of 0.01 .

Even though the sample code MOLCH gives satisfactory results for this PDE, users should be aware that for nonlinear problems, "shocks" can develop in the solution. The appearance of shocks may cause the code to fail in unpredictable ways. See Courant and Hilbert (1962), pages 488-490, for an introductory discussion of shocks in hyperbolic systems.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER LDY, NPDES, NX
PARAMETER (NPDES=2, NX=10, LDY=NPDES)
                            SPECIFICATIONS FOR PARAMETERS
INTEGER ICHAP, IGET, IPUT, KPARAM
PARAMETER (ICHAP=5, IGET=1, IPUT=2, KPARAM=11)
                            SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, IACT, IDO, IOPT(1), J, JGO, NOUT, NSTEP
REAL HINIT, PREC, PARAM(50), PI, T, TEND, TOL, XBREAK(NX), &
        Y(LDY,2*NX), ERROR(NX), ERRU
                            SPECIFICATIONS FOR INTRINSICS
INTRINSIC COS, FLOAT, SIN, SQRT
REAL COS, FLOAT, SIN, SQRT
                                    SPECIFICATIONS FOR SUBROUTINES
                                    SPECIFICATIONS FOR FUNCTIONS
EXTERNAL FCNBC, FCNUT
    Set breakpoints and initial
! conditions.
PI = CONST('pi')
IOPT(1) = KPARAM
DO 10 I=1, NX
        XBREAK(I) = FLOAT(I-1)/(NX-1)
        Y(1,I) = SIN(PI*XBREAK(I))
        Y(2,I) = 0.
                        Set first derivative values.
    Y(1,I+NX) = PI*COS(PI*XBREAK(I))
```

```
                Y(2,I+NX) = 0.0
    10 CONTINUE
!
    PREC = AMACH(4)
    TOL = 0.1*SQRT(PREC)
    HINIT = 0.01*TOL
    T = 0.0
    IDO = 1
    NSTEP = 200
    CALL UMACH (2, NOUT)
    J = 0
! Get and reset the PARAM array
    so that user-provided derivatives
    of the initial data are used.
    JGO = 1
    IACT = IGET
    GO TO 90
    20 CONTINUE
!
    PARAM(17) = 1.
    JGO = 2
    IACT = IPUT
    GO TO 90
    3 0 ~ C O N T I N U E
!
    ERRU = 0.
    TEND = 0.
    4 0 ~ C O N T I N U E ~
    J= J + 1
    TEND = TEND + 0.01
!
    CALL MOLCH (IDO, FCNUT, FCNBC, T, TEND, XBREAK, Y, NX=NX, &
                    HINIT=HINIT, TOL=TOL)
    DO 50 I=1, NX
        ERROR(I) = Y(1,I) - SIN(PI*XBREAK(I))*COS(PI*TEND)
    5 0 ~ C O N T I N U E
    IF (J .LE. NSTEP) THEN
        DO 60 I=1, NX
            ERRU = AMAX1 (ERRU,ABS (ERROR(I)))
    6 0 ~ C O N T I N U E ~
! Final call to release workspace
            IF (J .EQ. NSTEP) IDO = 3
            GO TO 40
        END IF
                            Show, for example, the maximum
                            step size used.
    JGO = 3
    IACT = IGET
    GO TO 90
    7 0 ~ C O N T I N U E
    WRITE (NOUT,*) ' Maximum error in u(x,t) divided by TOL: ', &
                ERRU/TOL
    WRITE (NOUT,*) ' Maximum step size used is: ', PARAM(33)
```

```
! Reset option to defaults
    JGO = 4
    IACT = IPUT
    IOPT(1) = -IOPT(1)
    GO TO 90
    8 0 ~ C O N T I N U E ~
    RETURN
    Internal routine to work options
    9 0 ~ C O N T I N U E ~
    CALL SUMAG ('math', ICHAP, IACT, IOPT, PARAM)
    GO TO (20, 30, 70, 80), JGO
    END
    SUBROUTINE FCNUT (NPDES, X, T, U, UX, UXX, UT)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NPDES
    REAL X, T, U(*), UX(*), UXX(*), UT(*)
                                    Define the PDE
    UT(1) = U(2)
    UT(2) = UXX(1)
! RETURN
    END
    SUBROUTINE FCNBC (NPDES, X, T, ALPHA, BTA, GAMP)
    SPECIFICATIONS FOR ARGUMENTS
    INTEGER NPDES
    REAL X, T, ALPHA(*), BTA(*), GAMP(*)
    ALPHA(1) = 1.0
    BTA(1) = 0.0
    GAMP(1) = 0.0
    ALPHA(2) = 1.0
    BTA(2) = 0.0
    GAMP (2) = 0.0
! RETURN
    END
```


## Output

Maximum error in $u(x, t)$ divided by TOL: 1.28094
Maximum step size used is: 9.99999E-02

## FPS2H

Solves Poisson's or Helmholtz's equation on a two-dimensional rectangle using a fast Poisson solver based on the HODIE finite-difference scheme on a uniform mesh.

## Required Arguments

PRHS - User-supplied FUNCTION to evaluate the right side of the partial differential equation. The form is $\operatorname{PRHS}(\mathrm{X}, \mathrm{Y})$, where
$X-X$-coordinate value. (Input)
$Y-Y$-coordinate value. (Input)
PRHS - Value of the right side at ( $\mathrm{X}, \mathrm{Y}$ ). (Output)
PRHS must be declared EXTERNAL in the calling program.
BRHS - User-supplied FUNCTION to evaluate the right side of the boundary conditions. The form is BRHS(ISIDE, $\mathrm{X}, \mathrm{Y}$ ), where

ISIDE - Side number. (Input)
See IвСТч below for the definition of the side numbers.
$X-X$-coordinate value. (Input)
$\mathrm{Y}-\mathrm{Y}$-coordinate value. (Input)
BRHS - Value of the right side of the boundary condition at ( $\mathrm{X}, \mathrm{Y}$ ). (Output)
BRHS must be declared EXTERNAL in the calling program.
COEFU - Value of the coefficient of U in the differential equation. (Input)
$N X$ - Number of grid lines in the X -direction. (Input)
NX must be at least 4. See Comment 2 for further restrictions on NX.
$N Y$ - Number of grid lines in the Y-direction. (Input)
NY must be at least 4 . See Comment 2 for further restrictions on NY.
$\boldsymbol{A X}$ - The value of x along the left side of the domain. (Input)
$\boldsymbol{B} \boldsymbol{X}$ - The value of x along the right side of the domain. (Input)
$\boldsymbol{A Y}$ - The value of Y along the bottom of the domain. (Input)
$\boldsymbol{B Y}$ - The value of Y along the top of the domain. (Input)
IBCTY - Array of size 4 indicating the type of boundary condition on each side of the domain or that the solution is periodic. (Input)
The sides are numbered 1 to 4 as follows:

## Side Location

1 - Right $\quad(X=B X)$
2 - Bottom $\quad(Y=A Y)$
3 - Left
$(X=A X)$

4 - Top
$(\mathrm{Y}=\mathrm{BY})$
There are three boundary condition types.

IBCTY
$1 \quad$ Value of u is given. (Dirichlet)
2 Value of dU/dX is given (sides 1 and/or 3). (Neumann) Value of dU/dY is given (sides 2 and/or 4).

3
Periodic.
$\boldsymbol{U}$ - Array of size NX by NY containing the solution at the grid points. (Output)

## Optional Arguments

IORDER - Order of accuracy of the finite-difference approximation. (Input)
It can be either 2 or 4 . Usually, IORDER $=4$ is used.
Default: $\operatorname{IORDER}=4$.
$\boldsymbol{L D U}$ - Leading dimension of U exactly as specified in the dimension statement of the calling program. (Input)
Default: LDU = size ( $\mathrm{U}, 1$ ).

## FORTRAN 90 Interface

Generic: CALL FPS2H (PRHS, BRHS, COEFU, NX, NY, AX, BX, AY, BY, IBCTY, U [,...])

Specific: The specific interface names are S_FPS2H and D_FPS2H.

## FORTRAN 77 Interface

Single: CALL FPS2H (PRHS, BRHS, COEFU, NX, NY, AX, BX, AY, BY, IBCTY, IORDER, U, LDU)

Double: The double precision name is DFPS2H.

## Description

Let $c=\operatorname{COEFU}, a_{x}=\mathrm{AX}, b_{x}=\mathrm{BX}, a_{y}=\mathrm{AY}, b_{y}=\mathrm{BY}, n_{x}=\mathrm{NX}$ and $n_{y}=\mathrm{NY}$.
FPS2H is based on the code HFFT2D by Boisvert (1984). It solves the equation

$$
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}+c u=p
$$

on the rectangular domain $\left(a_{x}, b_{x}\right) \times\left(a_{y}, b_{y}\right)$ with a user-specified combination of Dirichlet (solution prescribed), Neumann (first-derivative prescribed), or periodic boundary conditions. The sides are numbered clockwise, starting with the right side.


When $c=0$ and only Neumann or periodic boundary conditions are prescribed, then any constant may be added to the solution to obtain another solution to the problem. In this case, the solution of minimum $\infty$-norm is returned.

The solution is computed using either a second-or fourth-order accurate finite-difference approximation of the continuous equation. The resulting system of linear algebraic equations is solved using fast Fourier transform techniques. The algorithm relies upon the fact that $n_{x}-1$ is highly composite (the product of small primes). For details of the algorithm, see Boisvert (1984). If $n_{x}-1$ is highly composite then the execution time of FPS2H is proportional to $n_{x} n_{y} \log _{2} n_{x}$. If evaluations of $p(x, y)$ are inexpensive, then the difference in running time between IORDER $=2$ and $\operatorname{IORDER}=4$ is small.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2S2H/DF2S2H. The reference is:

CALL F2S2H (PRHS, BRHS, COEFU, NX, NY, AX, BX, AY, BY, IBCTY, IORDER, U, LDU, UWORK, WORK)

The additional arguments are as follows:
$\boldsymbol{U W O R K}-$ Work array of size NX +2 by NY +2 . If the actual dimensions of U are large enough, then U and UWORK can be the same array.

WORK - Work array of length $(N X+1)(N Y+1)($ IORDER -2$) / 2+$ $6(N X+N Y)+N X / 2+16$.
2. The grid spacing is the distance between the (uniformly spaced) grid lines. It is given by the formulas $H X=(B X-A X) /(N X-1)$ and $H Y=(B Y-A Y) /(N Y-1)$. The grid spacings in the $X$ and $Y$ directions must be the same, i.e., NX and NY must be such that

HX equals HY. Also, as noted above, NX and NY must both be at least 4. To increase the speed of the fast Fourier transform, NX -1 should be the product of small primes. Good choices are 17,33 , and 65 .
3. If-COEFU is nearly equal to an eigenvalue of the Laplacian with homogeneous boundary conditions, then the computed solution might have large errors.

## Example

In this example, the equation

$$
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}+3 u=-2 \sin (x+2 y)+16 e^{2 x+3 y}
$$

with the boundary conditions $\partial u / \partial y=2 \cos (x+2 y)+3 \exp (2 x+3 y)$ on the bottom side and $u=\sin (x+2 y)+\exp (2 x+3 y)$ on the other three sides. The domain is the rectangle $[0,1 / 4] \times[0$, $1 / 2$ ]. The output of FPS 2 H is a $17 \times 33$ table of $U$ values. The quadratic interpolation routine QD2VL is used to print a table of values.

```
    USE FPS2H_INT
    USE QD2VL_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER NCVAL, NX, NXTABL, NY, NYTABL
    PARAMETER (NCVAL=11, NX=17, NXTABL=5, NY=33, NYTABL=5)
    INTEGER I, IBCTY(4), IORDER, J, NOUT
    REAL AX, AY, BRHS, BX, BY, COEFU, ERROR, FLOAT, PRHS, &
        TRUE, U(NX,NY), UTABL, X, XDATA(NX), Y, YDATA(NY)
    INTRINSIC FLOAT
    EXTERNAL BRHS, PRHS
    AX = 0.0
    BX = 0.25
    AY = 0.0
    BY = 0.50
! Set boundary condition types
    IBCTY(1) = 1
    IBCTY(2) = 2
    IBCTY(3) = 1
    IBCTY(4) = 1
    COEFU = 3.0
    IORDER = 4
    CALL FPS2H (PRHS, BRHS, COEFU, NX, NY, AX, BX, AY, BY, IBCTY, U)
    Setup for quadratic interpolation
    DO 10 I=1, NX
        XDATA(I) = AX + (BX-AX)*FLOAT (I-1)/FLOAT (NX-1)
    1 0 ~ C O N T I N U E ~
```

```
        DO 20 J=1, NY
            YDATA(J) = AY + (BY-AY)*FLOAT (J-1)/FLOAT (NY-1)
    20 CONTINUE
    CALL UMACH (2, NOUT)
    WRITE (NOUT,'(8X,A,11X,A,11X,A,8X,A)') 'X', 'Y', 'U', 'Error'
    DO 40 J=1, NYTABL
        DO 30 I=1, NXTABL
            X = AX + (BX-AX)*FLOAT (I-1)/FLOAT (NXTABL-1)
            Y = AY + (BY-AY)*FLOAT (J-1)/FLOAT (NYTABL-1)
            UTABL = QD2VL(X,Y,XDATA,YDATA,U)
            TRUE = SIN (X+2.*Y) + EXP(2.*X+3.*Y)
            ERROR = TRUE - UTABL
            WRITE (NOUT,'(4F12.4)') X, Y, UTABL, ERROR
    3 0 ~ C O N T I N U E ~
    4 0 ~ C O N T I N U E
        END
!
        REAL FUNCTION PRHS (X, Y)
        REAL X, Y
!
        REAL EXP, SIN
        INTRINSIC EXP, SIN
        PRHS = -2.*SIN (X+2.*Y) + 16.*EXP(2.*X+3.*Y)
        RETURN
        END
!
    REAL FUNCTION BRHS (ISIDE, X, Y)
    INTEGER ISIDE
    REAL X, Y
    REAL COS, EXP, SIN
    INTRINSIC COS, EXP, SIN
                                    Define the boundary conditions
    IF (ISIDE .EQ. 2) THEN
        BRHS = 2.* COS (X+2.*Y) + 3.*EXP (2.*X+3.*Y)
    ELSE
        BRHS = SIN (X+2.*Y) + EXP(2.*X+3.*Y)
    END IF
    RETURN
    END
```


## Output

| X | Y | U | Error |
| :---: | :---: | :---: | :---: |
| 0.0000 | 0.0000 | 1.0000 | 0.0000 |
| 0.0625 | 0.0000 | 1.1956 | 0.0000 |
| 0.1250 | 0.0000 | 1.4087 | 0.0000 |
| 0.1875 | 0.0000 | 1.6414 | 0.0000 |
| 0.2500 | 0.0000 | 1.8961 | 0.0000 |
| 0.0000 | 0.1250 | 1.7024 | 0.0000 |
| 0.0625 | 0.1250 | 1.9562 | 0.0000 |
| 0.1250 | 0.1250 | 2.2345 | 0.0000 |
| 0.1875 | 0.1250 | 2.5407 | 0.0000 |


| 0.2500 | 0.1250 | 2.8783 | 0.0000 |
| :--- | :--- | :--- | :--- |
| 0.0000 | 0.2500 | 2.5964 | 0.0000 |
| 0.0625 | 0.2500 | 2.9322 | 0.0000 |
| 0.1250 | 0.2500 | 3.3034 | 0.0000 |
| 0.1875 | 0.2500 | 3.7148 | 0.0000 |
| 0.2500 | 0.2500 | 4.1720 | 0.0000 |
| 0.0000 | 0.3750 | 3.7619 | 0.0000 |
| 0.0625 | 0.3750 | 4.2163 | 0.0000 |
| 0.1250 | 0.3750 | 4.7226 | 0.0000 |
| 0.1875 | 0.3750 | 5.2878 | 0.0000 |
| 0.2500 | 0.3750 | 5.9199 | 0.0000 |
| 0.0000 | 0.5000 | 5.3232 | 0.0000 |
| 0.0625 | 0.5000 | 5.9520 | 0.0000 |
| 0.1250 | 0.5000 | 6.6569 | 0.0000 |
| 0.1875 | 0.5000 | 7.4483 | 0.0000 |
| 0.2500 | 0.5000 | 8.3380 | 0.0000 |

## FPS3H

Solves Poisson's or Helmholtz's equation on a three-dimensional box using a fast Poisson solver based on the HODIE finite-difference scheme on a uniform mesh.

## Required Arguments

PRHS - User-supplied FUNCTION to evaluate the right side of the partial differential equation. The form is PRHS ( $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ ), where

X - The $x$-coordinate value. (Input)
Y - The $y$-coordinate value. (Input)
Z - The $z$-coordinate value. (Input)
PRHS - Value of the right side at ( $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ ). (Output)
PRHS must be declared EXTERNAL in the calling program.
BRHS - User-supplied FUNCTION to evaluate the right side of the boundary conditions. The form is BRHS (ISIDE, $X, Y, Z$ ), where

ISIDE - Side number. (Input)
See IBCTY for the definition of the side numbers.
X - The $x$-coordinate value. (Input)
Y - The $y$-coordinate value. (Input)
Z - The $z$-coordinate value. (Input)
BRHS - Value of the right side of the boundary condition at ( $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ ). (Output)
BRHS must be declared EXTERNAL in the calling program.
COEFU - Value of the coefficient of U in the differential equation. (Input)
$\boldsymbol{N X}$ - Number of grid lines in the $x$-direction. (Input)
NX must be at least 4. See Comment 2 for further restrictions on NX.
$\boldsymbol{N Y}$ - Number of grid lines in the $y$-direction. (Input)
NY must be at least 4 . See Comment 2 for further restrictions on NY.
$N Z$ - Number of grid lines in the $y$-direction. (Input)
NZ must be at least 4 . See Comment 2 for further restrictions on NZ.
$\boldsymbol{A X}$ - Value of x along the left side of the domain. (Input)
$\boldsymbol{B} \boldsymbol{X}-$ Value of x along the right side of the domain. (Input)
$\boldsymbol{A Y}$ - Value of Y along the bottom of the domain. (Input)
$\boldsymbol{B Y}$ - Value of Y along the top of the domain. (Input)
$\boldsymbol{A Z}$ - Value of z along the front of the domain. (Input)
$\boldsymbol{B Z}$ - Value of z along the back of the domain. (Input)
IBCTY - Array of size 6 indicating the type of boundary condition on each face of the domain or that the solution is periodic. (Input)
The sides are numbers 1 to 6 as follows:

Side

## Location

1 - Right

$$
(X=B X)
$$

2 - Bottom

$$
(Y=A Y)
$$

3 - Left
$(X=A X)$

4 - Top
$(Y=B Y)$
5 - Front
$(Z=B Z)$
6 - Back
( $Z=A Z$ )

There are three boundary condition types.

## IBCTY Boundary Condition

$1 \quad$ Value of U is given. (Dirichlet)
2 Value of dU/dX is given (sides 1 and/or 3). (Neumann) Value of dU/dY is given (sides $2 \mathrm{and} /$ or 4 ). Value of $\mathrm{dU} / \mathrm{dZ}$ is given (sides $5 \mathrm{and} /$ or 6 ).
$\boldsymbol{U}$ - Array of size NX by NY by NZ containing the solution at the grid points. (Output)

## Optional Arguments

IORDER - Order of accuracy of the finite-difference approximation. (Input)
It can be either 2 or 4 . Usually, IORDER $=4$ is used.
Default: $\operatorname{IORDER}=4$.
$\boldsymbol{L D U}$ - Leading dimension of U exactly as specified in the dimension statement of the calling program. (Input)
Default: LDU $=\operatorname{size}(\mathrm{U}, 1)$.
MDU - Middle dimension of $U$ exactly as specified in the dimension statement of the calling program. (Input)
Default: MDU = size (U,2).

## FORTRAN 90 Interface

Generic: CALL FPS3H (PRHS, BRHS, COEFU, NX, NY, NZ, AX, BX, AY, BY, AZ, BZ, IBCTY, U [,...])

Specific: The specific interface names are S_FPS3H and D_FPS3H.

## FORTRAN 77 Interface

Single: CALL FPS3H (PRHS, BRHS, COEFU, NX, NY, NZ, AX, BX, AY, BY, AZ, BZ, IBCTY, IORDER, U, LDU, MDU)

Double: The double precision name is DFPS3H.

## Description

$$
\text { Let } c=\mathrm{COEFU}, a_{x}=\mathrm{AX}, b_{x}=\mathrm{BX}, n_{x}=\mathrm{NX}, a_{y}=\mathrm{AY}, b_{y}=\mathrm{BY}, n_{y}=\mathrm{NY}, a_{z}=\mathrm{AZ}, b_{z}=\mathrm{BZ} \text {, and } n_{z}=\mathrm{NZ} \text {. }
$$

FPS3H is based on the code HFFT3D by Boisvert (1984). It solves the equation

$$
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}+\frac{\partial^{2} u}{\partial z^{2}}+c u=p
$$

on the domain $\left(a_{x}, b_{x}\right) \times\left(a_{y}, b_{y}\right) \times\left(a_{z}, b_{z}\right)$ (a box) with a user-specified combination of Dirichlet (solution prescribed), Neumann (first derivative prescribed), or periodic boundary conditions. The six sides are numbered as shown in the following diagram.


When $c=0$ and only Neumann or periodic boundary conditions are prescribed, then any constant may be added to the solution to obtain another solution to the problem. In this case, the solution of minimum $\infty$-norm is returned.

The solution is computed using either a second-or fourth-order accurate finite-difference approximation of the continuous equation. The resulting system of linear algebraic equations is solved using fast Fourier transform techniques. The algorithm relies upon the fact that $n_{x}-1$ and $n_{z}-1$ are highly composite (the product of small primes). For details of the algorithm, see Boisvert (1984). If $n_{x}-1$ and $n_{z}-1$ are highly composite, then the execution time of FPS3H is proportional to

$$
n_{x} n_{y} n_{z}\left(\log _{2}^{2} n_{x}+\log _{2}^{2} n_{z}\right)
$$

If evaluations of $p(x, y, z)$ are inexpensive, then the difference in running time between IORDER $=2$ and IORDER $=4$ is small.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2S3H/DF2S3H. The reference is:

CALL F2S3H (PRHS, BRHS, COEFU, NX, NY, NZ, AX, BX, AY, BY, AZ, BZ, IBCTY, IORDER, U, LDU, MDU, UWORK, WORK)

The additional arguments are as follows:
$\boldsymbol{U W O R K}-$ Work array of size $N X+2$ by NY +2 by NZ +2 . If the actual dimensions of $U$ are large enough, then $U$ and UWORK can be the same array.

```
WORK - Work array of length (NX + 1) (NY + 1) (NZ + 1)(IORDER - 2)/2 +
    2(NX * NY + NX * NZ + NY * NZ) + 2(NX + NY + 1) +
    MAX(2 * NX * NY, 2 * NX + NY + 4 * NZ + (NX + NZ)/2 + 29)
```

2. The grid spacing is the distance between the (uniformly spaced) grid lines. It is given by the formulas
$H X=(B X-A X) /(N X-1)$,
$H Y=(B Y-A Y) /(N Y-1)$, and
$H Z=(B Z-A Z) /(N Z-1)$.
The grid spacings in the $X, Y$ and $Z$ directions must be the same, i.e., $N X, N Y$ and $N Z$ must be such that $H X=H Y=H Z$. Also, as noted above, $N X, N Y$ and $N Z$ must all be at least 4. To increase the speed of the Fast Fourier transform, NX -1 and nZ -1 should be the product of small primes. Good choices for NX and NZ are 17, 33 and 65.
3. If-COEFU is nearly equal to an eigenvalue of the Laplacian with homogeneous boundary conditions, then the computed solution might have large errors.

## Example

This example solves the equation

$$
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}+\frac{\partial^{2} u}{\partial z^{2}}+10 u=-4 \cos (3 x+y-2 z)+12 e^{x-z}+10
$$

with the boundary conditions $\partial u / \partial z=-2 \sin (3 x+y-2 z)-\exp (x-z)$ on the front side and $u=\cos (3 x+y-2 z)+\exp (x-z)+1$ on the other five sides. The domain is the box $[0,1 / 4] \times[0$, $1 / 2] \times[0,1 / 2]$. The output of FPS3H is a $9 \times 17 \times 17$ table of $U$ values. The quadratic interpolation routine $Q D 3 V L$ is used to print a table of values.

```
USE FPS3H INT
USE UMACH_INT
USE QD3VL_INT
IMPLICIT NONE
INTEGER LDU, MDU, NX, NXTABL, NY, NYTABL, NZ, NZTABL
PARAMETER (NX=5, NXTABL=4, NY=9, NYTABL=3, NZ=9, &
    NZTABL=3, LDU=NX, MDU=NY)
INTEGER I, IBCTY(6), IORDER, J, K, NOUT
REAL AX, AY, AZ, BRHS, BX, BY, BZ, COEFU, FLOAT, PRHS, &
    U(LDU,MDU,NZ), UTABL, X, ERROR, TRUE, &
    XDATA(NX), Y, YDATA(NY), Z, ZDATA(NZ)
INTRINSIC COS, EXP, FLOAT
EXTERNAL BRHS, PRHS
                                    Define domain
AX = 0.0
BX = 0.125
AY = 0.0
BY = 0.25
AZ = 0.0
```

$!$

```
    BZ = 0.25
! Set boundary condition types
    IBCTY(1) = 1
    IBCTY(2) = 1
    IBCTY(3) = 1
    IBCTY(4) = 1
    IBCTY(5) = 2
    IBCTY(6) = 1
    COEFU = 10.0
    IORDER = 4
    CALL FPS3H (PRHS, BRHS, COEFU, NX, NY, NZ, AX, BX, AY, BY, AZ, &
                BZ, IBCTY, U)
    DO 10 I=1, NX
        XDATA(I) = AX + (BX-AX)*FLOAT(I-1)/FLOAT (NX-1)
    1 0 ~ C O N T I N U E ~
    DO 20 J=1, NY
        YDATA(J) = AY + (BY-AY)*FLOAT(J-1)/FLOAT (NY-1)
    2 0 ~ C O N T I N U E ~
    DO 30 K=1, NZ
        ZDATA (K) = AZ + (BZ-AZ)*FLOAT (K-1)/FLOAT (NZ-1)
    3 0
!
    CONTINUE
    CALL UMACH (2, NOUT)
        WRITE (NOUT,'(8X,5(A,11X))') 'X', 'Y', 'Z', 'U', 'Error'
        DO 60 K=1, NZTABL
        DO 50 J=1, NYTABL
            DO 40 I=1, NXTABL
                    X = AX + (BX-AX)*FLOAT (I-1)/FLOAT (NXTABL-1)
                    Y = AY + (BY-AY)*FLOAT (J-1)/FLOAT (NYTABL-1)
                    Z = AZ + (BZ-AZ)*FLOAT (K-1)/FLOAT (NZTABL-1)
                    UTABL = QD3VL(X,Y,Z,XDATA,YDATA, ZDATA,U, CHECK=.false.)
                    TRUE = COS(3.0*X+Y-2.0*Z) + EXP(X-Z) + 1.0
                    ERROR = UTABL - TRUE
                    WRITE (NOUT,'(5F12.4)') X, Y, Z, UTABL, ERROR
                CONTINUE
    40 CONTIN
    6 0 ~ C O N T I N U E ~
    END
!
    REAL FUNCTION PRHS (X, Y, Z)
    REAL X, Y, Z
    REAL COS, EXP
    INTRINSIC COS, EXP
    PRHS = -4.0*COS(3.0*X+Y-2.0*Z) + 12*EXP(X-Z) + 10.0
    RETURN
    END
    REAL FUNCTION BRHS (ISIDE, X, Y, Z)
    INTEGER ISIDE
```



## Output

| X | Y | Z | U | Error |
| :---: | :---: | :---: | ---: | ---: |
| 0.0000 | 0.0000 | 0.0000 | 3.0000 | 0.0000 |
| 0.0417 | 0.0000 | 0.0000 | 3.0348 | 0.0000 |
| 0.0833 | 0.0000 | 0.0000 | 3.0558 | 0.0001 |
| 0.1250 | 0.0000 | 0.0000 | 3.0637 | 0.0001 |
| 0.0000 | 0.1250 | 0.0000 | 2.9922 | 0.0000 |
| 0.0417 | 0.1250 | 0.0000 | 3.0115 | 0.0000 |
| 0.0833 | 0.1250 | 0.0000 | 3.0175 | 0.0000 |
| 0.1250 | 0.1250 | 0.0000 | 3.0107 | 0.0000 |
| 0.0000 | 0.2500 | 0.0000 | 2.9690 | 0.0001 |
| 0.0417 | 0.2500 | 0.0000 | 2.9731 | 0.0000 |
| 0.0833 | 0.2500 | 0.0000 | 2.9645 | 0.0000 |
| 0.1250 | 0.2500 | 0.0000 | 2.9440 | -0.0001 |
| 0.0000 | 0.0000 | 0.1250 | 2.8514 | 0.0000 |
| 0.0417 | 0.0000 | 0.1250 | 2.9123 | 0.0000 |
| 0.0833 | 0.0000 | 0.1250 | 2.9592 | 0.0000 |
| 0.1250 | 0.0000 | 0.1250 | 2.9922 | 0.0000 |
| 0.0000 | 0.1250 | 0.1250 | 2.8747 | 0.0000 |
| 0.0417 | 0.1250 | 0.1250 | 2.9211 | 0.0010 |
| 0.0833 | 0.1250 | 0.1250 | 2.9524 | 0.0010 |
| 0.1250 | 0.1250 | 0.1250 | 2.9689 | 0.0000 |
| 0.0000 | 0.2500 | 0.1250 | 2.8825 | 0.0000 |
| 0.0417 | 0.2500 | 0.1250 | 2.9123 | 0.0000 |
| 0.0833 | 0.2500 | 0.1250 | 2.9281 | 0.0000 |
| 0.1250 | 0.2500 | 0.1250 | 2.9305 | 0.0000 |
| 0.0000 | 0.0000 | 0.2500 | 2.6314 | -0.0249 |
| 0.0417 | 0.0000 | 0.2500 | 2.7420 | -0.0004 |
| 0.0833 | 0.0000 | 0.2500 | 2.8112 | -0.0042 |
| 0.1250 | 0.0000 | 0.2500 | 2.8609 | -0.0138 |
| 0.0000 | 0.1250 | 0.2500 | 2.7093 | 0.0000 |
| 0.0417 | 0.1250 | 0.2500 | 2.8153 | 0.0344 |
| 0.0833 | 0.1250 | 0.2500 | 2.8628 | 0.0237 |
| 0.1250 | 0.1250 | 0.2500 | 2.8825 | 0.0000 |
| 0.0000 | 0.2500 | 0.2500 | 2.7351 | -0.0127 |
| 0.0417 | 0.2500 | 0.2500 | 2.8030 | -0.0011 |
| 0.0833 | 0.2500 | 0.2500 | 2.8424 | -0.0040 |
| 0.1250 | 0.2500 | 0.2500 | 2.8735 | -0.0012 |
| 0 |  |  |  | 0 |

## SLEIG

Determines eigenvalues, eigenfunctions and/or spectral density functions for Sturm-Liouville problems in the form

$$
-\frac{d}{d x}\left(p(x) \frac{d u}{d x}\right)+q(x) u=\lambda r(x) u \text { for } x \text { in }(a, b)
$$

with boundary conditions (at regular points)

$$
\begin{gathered}
a_{1} u-a_{2}\left(p u^{\prime}\right)=\lambda\left(a_{1}^{\prime} u-a_{2}^{\prime}\left(p u^{\prime}\right)\right) \text { at } a \\
b_{1} u+b_{2}\left(p u^{\prime}\right)=0 \text { at } b
\end{gathered}
$$

## Required Arguments

CONS - Array of size eight containing

$$
a_{1}, a_{1}^{\prime}, a_{2}, a_{2}^{\prime}, b_{1}, b_{2}, a \text { and } b
$$

in locations $\operatorname{CONS}(1)$ through $\operatorname{CONS}(8)$, respectively. (Input)
COEFFN - User-supplied SUBROUTINE to evaluate the coefficient functions. The usage is
CALL COEFFN ( $\mathrm{X}, \mathrm{PX}, \mathrm{QX}, \mathrm{RX}$ )
$X$ - Independent variable. (Input)
PX - The value of $p(x)$ at X . (Output)
QX - The value of $q(x)$ at X . (Output)
RX - The value of $r(x)$ at X . (Output)
COEFFN must be declared EXTERNAL in the calling program.
ENDFIN - Logical array of size two. ENDFIN (1) = .true. if the endpoint $a$ is finite. $\operatorname{ENDFIN}(2)=$. true. if endpoint $b$ is finite. (Input)

INDEX - Vector of size NUMEIG containing the indices of the desired eigenvalues. (Input)
$\boldsymbol{E V A L}$ - Array of length NUMEIG containing the computed approximations to the eigenvalues whose indices are specified in INDEX. (Output)

## Optional Arguments

NUMEIG - The number of eigenvalues desired. (Input) Default: NUMEIG = size (INDEX, 1 ).

TEVLAB - Absolute error tolerance for eigenvalues. (Input)
Default: TEVLAB $=10$.* machine precision.
TEVLRL - Relative error tolerance for eigenvalues. (Input)
Default: TEVLRL $=\operatorname{SQRT}$ (machine precision).

## FORTRAN 90 Interface

Generic: CALL SLEIG (CONS, COEFFN, ENDFIN, INDEX, EVAL [,...])
Specific: The specific interface names are S_SLEIG and D_SLEIG.

## FORTRAN 77 Interface

Single: CALL SLEIG (CONS, COEFFN, ENDFIN, NUMEIG, INDEX, TEVLAB, TEVLRL, EVAL)

Double: The double precision name is DSLEIG.

## Description

This subroutine is designed for the calculation of eigenvalues, eigenfunctions and/or spectral density functions for Sturm-Liouville problems in the form

$$
-\frac{d}{d x}\left(p(x) \frac{d u}{d x}\right)+q(x) u=\lambda r(x) u \text { for } x \text { in }(a, b)(1)
$$

with boundary conditions (at regular points)

$$
\begin{gathered}
a_{1} u-a_{2}\left(p u^{\prime}\right)=\lambda\left(a_{1}^{\prime} u-a_{2}^{\prime}\left(p u^{\prime}\right)\right) \text { at } a \\
b_{1} u+b_{2}\left(p u^{\prime}\right)=0 \text { at } b
\end{gathered}
$$

We assume that

$$
a_{1}^{\prime} a_{2}-a_{1} a_{2}^{\prime}>0
$$

when $a^{\prime}{ }_{1} \neq 0$ and $a_{2} \neq 0$. The problem is considered regular if and only if

- $\quad a$ and $b$ are finite,
- $\quad p(x)$ and $r(x)$ are positive in $(a, b)$,
- $\quad 1 / p(x), q(x)$ and $r(x)$ are locally integrable near the endpoints.

Otherwise the problem is called singular. The theory assumes that $p, p^{\prime}, q$, and $r$ are at least continuous on $(a, b)$, though a finite number of jump discontinuities can be handled by suitably defining an input mesh.
For regular problems, there are an infinite number of eigenvalues

$$
\lambda_{0}<\lambda_{1}<\ldots<\lambda_{k}, k \rightarrow \infty
$$

Each eigenvalue has an associated eigenfunction which is unique up to a constant. For singular problems, there is a wide range in the behavior of the eigenvalues.
As presented in Pruess and Fulton (1993) the approach is to replace (1) by a new problem

$$
\begin{equation*}
-\left(\hat{p} \hat{u}^{\prime}\right)^{\prime}+\hat{q} \hat{u}=\hat{\lambda} \hat{r} \hat{u} \tag{2}
\end{equation*}
$$

with analogous boundary conditions

$$
\begin{gathered}
a_{1} \hat{u}(a)-a_{2}\left(\hat{p} \hat{u}^{\prime}\right)(a)=\hat{\lambda}\left[a_{1}^{\prime} \hat{u}(a)-a_{2}^{\prime}\left(\hat{p} \hat{u}^{\prime}\right)(a)\right] \\
b_{1} \hat{u}(b)+b_{2}\left(\hat{p} \hat{u}^{\prime}\right)(b)=0
\end{gathered}
$$

where

$$
\hat{p}, \hat{q} \text { and } \hat{r}
$$

are step function approximations to $p, q$, and $r$, respectively. Given the mesh $a=x_{1}<x_{2}<\ldots<x_{N+1}=b$, the usual choice for the step functions uses midpoint interpolation, i. e.,

$$
\hat{p}(x)=p_{n} \equiv p\left(\frac{x_{n}+x_{n+1}}{2}\right)
$$

for $x$ in $\left(x_{n}, x_{n+1}\right)$ and similarly for the other coefficient functions. This choice works well for regular problems. Some singular problems require a more sophisticated technique to capture the asymptotic behavior. For the midpoint interpolants, the differential equation (2) has the known closed form solution in
$\left(x_{n}, x_{n+1}\right)$

$$
\hat{u}(x)=\hat{u}\left(x_{n}\right) \phi_{n}^{\prime}\left(x-x_{n}\right)+\left(\hat{p} \hat{u}^{\prime}\right)\left(x_{n}\right) \phi_{n}\left(x-x_{n}\right) / p_{n}
$$

with

$$
\phi_{n}(t)=\left\{\begin{array}{l}
\sin \omega_{n} t / \omega_{n}, \tau_{n}>0 \\
\sinh \omega_{n} t / \omega_{n}, \tau_{n}<0 \\
t, \tau=0
\end{array}\right.
$$

where

$$
\tau_{n}=\left(\hat{\lambda} r_{n}-q_{n}\right) / p_{n}
$$

and

$$
\omega_{n}=\sqrt{\left|\tau_{n}\right|}
$$

Starting with,

$$
\hat{u}(a) \text { and }\left(\hat{p} \hat{u}^{\prime}\right)(a)
$$

consistent with the boundary condition,

$$
\begin{aligned}
& \hat{u}(a)=a_{2}-a_{2}^{\prime} \hat{\lambda} \\
& \left(\hat{p} \hat{u}^{\prime}\right)(a)=a_{1}-a_{1}^{\prime} \hat{\lambda}
\end{aligned}
$$

an algorithm is to compute for $n=1,2, \ldots, N$,

$$
\begin{aligned}
& \hat{u}\left(x_{n+1}\right)=\hat{u}\left(x_{n}\right) \phi_{n}^{\prime}\left(h_{n}\right)+\left(\hat{p} \hat{u}^{\prime}\right)\left(x_{n}\right) \phi_{n}\left(h_{n}\right) / p_{n} \\
& \left(\hat{p} \hat{u}^{\prime}\right)\left(x_{n+1}\right)=-\tau_{n} p_{n} \hat{u}\left(x_{n}\right) \phi_{n}^{\prime}\left(h_{n}\right)+\left(\hat{p} \hat{u}^{\prime}\right)\left(x_{n}\right) \phi_{n}\left(h_{n}\right)
\end{aligned}
$$

which is a shooting method. For a fixed mesh we can iterate on the approximate eigenvalue until the boundary condition at $b$ is satisfied. This will yield an $O\left(h^{2}\right)$ approximation

$$
\hat{\lambda}_{k}
$$

to some $\lambda_{k}$.
The problem (2) has a step spectral function given by

$$
\hat{\rho}(t)=\sum \frac{1}{\int \hat{r}(x) \hat{u}_{k}^{2}(x) d x+\alpha}
$$

where the sum is taken over $k$ such that

$$
\hat{\lambda}_{k} \leq t
$$

and

$$
\alpha=a_{1}^{\prime} a_{2}-a_{1} a_{2}^{\prime}
$$

## Comments

1. Workspace may be explicitly provided, if desired, by use of S2EIG/DS2EIG. The reference is:

CALL S2EIG (CONS, COEFFN, ENDFIN, NUMEIG, INDEX, TEVLAB, TEVLRL, EVAL, JOB, IPRINT, TOLS, NUMX, XEF, NRHO, T, TYPE, EF, PDEF, RHO, IFLAG, WORK, IWORK)

The additional arguments are as follows:
$J O B$ - Logical array of length five. (Input)
$\operatorname{JOB}(1)=$.true. if a set of eigenvalues are to be computed but not their eigenfunctions.
$\operatorname{JOB}(2)=$.true. if a set of eigenvalue and eigenfunction pairs are to be computed.
$\operatorname{JOB}(3)=$.true. if the spectral function is to be computed over some subinterval of the essential spectrum.
$\operatorname{JOB}(4)=. \operatorname{true}$. if the normal automatic classification is overridden. If $\operatorname{JOB}(4)=$.true. then $\operatorname{TYPE}\left(*,{ }^{*}\right)$ must be entered correctly. Most users will not want to override the classification process, but it might be appropriate for users experimenting with problems for which the coefficient functions do not have power-like behavior near the singular endpoints. The classification is considered
sufficiently important for spectral density function calculations that $\operatorname{JOB(4)~is~}$ ignored with $\operatorname{JOB}(3)=$.true..
$\operatorname{JOB}(5)=$. true. if mesh distribution is chosen by SLEIG. If JOB(5) = .true. and NUMX is zero, the number of mesh points are also chosen by SLEIG. If NUMX $>0$ then NUMX mesh points will be used. If JOB(5) = .false., the number NUMX and distribution $\operatorname{XEF}(*)$ must be input by the user.

IPRINT - Control levels of internal printing. (Input)
No printing is performed if IPRINT $=0$. If either $\operatorname{JOB(1)}$ or $\operatorname{JOB}(2)$ is true:

## IPRINT Printed Output

1 initial mesh (the first 51 or fewer points), eigenvalue estimate at each level
4 the above and at each level matching point for
eigenfunction shooting, $\mathrm{X}\left({ }^{*}\right), \operatorname{EF}\left(^{*}\right)$ and $\operatorname{PDEF}\left({ }^{*}\right)$ values
5 the above and at each level the brackets for the eigenvalue search, intermediate shooting information for the eigenfunction and eigenfunction norm.

If $\operatorname{JOB}(3)=$.true.
IPRINT Printed Output
1 the actual $(\mathrm{a}, \mathrm{b})$ used at each iteration and the total number of eigenvalues computed
2 the above and switchover points to the asymptotic formulas, and some intermediate $\rho(t)$ approximations
4 the above and initial meshes for each iteration, the index of the largest eigenvalue which may be computed, and various eigenvalue and $R_{N}$ values
4 the above and

$$
\hat{\rho}
$$

values at each level
$5 \quad$ the above and $R_{N}$ add eigenvalues below the switchover point
If $\operatorname{JOB}(4)=. f a l$ se.

## IPRINT Printed Output

2 output a description of the spectrum
3 the above and the constants for the Friedrichs' boundary condition(s)
5 the above and intermediate details of the classification calculation

TOLS - Array of length 4 containing tolerances. (Input)
TOLS (1) - absolute error tolerance for eigenfunctions
TOLS (2) - relative error tolerance for eigenfunctions
TOLS (3) - absolute error tolerance for eigenfunction derivatives
TOLS(4) - relative error tolerance for eigenfunction derivatives
The absolute tolerances must be positive.
The relative tolerances must be at least 100 *amach(4)

NUMX - Integer whose value is the number of output points where each eigenfunction is to be evaluated (the number of entries in $\operatorname{XEF}(*)$ ) when $\operatorname{JOB}(2)=$. true.. If $\operatorname{JOB}(5)=$. false. and NUMX is greater than zero, then NUMX is the number of points in the initial mesh used. If $\operatorname{JOB}(5)=$.false., the points in $\operatorname{XEF}\left({ }^{*}\right)$ should be chosen with a reasonable distribution. Since the endpoints $a$ and $b$ must be part of any mesh, NUMX cannot be one in this case. If $\operatorname{JOB}(5)=$.false. and $\operatorname{JOB}(3)=$.true., then NUMX must be positive. On output, NUMX is set to the number of points for eigenfunctions when input NUMX $=0$, and $\operatorname{JOB}(2)$ or $\operatorname{JOB~}(5)=$.true.. (Input/Output)
$\boldsymbol{X E F}$ - Array of points on input where eigenfunction estimates are desired, if $\mathrm{JOB}(2)=$.true.. Otherwise, if $\operatorname{JOB}(5)=$.false. and NUMX is greater than zero, the user's initial mesh is entered. The entries must be ordered so that $a=\operatorname{XEF}(1)<\operatorname{XEF}(2)<\ldots<\operatorname{XEF}(\operatorname{NUMX})=b$. If either endpoint is infinite, the corresponding $\operatorname{XEF}(1)$ or $\operatorname{XEF}$ (NUMX) is ignored. However, it is required that $\operatorname{XEF}(2)$ be negative when $\operatorname{ENDFIN}(1)=$.false., and that $\operatorname{XEF}(\operatorname{NUMX}-1)$ be positive when $\operatorname{ENDFIN}(2)=$.false.. On output, $\operatorname{XEF}(*)$ is changed only if $\operatorname{JOB(2)~and~} \operatorname{JOB(5)}$ are true. If $\operatorname{JOB}(2)=$.false., this vector is not referenced. If $\operatorname{JOB}(2)=$.true. and NUMX is greater than zero on input, $\operatorname{XEF}(*)$ should be dimensioned at least NUMX +16 . If JOB( 2 ) is true and NUMX is zero on input, $\operatorname{XEF}(*)$ should be dimensioned at least 31 .

NRHO - The number of output values desired for the array RHO $\left(^{*}\right)$. NRHO is not used if $\operatorname{JOB}(3)=$.false.. (Input)

T-Real vector of size NRHO containing values where the spectral function RHO(*) is desired. The entries must be sorted in increasing order. The existence and location of a continuous spectrum can be determined by calling SLEIG with the first four entries of JOB set to false and IPRINT set to $1 . T\left({ }^{*}\right)$ is not used if $\mathrm{JOB}(3)=$.false.. (Input)

TYPE - 4 by 2 logical matrix. Column 1 contains information about endpoint $a$ and column 2 refers to endpoint $b$.
$\operatorname{TYPE}(1, *)=$. true. if and only if the endpoint is regular
$\operatorname{TYPE}\left(2,{ }^{*}\right)=$.true. if and only if the endpoint is limit circle
$\operatorname{TYPE}\left(3,{ }^{*}\right)=$.true. if and only if the endpoint is nonoscillatory for all eigenvalues
$\operatorname{TYPE}(4, *)=$.true. if and only if the endpoint is oscillatory for all eigenvalues
Note: all of these values must be correctly input if $\mathrm{JOB}(4)=$.true.
Otherwise, $\operatorname{TYPE}(*, *)$ is output. (Input/Output)
$\boldsymbol{E F}$ - Array of eigenfunction values. $\operatorname{EF}\left((\mathrm{k}-1)^{*} \mathrm{NUMX}+\mathrm{i}\right)$ is the estimate of $\mathrm{u}(\mathrm{XEF}(\mathrm{i}))$ corresponding to the eigenvalue in $\operatorname{EV}(\mathrm{k})$. If $\mathrm{JOB}(2)=$.false. then this vector is not referenced. If $\operatorname{JOB}(2)=$.true. and NUMX is greater than zero on entry, then $\mathrm{EF}\left(^{*}\right)$ should be dimensioned at least NUMX * NUMEIG. If $\operatorname{JOB}(2)=$.true. and NUMX is zero on input, then $\operatorname{EF}\left({ }^{*}\right)$ should be dimensioned 31 * nUMEIG. (Output)

PDEF - Array of eigenfunction derivative values. $\operatorname{PDEF}((\mathrm{k}-1) * \operatorname{NUMX}+\mathrm{i})$ is the estimate of $\left(p u^{\prime}\right)(\operatorname{XEF}(\mathrm{i}))$ corresponding to the eigenvalue in $\operatorname{EV}(\mathrm{k})$. If $\mathrm{JOB}(2)=$.false. this vector is not referenced. If $\operatorname{JOB}(2)=$.true., it must be dimensioned the same as $E F(*)$. (Output)
$\boldsymbol{R H O}$ - Array of size NRHO containing values for the spectral density function $\rho(t)$, $R H O(I)=\rho(T(I))$. This vector is not referenced if $\operatorname{JOB}(3)$ is false. (Output)

IFLAG - Array of size max(1, numeig) containing information about the output. IFLAG(K) refers to the K -th eigenvalue, when $\mathrm{JOB}(1)$ or $\operatorname{JOB}(2)=$.true.. Otherwise, only IFLAG(1) is used. Negative values are associated with fatal errors, and the calculations are ceased. Positive values indicate a warning. (Output) IFLAG(K)

## IFLAG(K) Description

-1 too many levels needed for the eigenvalue calculation; problem seems too difficult at this tolerance. Are the coefficient functions nonsmooth?
-2 too many levels needed for the eigenfunction calculation; problem seems too difficult at this tolerance. Are the eigenfunctions ill-conditioned?
-3 too many levels needed for the spectral density calculation; problem seems too difficult at this tolerance.
-4 the user has requested the spectral density function for a problem which has no continuous spectrum.
-5 the user has requested the spectral density function for a problem with both endpoints generating essential spectrum, i.e. both endpoints either OSC or O-NO.
-6 the user has requested the spectral density function for a problem in spectral category 2 for which a proper normalization of the solution at the NONOSC endpoint is not known; for example, problems with an irregular singular point or infinite endpoint at one end and continuous spectrum generated at the other.
-7 problems were encountered in obtaining a bracket.
-8 too small a step was used in the integration. The TOLS (*) values may be too small for this problem.
-9 too small a step was used in the spectral density function calculation for which the continuous spectrum is generated by a finite endpoint.
-10 an argument to the circular trig functions is too large. Try running the problem again with a finer initial mesh or, for singular problems, use interval truncation.
-15 $\quad p(x)$ and $r(x)$ are not positive in the interval $(a, b)$.
-20 eigenvalues and/or eigenfunctions were requested for a problem with an OSC singular endpoint. Interval truncation must be used on such problems.

1 Failure in the bracketing procedure probably due to a cluster of eigenvalues which the code cannot separate. Calculations have continued but any eigenfunction results are suspect. Try running the problem again with tighter input tolerances to separate the cluster.
2 there is uncertainty in the classification for this problem. Because of the limitations of floating point arithmetic, and the nature of the finite sampling, the routine cannot be certain about the classification information at the requested tolerance.
3 there may be some eigenvalues embedded in the essential spectrum. Use of IPRINT greater than zero will provide additional output giving the location of the approximating eigenvalues for the step function problem. These could be extrapolated to estimate the actual eigenvalue embedded in the essential spectrum. a change of variables was made to avoid potentially slow convergence. However, the global error estimates may not be as reliable. Some experimentation using different tolerances is recommended.
there were problems with eigenfunction convergence in a spectral density calculation. The output $\rho(t)$ may not be accurate.

WORK - Array of size MAX (1000, NUMEIG + 22) used for workspace.
IWORK - Integer array of size NUMEIG +3 used for workspace.

## Example 1

This example computes the first ten eigenvalues of the problem from Titchmarsh (1962) given by

```
p(x)=r(x)=1
q(x) = x
[a, b] = [0, \infty]
u(a) =u(b)=0
```

The eigenvalues are known to be the zeros of

$$
f(\lambda)=J_{1 / 3}\left(\frac{2}{3} \lambda^{3 / 2}\right)+J_{-1 / 3}\left(\frac{2}{3} \lambda^{3 / 2}\right)
$$

For each eigenvalue $\lambda_{k}$, the program prints $k, \lambda_{k}$ and $f\left(\lambda_{k}\right)$.

```
USE SLEIG INT
USE CBJS_INT
IMPLICIT NONE
```

```
! SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER I, INDEX(10), NUMEIG, NOUT
    REAL CONS(8), EVAL(10), LAMBDA, TEVLAB,&
        TEVLRL, XNU
    COMPLEX CBS1(1), CBS2(1), Z
    LOGICAL ENDFIN(2)
    INTRINSIC CMPLX, SQRT
    REAL SQRT
    COMPLEX CMPLX
        SPECIFICATIONS FOR SUBROUTINES
        SPECIFICATIONS FOR FUNCTIONS
    EXTERNAL COEFF
    CALL UMACH (2, NOUT)
        Define boundary conditions
    CONS(1) = 1.0
    CONS (2) = 0.0
    CONS (3) = 0.0
    CONS (4) = 0.0
    CONS(5) = 1.0
    CONS (6) = 0.0
    CONS (7) = 0.0
    CONS (8) = 0.0
!
    ENDFIN(1) = .TRUE.
    ENDFIN(2) = .FALSE.
!
    NUMEIG = 10
    DO 10 I=1, NUMEIG
        INDEX(I) = I - 1
    1 0 ~ C O N T I N U E
!
    CALL SLEIG (CONS, COEFF, ENDFIN, INDEX, EVAL)
    XNU = -1.0/3.0
    WRITE (NOUT, 99998)
    DO 20 I=1, NUMEIG
        LAMBDA = EVAL(I)
        Z = CMPLX(2.0/3.0*LAMBDA*SQRT (LAMBDA),0.0)
        CALL CBJS (XNU, Z, 1, CBS1)
        CALL CBJS (-XNU, Z, 1, CBS2)
        WRITE (NOUT,99999) I-1, LAMBDA, REAL(CBS1(1) + CBS2(1))
    2 0 ~ C O N T I N U E
!
99998 FORMAT(/, 2X, 'index', 5X, 'lambda', 5X, 'f(lambda)',/)
99999 FORMAT(I5, F13.4, E15.4)
    END
    SUBROUTINE COEFF (X, PX, QX, RX)
        SPECIFICATIONS FOR ARGUMENTS
    REAL X, PX, QX, RX
    PX = 1.0
```

```
QX = X
RX = 1.0
RETURN
END
```


## Output

| index | lambda | $f($ lambda $)$ |
| :---: | ---: | ---: |
|  |  |  |
| 0 | 2.3381 | $-0.8285 \mathrm{E}-05$ |
| 1 | 4.0879 | $-0.1651 \mathrm{E}-04$ |
| 2 | 5.5205 | $0.6843 \mathrm{E}-04$ |
| 3 | 6.7867 | $-0.4523 \mathrm{E}-05$ |
| 4 | 7.9440 | $0.8952 \mathrm{E}-04$ |
| 5 | 9.0227 | $0.1123 \mathrm{E}-04$ |
| 6 | 10.0401 | $0.1031 \mathrm{E}-03$ |
| 7 | 11.0084 | $-0.7913 \mathrm{E}-04$ |
| 8 | 11.9361 | $-0.5095 \mathrm{E}-04$ |
| 9 | 12.8293 | $0.4645 \mathrm{E}-03$ |

## Additional Examples

## Example 2

In this problem from Scott, Shampine and Wing (1969),

$$
\begin{aligned}
& p(x)=r(x)=1 \\
& q(x)=x^{2}+x^{4} \\
& {[a, b]=[-\infty, \infty]} \\
& u(a)=u(b)=0
\end{aligned}
$$

the first eigenvalue and associated eigenfunction, evaluated at selected points, are computed. As a rough check of the correctness of the results, the magnitude of the residual

$$
-\frac{d}{d x}\left(p(x) \frac{d u}{d x}\right)+q(x) u-\lambda r(x) u
$$

is printed. We compute a spline interpolant to $u^{\prime}$ and use the function CSDER to estimate the quantity $-\left(p(x) u^{\prime}\right)^{\prime}$.

```
USE S2EIG_INT
USE CSDER ' INT
USE UMACH_INT
USE CSAKM_INT
IMPLICIT NONE
INTEGER I, IFLAG(1), INDEX(1), IWORK(100), NINTV, NOUT, NRHO, &
    NUMEIG, NUMX
REAL BRKUP(61), CONS(8), CSCFUP(4,61), EF(61), EVAL(1), &
    LAMBDA, PDEF(61), PX, QX, RESIDUAL, RHO(1), RX, T(1), &
```

```
            TEVLAB, TEVLRL, TOLS(4), WORK(3000), X, XEF(61)
    LOGICAL ENDFIN(2), JOB(5), TYPE (4,2)
                                    SPECIFICATIONS FOR INTRINSICS
    INTRINSIC ABS, REAL
    REAL ABS, REAL
    SPECIFICATIONS FOR SUBROUTINES
    EXTERNAL COEFF
    CONS (1) = 1.0
    CONS (2) = 0.0
    CONS (3) = 0.0
    CONS(4) = 0.0
    CONS (5) = 1.0
    CONS (6) = 0.0
    CONS (7) = 0.0
    CONS(8) = 0.0
    Compute eigenvalue and eigenfunctions
    JOB(1) = .FALSE.
    JOB(2) = .TRUE.
    JOB(3) = .FALSE.
    JOB(4) = .FALSE.
    JOB(5) = .FALSE.
!
    ENDFIN(1) = .FALSE.
    ENDFIN(2) = .FALSE.
    NUMEIG = 1
    INDEX(1) = 0
    TEVLAB = 1.0E-3
    TEVLRL = 1.0E-3
    TOLS(1) = TEVLAB
    TOLS(2) = TEVLRL
    TOLS(3) = TEVLAB
    TOLS(4) = TEVLRL
    NRHO = 0
! Set up mesh, points at which u and
! u' will be computed
        NUMX = 61
        DO 10 I=1, NUMX
            XEF(I) = 0.05*REAL (I-31)
    10 CONTINUE
!
        CALL S2EIG (CONS, COEFF, ENDFIN, NUMEIG, INDEX, TEVLAB, TEVLRL, &
                EVAL, JOB, 0, TOLS, NUMX, XEF, NRHO, T, TYPE, EF, &
                PDEF, RHO, IFLAG, WORK, IWORK)
        LAMBDA = EVAL(1)
    20 CONTINUE
Compute spline interpolant to u'
    CALL CSAKM (XEF, PDEF, BRKUP, CSCFUP)
    NINTV = NUMX - 1
    CALL UMACH (2, NOUT)
```

```
    WRITE (NOUT,99997) ' lambda = ', LAMBDA
    WRITE (NOUT,99999)
!
    At a subset of points from the
    input mesh, compute residual =
    abs( -(u')' + q(x)u - lambda*u ).
    We know }\textrm{p}(\textrm{x})=1\mathrm{ and }r(x)=1
!
                                    Use the spline fit to u' to
                                    estimate u'' with CSDER
    RESIDUAL = ABS (-CSDER (1,X,BRKUP,CSCFUP)+QX*EF}(I+10)-
                                    LAMBDA*EF(I+10))
        WRITE (NOUT,99998) X, EF(I+10), PDEF(I+10), RESIDUAL
    3 0 ~ C O N T I N U E ~
!
9 9 9 9 7 ~ F O R M A T ~ ( / , ~ A 1 4 , ~ F 1 0 . 5 , ~ / ) ~
99998 FORMAT (5X, F4.1, 3F15.5)
99999 FORMAT (7X, 'x', 11X, 'u(x)', 10X, 'u''(x)', 9X, 'residual', /)
        END
!
    SUBROUTINE COEFF (X, PX, QX, RX)
                                    SPECIFICATIONS FOR ARGUMENTS
    REAL X, PX, QX, RX
PX = 1.0
QX = X*X + X* X*X*X
RX = 1.0
RETURN
END
```


## Output

| lambda $=$ | 1.39247 |  |  |
| :---: | :---: | ---: | :---: |
| $x$ | $u(x)$ | $u^{\prime}(x)$ | residual |
| -1.0 | 0.38632 | 0.65019 | 0.00189 |
| -0.9 | 0.45218 | 0.66372 | 0.00081 |
| -0.8 | 0.51837 | 0.65653 | 0.00023 |
| -0.7 | 0.58278 | 0.62827 | 0.00113 |
| -0.6 | 0.64334 | 0.57977 | 0.00183 |
| -0.5 | 0.69812 | 0.51283 | 0.00230 |
| -0.4 | 0.74537 | 0.42990 | 0.00273 |
| -0.3 | 0.78366 | 0.33393 | 0.00265 |
| -0.2 | 0.81183 | 0.22811 | 0.00273 |
| -0.1 | 0.82906 | 0.11570 | 0.00278 |
| 0.0 | 0.83473 | 0.00000 | 0.00136 |
| 0.1 | 0.82893 | -0.11568 | 0.00273 |
| 0.2 | 0.81170 | -0.22807 | 0.00273 |
| 0.3 | 0.78353 | -0.33388 | 0.00267 |
| 0.4 | 0.74525 | -0.42983 | 0.00265 |
| 0.5 | 0.69800 | -0.51274 | 0.00230 |
| 0.6 | 0.64324 | -0.57967 | 0.00182 |
| 0.7 | 0.58269 | -0.62816 | 0.00113 |


| 0.8 | 0.51828 | -0.65641 | 0.00023 |
| :--- | :--- | :--- | :--- |
| 0.9 | 0.45211 | -0.66361 | 0.00081 |
| 1.0 | 0.38626 | -0.65008 | 0.00189 |

## SLCNT

Calculates the indices of eigenvalues of a Sturm-Liouville problem of the form for

$$
-\frac{d}{d x}\left(p(x) \frac{d u}{d x}\right)+q(x) u=\lambda r(x) u \text { for } x \text { in }[a, b]
$$

with boundary conditions (at regular points)

$$
\begin{gathered}
a_{1} u-a_{2}\left(p u^{\prime}\right)=\lambda\left(a_{1}^{\prime} u-a_{2}^{\prime}\left(p u^{\prime}\right)\right) \text { at } a \\
b_{1} u+b_{2}\left(p u^{\prime}\right)=0 \text { at } b
\end{gathered}
$$

in a specified subinterval of the real line, $[\alpha, \beta]$.

## Required Arguments

ALPHA - Value of the left end point of the search interval. (Input)
BETAR - Value of the right end point of the search interval. (Input)
CONS - Array of size eight containing
$a_{1}, a_{1}^{\prime}, a_{2}, a_{2}^{\prime}, b_{1}, b_{2}, a$ and $b$
in locations CONS (1) $\ldots$ CONS (8), respectively. (Input)
COEFFN - User-supplied SUBROUTINE to evaluate the coefficient functions. The usage is CALL COEFFN ( $\mathrm{X}, \mathrm{PX}, \mathrm{QX}, \mathrm{RX}$ ) X - Independent variable. (Input)
PX - The value of $p(x)$ at X . (Output)
QX - The value of $q(x)$ at X . (Output)
RX - The value of $r(x)$ at X . (Output)
COEFFN must be declared EXTERNAL in the calling program.
ENDFIN - Logical array of size two. ENDFIN = . true. if and only if the endpoint $a$ is finite. ENDFIN $(2)=$.true. if and only if endpoint $b$ is finite. (Input)

IFIRST - The index of the first eigenvalue greater than $\alpha$. (Output)
NTOTAL - Total number of eigenvalues in the interval $[\alpha, \beta]$. (Output)

## FORTRAN 90 Interface

Generic: CALL SLCNT (ALPHA, BETAR, CONS, COEFFN, ENDFIN, IFIRST,
Specific: The specific interface names are S_SLCNT and D_SLCNT.

## FORTRAN 77 Interface

Single: CALL SLCNT (ALPHA, BETAR, CONS, COEFFN, ENDFIN, IFIRST, NTOTAL)

Double: The double precision name is DSLCNT.

## Description

This subroutine computes the indices of eigenvalues, if any, in a subinterval of the real line for Sturm-Liouville problems in the form

$$
-\frac{d}{d x}\left(p(x) \frac{d u}{d x}\right)+q(x) u=\lambda r(x) u \text { for } x \text { in }[a, b]
$$

with boundary conditions (at regular points)

$$
\begin{gathered}
a_{1} u-a_{2}\left(p u^{\prime}\right)=\lambda\left(a_{1}^{\prime} u-a_{2}^{\prime}\left(p u^{\prime}\right)\right) \text { at } a \\
b_{1} u+b_{2}\left(p u^{\prime}\right)=0 \text { at } b
\end{gathered}
$$

It is intended to be used in conjunction with SLEIG. SLCNT is based on the routine INTERV from the package SLEDGE.

## Example

Consider the harmonic oscillator (Titchmarsh) defined by

$$
\begin{aligned}
& p(x)=1 \\
& q(x)=x^{2} \\
& r(x)=1 \\
& {[a, b]=[-\infty, \infty]} \\
& u(a)=0 \\
& u(b)=0
\end{aligned}
$$

The eigenvalues of this problem are known to be

$$
\lambda_{k}=2 k+1, k=0,1, \ldots
$$

Therefore in the interval $[10,16]$ we expect SLCNT to note three eigenvalues, with the first of these having index five.

```
USE SLCNT_INT
```

```
USE UMACH_INT
IMPLICIT NONE
! SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER IFIRST, NOUT, NTOTAL
REAL ALPHA, BETAR, CONS (8)
LOGICAL ENDFIN(2)
SPECIFICATIONS FOR SUBROUTINES
SPECIFICATIONS FOR FUNCTIONS
EXTERNAL COEFFN
CALL UMACH (2, NOUT)
CONS (1) = 1.0E0
CONS (2) = 0.0EO
CONS (3) = 0.0E0
CONS (4) = 0.0EO
CONS(5) = 1.0E0
CONS(6) = 0.0EO
CONS (7) = 0.0EO
CONS (8) = 0.0EO
!
ENDFIN(1) = .FALSE.
ENDFIN(2) = .FALSE.
!
ALPHA = 10.0
BETAR = 16.0
CALL SLCNT (ALPHA, BETAR, CONS, COEFFN, ENDFIN, IFIRST, NTOTAL)
WRITE (NOUT,99998) ALPHA, BETAR, IFIRST
WRITE (NOUT,99999) NTOTAL
!
99998 FORMAT (/, 'Index of first eigenvalue in [', F5.2, ',', F5.2, &
    '] IS ', I2)
99999 FORMAT ('Total number of eigenvalues in this interval: ', I2)
!
    END
SUBROUTINE COEFFN (X, PX, QX, RX)
REAL X, PX, QX, RX
PX = 1.0E0
QX = X*X
RX = 1.0E0
RETURN
END
```


## Output

```
Index of first eigenvalue in [10.00,16.00] is 5
```

Total number of eigenvalues in this interval: 3

## Chapter 6: Transforms

## Routines

6.1. Real Trigonometric FFT
Computes the Discrete Fourier Transform
of a rank-1 complex array, $x$. ........................................FAST DFT ..... 1086
Computes the Discrete Fourier Transform (2DFT) of a rank-2 complex array, $x$.........................................FAST_2DFT ..... 1093
Computes the Discrete Fourier Transform 2DFT)
of a rank-3 complex array, $x$......................................FAST_3DFT ..... 1099
Forward transform FFTRF ..... 1103
Backward or inverse transform FFTRB ..... 1106
Initialization routine for FFTR* ..... 1109
6.2. Complex Exponential FFT
Forward transform FFTCF ..... 1111
Backward or inverse transform FFTCB ..... 1113
Initialization routine for FFTC* FFTCI ..... 1116
6.3. Real Sine and Cosine FFTs
Forward and inverse sine transform ..... FSINT ..... 1118
Initialization routine for FSINT ..... FSINI ..... 1120
Forward and inverse cosine transform FCOST ..... 1122
Initialization routine for FCOST FCOSI ..... 1124
6.4. Real Quarter Sine and Quarter Cosine FFTs
Forward quarter sine transform QSINF ..... 1126
Backward or inverse transform ..... QSINB ..... 1129
Initialization routine for QSIN*. ..... QSINI ..... 1131
Forward quarter cosine transform QCOSF ..... 1133
Backward or inverse transform QCOSB ..... 1135
Initialization routine for QCOS* QCOSI ..... 1137
6.5. Two- and Three-Dimensional Complex FFTs
Forward transform FFT2D ..... 1139
Backward or inverse transform FFT2B ..... 1142
Forward transform FFT3F ..... 1145
Backward or inverse transform FFT3B ..... 1149
6.6. Convolutions and Correlations
Real convolution................................................................. RCONV
1153
Complex convolution .......................................................... CCONV 1158
Real correlation ...................................................................RCORL 1163
Complex correlation ...........................................................CCORL 1168
6.7. Laplace Transform
Inverse Laplace transform...................................................INLAP 1172
Inverse Laplace transform for smooth functions ................... SINLP 1175

## Usage Notes

## Fast Fourier Transforms

A Fast Fourier Transform (FFT) is simply a discrete Fourier transform that can be computed efficiently. Basically, the straightforward method for computing the Fourier transform takes approximately $N^{2}$ operations where $N$ is the number of points in the transform, while the FFT (which computes the same values) takes approximately $N \log N$ operations. The algorithms in this chapter are modeled on the Cooley-Tukey (1965) algorithm; hence, the computational savings occur, not for all integers $N$, but for $N$ which are highly composite. That is, $N$ (or in certain cases $N+1$ or $N-1$ ) should be a product of small primes.

All of the FFT routines compute a discrete Fourier transform. The routines accept a vector $x$ of length $N$ and return a vector

$$
\hat{x}
$$

defined by

$$
\hat{x}_{m}:=\sum_{n=1}^{N} x_{n} \omega_{n m}
$$

The various transforms are determined by the selection of $\omega$. In the following table, we indicate the selection of $\omega$ for the various transforms. This table should not be mistaken for a definition since the precise transform definitions (at times) depend on whether $N$ or $m$ is even or odd.

| Routine | $\omega_{n m}$ |
| :---: | :---: |
| FFTRF | $\cos$ or $\sin \frac{(m-1)(n-1) 2 \pi}{N}$ |
| FFTRB | $\cos$ or $\sin \frac{(m-1)(n-1) 2 \pi}{N}$ |
| FFTCF | $\exp ^{-2 \pi i(n-1)(m-1) / N}$ |
| FFTCB | $\exp ^{2 \pi i(n-1)(m-1) / N}$ |
| FSINT | $\sin \frac{n m \pi}{N+1}$ |
| FCOST | $\cos \frac{(n-1)(m-1) \pi}{N-1}$ |
| QSINF | $2 \sin \frac{(2 m-1) n \pi}{2 N}$ |
| QSINB | $4 \sin \frac{(2 n-1) m \pi}{2 N}$ |
| QCOSF | $2 \cos \frac{(2 m-1)(n-1) \pi}{2 N}$ |
| QCOSB | $4 \cos \frac{(2 n-1)(m-1) \pi}{2 N}$ |

For many of the routines listed above, there is a corresponding " $I$ " (for initialization) routine. Use these routines only when repeatedly transforming sequences of the same length. In this situation, the "I" routine will compute the initial setup once, and then the user will call the corresponding " 2 " routine. This can result in substantial computational savings. For more information on the usage of these routines, the user should consult the documentation under the appropriate routine name.

In addition to the one-dimensional transformations described above, we also provide complex two and three-dimensional FFTs and their inverses based on calls to either FFTCF or FFTCB. If you need a higher dimensional transform, then you should consult the example program for FFTCI, which suggests a basic strategy one could employ.

## Continuous versus Discrete Fourier Transform

There is, of course, a close connection between the discrete Fourier transform and the continuous Fourier transform. Recall that the continuous Fourier transform is defined (Brigham, 1974) as

$$
\hat{f}(\omega)=(F f)(\omega)=\int_{-\infty}^{\infty} f(t) e^{-2 \pi i \omega t} d t
$$

We begin by making the following approximation:

$$
\begin{aligned}
\hat{f}(\omega) & \approx \int_{-T / 2}^{T / 2} f(t) e^{-2 \pi i \omega t} d t \\
& =\int_{0}^{T} f(t-T / 2) e^{-2 \pi i \omega(t-T / 2)} d t \\
& =e^{\pi i \omega T} \int_{0}^{T} f(t-T / 2) e^{-2 \pi i \omega t} d t
\end{aligned}
$$

If we approximate the last integral using the rectangle rule with spacing $h=T / N$, we have

$$
\hat{f}(\omega) \approx e^{\pi i \omega T} h \sum_{k=0}^{N-1} e^{-2 \pi i \omega k h} f(k h-T / 2)
$$

Finally, setting $\omega=j / T$ for $j=0, \ldots, N-1$ yields

$$
\hat{f}(j / T) \approx e^{\pi i j} h \sum_{k=0}^{N-1} e^{-2 \pi i j k / N} f(k h-T / 2)=(-1)^{j} h \sum_{k=0}^{N-1} e^{-2 \pi i j k / N} f_{k}^{h}
$$

where the vector $f^{h}=(f(-T / 2), \ldots, f((N-1) h-T / 2))$. Thus, after scaling the components by $(-1)^{j} h$, the discrete Fourier transform as computed in FFTCF (with input $f^{h}$ ) is related to an approximation of the continuous Fourier transform by the above formula. This is seen more clearly by making a change of variables in the last sum. Set

$$
n=k+1, m=j+1, \text { and } f_{k}^{h}=x_{n}
$$

then, for $m=1, \ldots, N$ we have

$$
\hat{f}((m-1) / T) \approx-(-1)^{m} h \hat{x}_{m}=-(-1)^{m} h \sum_{n=1}^{N} e^{-2 \pi i(m-1)(n-1) / N} x_{n}
$$

If the function $f$ is expressed as a FORTRAN function routine, then the continuous Fourier transform

$$
\hat{f}
$$

can be approximated using the IMSL routine QDAWF (see Chapter 4, Integration and Differentiation).

## Inverse Laplace Transform

The last two routines described in this chapter, INLAP and SINLP, compute the inverse Laplace transforms.

## FAST_DFT

Computes the Discrete Fourier Transform (DFT) of a rank-1 complex array, $x$.

## Required Arguments

No required arguments; pairs of optional arguments are required. These pairs are forward_in and forward_out or inverse_in and inverse_out.

## Optional Arguments

```
forward in = x (Input)
```

Stores the input complex array of rank-1 to be transformed.
forward_out = y (Output)
Stores the output complex array of rank-1 resulting from the transform.
inverse_in = y (Input)
Stores the input complex array of rank-1 to be inverted.
inverse_out = x (Output)
Stores the output complex array of rank-1 resulting from the inverse transform.
ndata $=\mathrm{n}$ (Input)
Uses the sub-array of size $n$ for the numbers.
Default value: $n=\operatorname{size}(x)$.
ido = ido (Input/Output)
Integer flag that directs user action. Normally, this argument is used only when the working variables required for the transform and its inverse are saved in the calling program unit. Computing the working variables and saving them in internal arrays within fast_dft is the default. This initialization step is expensive.

There is a two-step process to compute the working variables just once. Example 3 illustrates this usage. The general algorithm for this usage is to enter fast_dft with ido $=0$. A return occurs thereafter with ido $<0$. The optional rank-1 complex array $\mathrm{w}(:)$ with size $(\mathrm{w})>=-$ ido must be re-allocated. Then, re-enter fast dft. The next return from fast dft has the output value ido $=1$. The variables required for the transform and its inverse are saved in $w(:)$. Thereafter, when the routine is entered with ido $=1$ and for the same value of $n$, the contents of $w(:)$ will be used for the working variables. The expensive initialization step is avoided. The optional arguments "ido=" and "work_array=" must be used together.
work_array = w(:) (Output/Input)
Complex array of rank-1 used to store working variables and values between calls to fast_dft. The value for size(w) must be at least as large as the value - ido for the value of ido $<0$.
iopt = iopt(:) (Input/Output)
Derived type array with the same precision as the input array; used for passing optional data to fast_dft. The options are as follows:

| Packaged Options for FAST_DFT |  |  |
| :---: | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| $c_{-}, z_{-}$ | fast_dft_scan_for_NaN | 1 |
| $c_{-}, z_{-}$ | fast_dft_near_power_of_2 | 2 |
| $c_{-}, z_{-}$ | fast_dft_scale_forward | 3 |
| $c_{-}, z_{-}$ | Fast_dft_scale_inverse | 4 |

iopt(IO) = ?_options(?_fast_dft_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(x(i)) ==.true.
See the isnan () function, Chapter 10.
Default: Does not scan for NaNs.

```
iopt(IO) = ?_options(?_fast_dft_near_power_of_2, ?_dummy)
```

Nearest power of $2 \geq n$ is returned as an output in iopt (IO +1 ) \%idummy.

```
iopt(IO) = ?_options(?_fast_dft_scale_forward, real_part_of_scale)
```

iopt(IO+1) = ?_options(?_dummy, imaginary_part_of_scale)
Complex number defined by the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is multiplied by the forward transformed array.
Default value is 1 .
iopt(IO) = ?_options(?_fast_dft_scale_inverse, real_part_of_scale)
iopt(IO+1) = ?_options(?_dummy, imaginary_part_of_scale)
Complex number defined by the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is
multiplied by the inverse transformed array.
Default value is 1 .

## FORTRAN 90 Interface

Generic: None

Specific: The specific interface names are S_FAST_DFT, D_FAST_DFT, C_FAST_DFT, and Z_FAST_DFT.

## Description

The fast_dft routine is a Fortran 90 version of the FFT suite of IMSL (1994, pp. 772-776). The maximum computing efficiency occurs when the size of the array can be factored in the form

$$
n=2^{i_{i}} 3^{i_{2}} 4^{i_{5}} 5^{i_{4}}
$$

using non-negative integer values $\left\{i_{1}, i_{2}, i_{3}, i_{4}\right\}$. There is no further restriction on $n \geq 1$.

## Fatal and Terminal Messages

See the messages.gls file for error messages for FAST_DFT. These error messages are numbered 651-661; 701-711.

## Example 1: Transforming an Array of Random Complex Numbers

An array of random complex numbers is obtained. The transform of the numbers is inverted and the final results are compared with the input array.

```
    use fast_dft_int
    use rand_gen_int
    implicit none
! This is Example 1 for FAST_DFT.
integer, parameter :: n=1024
real(kind(1e0)), parameter : : one=1e0
real(kind(le0)) err, y(2*n)
complex(kind(le0)), dimension(n) :: a, b, c
! Generate a random complex sequence.
    call rand_gen(y)
    a = cmplx(y(1:n),y(n+1:2*n),kind(one))
    c = a
! Transform and then invert the sequence back.
    call c_fast_dft(forward_in=a, &
        forward_out=b)
    call c_fast_dft(inverse_in=b, &
        inverse_out=a)
! Check that inverse(transform(sequence)) = sequence.
    err = maxval(abs(c-a))/maxval(abs(c))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for FAST_DFT is correct.'
    end if
    end
```


## Output

Example 1 for FAST_DFT is correct.

## Additional Examples

## Example 2: Cyclical Data with a Linear Trend

This set of data is sampled from a function $x(t)=a t+b+y(t)$, where $y(t)$ is a harmonic series. The independent variable is normalized as $-1 \leq t \leq 1$. Thus, the data is said to have cyclical components plus a linear trend. As a first step, the linear terms are effectively removed from the
data using the least-squares system solver LIN_SOL_LSQ, Chapter 1. Then, the residuals are transformed and the resulting frequencies are analyzed.

```
    use fast_dft_int
    use lin sol l
    use rand_gen_int
    use sort_real_int
    implicit none
! This is Example 2 for FAST_DFT.
    integer i
    integer, parameter :: n=64, k=4
    integer ip(n)
    real(kind(1e0)), parameter :: one=1e0, two=2e0, zero=0e0
    real(kind(1e0)) delta_t, pi
    real(kind(le0)) y(k), z(2), indx(k), t(n), temp(n)
    complex(kind(le0)) a_trend(n,2), a, b_trend(n,1), b, c(k), f(n),&
                r(n), x(n), x_trend (2,1)
! Generate random data for linear trend and harmonic series.
    call rand_gen(z)
    a = z(1); b = z(2)
    call rand_gen(y)
! This emphasizes harmonics 2 through k+1.
    c = y + one
! Determine sampling interval.
    delta_t = two/n
    t=(/(-one+i*delta_t, i=0,n-1)/)
! Compute pi.
    pi = atan(one)*4E0
    indx=(/(i*pi,i=1,k)/)
! Make up data set as a linear trend plus harmonics.
    x = a + b*t + &
        matmul(exp(cmplx(zero, spread(t, 2,k)*spread(indx,1,n),kind(one))), c)
! Define least-squares matrix data for a linear trend.
    a_trend(1:,1) = one
    a_trend (1:,2) = t
    b_trend (1:,1) = x
! Solve for a linear trend.
    call lin_sol_lsq(a_trend, b_trend, x_trend)
! Compute harmonic residuals.
    r = x - reshape(matmul(a_trend,x_trend),(/n/))
! Transform harmonic residuals.
    call c_fast_dft(forward_in=r, forward_out=f)
    ip=(/(\overline{i},i=1,
```

```
! The dominant frequencies should be 2 through k+1.
! Sort the magnitude of the transform first.
    call s_sort_real(-(abs(f)), temp, iperm=ip)
! The dominant frequencies are output in ip(1:k).
! Sort these values to compare with 2 through k+1.
    call s_sort_real(real(ip(1:k)), temp)
    ip(1:k)=(/(i,i=2,k+1)/)
! Check the results.
    if (count(int(temp(1:k)) /= ip(1:k)) == 0) then
        write (*,*) 'Example 2 for FAST_DFT is correct.'
    end if
    end
```


## Output

Example 2 for FAST_DFT is correct.

## Example 3: Several Transforms with Initialization

In this example, the optional arguments ido and work_array are used to save working variables in the calling program unit. This results in maximum efficiency of the transform and its inverse since the working variables do not have to be precomputed following each entry to routine fast_dft.
use fast_dft_int
use rand_gen_int
implicit none
! This is Example 3 for FAST_DFT.
! The value of the array size for work(:) is computed in the
! routine fast_dft as a first step.
integer, parameter :: n=64
integer ido_value
real(kind(1-0)) :: one=1e0
real(kind(1e0)) err, y(2*n)
complex(kind(le0)), dimension(n) :: a, b, save_a
complex(kind(1e0)), allocatable :: work(:)
! Generate a random complex array.
call rand_gen(y)
$a=\operatorname{cmplx} \overline{(y}(1: n), y(n+1: 2 * n)$, kind(one))
save_a = a
! Transform and then invert the sequence using the pre-computed
! working values.
ido_value = 0
do
if(allocated(work)) deallocate(work)

```
! Allocate the space required for work(:).
    if (ido value <= 0) allocate(work(-ido value))
    call c_fast_dft(forward_in=a, forward_out=b, &
```



```
    if (ido_value == 1) exit
    end do
! Re-enter routine with working values available in work(:).
    call c_fast_dft(inverse_in=b, inverse_out=a, &
            ido=ido_value, work_array=work)
! Deallocate the space used for work(:).
    if (allocated(work)) deallocate(work)
! Check the results.
    err = maxval(abs(save a-a))/maxval(abs(save a))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 3 for FAST_DFT is correct.'
    end if
    end
```


## Output

Example 3 for FAST_DFT is correct.

## Example 4: Convolutions using Fourier Transforms

In this example we compute sums

$$
c_{k}=\sum_{j=0}^{n-1} a_{j} b_{k-j}, k=0, \ldots, n-1
$$

The definition implies a matrix-vector product. A direct approach requires about $n^{2}$ operations consisisting of an add and multiply. An efficient method consisting of computing the products of the transforms of the

$$
\left\{a_{j}\right\} \text { and }\left\{b_{j}\right\}
$$

then inverting this product, is preferable to the matrix-vector approach for large problems. The example is also illustrated in operator_ex37, Chapter 10 using the generic function interface FFT and IFFT.
use fast_dft_int
use rand_gen_int
implicit none
! This is Example 4 for FAST_DFT.
integer j

```
    integer, parameter :: n=40
    real(kind(1e0)) :: one=1e0
    real(kind(1e0)) err
    real(kind(le0)), dimension(n) :: x, y, yy(n,n)
    complex(kind(le0)), dimension(n) :: a, b, c, d, e, f
! Generate two random complex sequence 'a' and 'b'.
    call rand_gen(x)
    call rand_gen(y)
    a=x; b=y
! Compute the convolution 'c' of 'a' and 'b'.
! Use matrix times vector for test results.
    Yy(1:,1)=y
    do j=2,n
        YY (2:,j)=YY (1:n-1,j-1)
        YY(1,j)=YY(n,j-1)
    end do
    c=matmul (yy,x)
! Transform the 'a' and 'b' sequences into 'd' and 'e'.
    call c_fast_dft(forward_in=a, &
        forwar\overline{d_out=d)}
    call c_fast_\overline{dft(forward_in=b, &}
        forwar\overline{d_out=e)}
! Invert the product d*e.
    call c_fast_dft(inverse_in=d*e, &
        in
! Check the Convolution Theorem:
! inverse(transform(a)*transform(b)) = convolution(a,b).
    err = maxval(abs(c-f))/maxval(abs(c))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 4 for FAST_DFT is correct.'
    end if
    end
```


## Output

Example 4 for FAST_DFT is correct.

## FAST_2DFT

Computes the Discrete Fourier Transform (2DFT) of a rank-2 complex array, $x$.

## Required Arguments

No required arguments; pairs of optional arguments are required. These pairs are forward_in and forward_out or inverse_in and inverse_out.

## Optional Arguments

```
forward_in = x (Input)
```

Stores the input complex array of rank-2 to be transformed.
forward_out = y (Output)
Stores the output complex array of rank-2 resulting from the transform.
inverse_in = y (Input)
Stores the input complex array of rank-2 to be inverted.
inverse_out $=x$ (Output)
Stores the output complex array of rank-2 resulting from the inverse transform.
mdata $=m$ (Input)
Uses the sub-array in first dimension of size $m$ for the numbers.
Default value: $m=\operatorname{size}(x, 1)$.
ndata $=\mathrm{n}$ (Input)
Uses the sub-array in the second dimension of size $n$ for the numbers.
Default value: $n=\operatorname{size}(x, 2)$.
ido = ido (Input/Output)
Integer flag that directs user action. Normally, this argument is used only when the working variables required for the transform and its inverse are saved in the calling program unit. Computing the working variables and saving them in internal arrays within fast_2dft is the default. This initialization step is expensive.

There is a two-step process to compute the working variables just once. Example 3 illustrates this usage. The general algorithm for this usage is to enter fast_2dft with $\mathrm{ido}=0$. A return occurs thereafter with ido $<0$. The optional rank-1 complex array $\mathrm{w}(:)$ with $\operatorname{size}(\mathrm{w})>=-$ ido must be re-allocated. Then, re-enter fast_2dft. The next return from fast_2dft has the output value ido $=1$. The variables required for the transform and its inverse are saved in $\mathrm{w}(:)$. Thereafter, when the routine is entered with ido $=1$ and for the same values of m and n , the contents of $\mathrm{w}(:)$ will be used for the working variables. The expensive initialization step is avoided. The optional arguments "ido=" and "work_array=" must be used together.
work_array = w(:) (Output/Input)
Complex array of rank-1 used to store working variables and values between calls to fast_2dft. The value for size(w) must be at least as large as the value - ido for the value of ido $<0$.
iopt $=$ iopt (:) (Input/Output)
Derived type array with the same precision as the input array; used for passing optional data to fast_2dft. The options are as follows:

| Packaged Options for FAST_2DFT |  |  |
| :---: | :--- | :---: |
| Option Prefix =? | Option Name | Option Value |
| $c_{-}, z_{-}$ | fast_2dft_scan_for_NaN | 1 |
| $c_{-}, z_{-}$ | fast_2dft_near_power_of_2 | 2 |
| $c_{-}, z_{-}$ | fast_2dft_scale_forward | 3 |
| $c_{-}, z_{-}$ | fast_2dft_scale_inverse | 4 |

iopt(IO) $=$ ?_options(?_fast_2dft_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that isNaN(x(i,j)) ==.true.
See the isNaN () function, Chapter 10.
Default: Does not scan for NaNs.
iopt (IO) $=$ ? _options(?_fast_2dft_near_power_of_2, ?_dummy)
Nearest powers of $2 \geq m$ and $\geq n$ are returned as an outputs in iopt (IO +

1) \%idummy and iopt (IO + 2) \%idummy.
iopt(IO) = ?_options(?_fast_2dft_scale_forward, real_part_of_scale)
iopt(IO+1) = ?_options(?_dummy, imaginary_part_of_scale)
Complex number defined $\bar{b} y$ the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is
multiplied by the forward transformed array.
Default value is 1 .
iopt(IO) $=$ ? _options(?_fast_2dft_scale_inverse, real_part_of_scale)
iopt(IO+1) = ?_options(?_dummy, imaginary_part_of_scale)
Complex number defined by the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is
multiplied by the inverse transformed array.
Default value is 1 .

## FORTRAN 90 Interface

Generic: None

Specific: The specific interface names are S_FAST_2DFT, D_FAST_2DFT, C_FAST_2DFT, and Z_FAST_2DFT.

## Description

The fast_2dft routine is a Fortran 90 version of the FFT suite of IMSL (1994, pp. 772-776).

## Fatal and Terminal Messages

See the messages.gls file for error messages for FAST_2DET. These error messages are numbered 670-680; 720-730.

## Example 1: Transforming an Array of Random Complex Numbers

An array of random complex numbers is obtained. The transform of the numbers is inverted and the final results are compared with the input array.

```
    use fast_2dft_int
    use rand_int
    implicit none
! This is Example 1 for FAST_2DFT.
    integer, parameter :: n=24
    integer, parameter :: m=40
    real(kind(1e0)) :: err, one=1e0
    complex(kind(le0)), dimension(n,m) :: a, b, c
! Generate a random complex sequence.
    a=rand(a); c=a
! Transform and then invert the transform.
    call c_fast_2dft(forward_in=a, &
        forward out=b)
    call c_fast_2dft(inverse_in=b, &
            inverse_out=a)
! Check that inverse(transform(sequence)) = sequence.
    err = maxval(abs(c-a))/maxval(abs(c))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for FAST_2DFT is correct.'
    end if
    end
```


## Output

Example 1 for FAST_2DFT is correct.

## Additional Examples

## Example 2: Cyclical 2D Data with a Linear Trend

This set of data is sampled from a function $x(s, t)=a+b s+c t+y(s, t)$, where $y(s, t)$ is an harmonic series. The independent variables are normalized as $-1 \leq s \leq 1$ and $-1 \leq t \leq 1$. Thus, the data is said to have cyclical components plus a linear trend. As a first step, the linear terms are effectively removed from the data using the least-squares system solver . Then, the residuals are transformed and the resulting frequencies are analyzed.

```
    use fast 2dft int
    use lin_sol_lsq_int
use sort real int
use rand_int
implicit none
! This is Example 2 for FAST_2DFT.
    integer i
    integer, parameter :: n=8, k=15
    integer ip(n*n), order(k)
    real(kind(1e0)), parameter :: one=1e0, two=2e0, zero=0e0
    real(kind(1e0)) delta_t
    real(kind(le0)) rn(3), s(n), t(n), temp(n*n), new_order(k)
    complex(kind(le0)) a, b, c, a_trend(n*n,3), b_trend(n*n,1), &
            f(n,n), r(n,n), x(n,n), x_trend(3,1)
    complex(kind(le0)), dimension(n,n)
! Generate random data for planar trend.
    rn = rand(rn)
    a = rn(1)
    b}=rn(2
    c}=rn(3
! Generate the frequency components of the harmonic series.
! Non-zero random amplitudes given on two edges of the square domain.
    g(1:,1)=rand (g(1:,1))
    g(1,1:)=rand (g(1,1:))
! Invert 'g' into the harmonic series 'h' in time domain.
    call c_fast_2dft(inverse_in=g, inverse_out=h)
! Compute sampling interval.
    delta t = two/n
    s = (// (-one + (i-1)*delta_t, i=1,n)/)
    t = (/(-one + (i-1)*delta_t, i=1,n)/)
! Make up data set as a linear trend plus harmonics.
    x = a + b*spread(s,dim=2,ncopies=n) + &
        c*spread(t,dim=1,ncopies=n) + h
! Define least-squares matrix data for a planar trend.
    a_trend(1:,1) = one
```

```
    a_trend(1:,2) = reshape(spread(s,dim=2,ncopies=n),(/n*n/))
    a_trend(1:,3) = reshape(spread(t,dim=1,ncopies=n),(/n*n/))
    b_trend(1:,1) = reshape (x,(/n*n/))
! Solve for a linear trend.
    call lin_sol_lsq(a_trend, b_trend, x_trend)
! Compute harmonic residuals.
    r = x - reshape(matmul(a_trend,x_trend),(/n,n/))
! Transform harmonic residuals.
    call c_fast_2dft(forward_in=r, forward_out=f)
    ip = (/(i,i=1,n**2)/)
! Sort the magnitude of the transform.
    call s_sort_real(-(abs(reshape(f,(/n*n/)))), &
                                    temp, iperm=ip)
! The dominant frequencies are output in ip(1:k).
! Sort these values to compare with the original frequency order.
    call s_sort_real(real(ip(1:k)), new_order)
    order(1:n) = (/ (i,i=1,n)/)
    order (n+1:k) = (/ ((i-n)*n+1,i=n+1,k)/)
! Check the results.
    if (count(order /= int(new_order)) == 0) then
        write (*,*) 'Example 2 for FAST_2DFT is correct.'
    end if
    end
```


## Output

Example 2 for FAST_2DFT is correct.

## Example 3: Several 2D Transforms with Initialization

In this example, the optional arguments ido and work_array are used to save working variables in the calling program unit. This results in maximum efficiency of the transform and its inverse since the working variables do not have to be precomputed following each entry to routine fast_2dft.
use fast_2dft_int
implicit none
! This is Example 3 for FAST 2DFT.
integer i, j
integer, parameter :: n=256
real(kind(1e0)), parameter :: one=1e0, zero=0e0
real(kind(1e0)) r(n,n), err

```
    complex(kind(1e0)) a(n,n), b(n,n), c(n,n)
! The value of the array size for work(:) is computed in the
! routine fast_dft as a first step.
    integer ido_value
    complex(kind(le0)), allocatable :: work(:)
! Fill in value one for points inside the circle with r=64.
    a = zero
    r = reshape((/(((i-n/2)**2 + (j-n/2)**2, i=1,n), &
        j=1,n)/),(/n,n/))
    where (r <= (n/4)**2) a = one
    c = a
! Transform and then invert the sequence using the pre-computed
! working values.
    ido value = 0
    do
        if(allocated(work)) deallocate(work)
! Allocate the space required for work(:).
            if (ido_value <= 0) allocate(work(-ido_value))
! Transform the image and then invert it back.
    call c_fast_2dft(forward_in=a, &
            forrwar\overline{d_out=b, IDO=\overline{i}do_value, work_array=work)}
            if (ido_value == 1) exit
    end do
    call c_fast_2dft(inverse_in=b, &
            inverse_out=a, IDO=ido_value, work_array=work)
! Deallocate the space used for work(:).
    if (allocated(work)) deallocate(work)
! Check that inverse(transform(image)) = image.
    err = maxval(abs(c-a))/maxval(abs(c))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 3 for FAST_2DFT is correct.'
    end if
    end
```


## Output

Example 3 for FAST_2DFT is correct.

## FAST_3DFT

Computes the Discrete Fourier Transform (2DFT) of a rank-3 complex array.

## Required Arguments

No required arguments; pairs of optional arguments are required. These pairs are forward_in and forward_out or inverse_in and inverse_out.

## Optional Arguments

```
forward_in = x (Input)
```

Stores the input complex array of rank-3 to be transformed.
forward_out = y (Output)
Stores the output complex array of rank-3 resulting from the transform.
inverse_in = y (Input)
Stores the input complex array of rank- 3 to be inverted.
inverse_out $=x$ (Output)
Stores the output complex array of rank-3 resulting from the inverse transform.
mdata $=m$ (Input)
Uses the sub-array in first dimension of size $m$ for the numbers.
Default value: $\mathrm{m}=\operatorname{size}(\mathrm{x}, 1)$.
ndata $=n$ (Input)
Uses the sub-array in the second dimension of size $n$ for the numbers.
Default value: $\mathrm{n}=\operatorname{size}(\mathrm{x}, 2)$.
kdata $=\mathrm{k}$ (Input)
Uses the sub-array in the third dimension of size $k$ for the numbers.
Default value: $k=\operatorname{size}(x, 3)$.
ido = ido (Input/Output)
Integer flag that directs user action. Normally, this argument is used only when the working variables required for the transform and its inverse are saved in the calling program unit. Computing the working variables and saving them in internal arrays within fast_3dft is the default. This initialization step is expensive.

There is a two-step process to compute the working variables just once. The general algorithm for this usage is to enter fast_3dft with ido $=0$. A return occurs thereafter with ido $<0$. The optional rank-1 complex array $\mathrm{w}(:)$ with $\operatorname{size}(\mathrm{w})>=-$ ido must be re-allocated. Then, re-enter fast_3dft. The next return from fast_3dft has the output value ido $=1$. The variables required for the transform and its inverse are saved in $\mathrm{w}(:)$. Thereafter, when the routine is entered with ido $=1$ and for the same values of $m$ and $n$, the contents of $w(:)$ will be used for the working variables. The expensive initialization step is avoided. The optional arguments "ido=" and "work_array=" must be used together.
work_array = w(:) (Output/Input)
Complex array of rank-1 used to store working variables and values between calls to
fast_3dft. The value for size(w) must be at least as large as the value - ido for the value of ido $<0$.
iopt $=$ iopt (:) (Input/Output)
Derived type array with the same precision as the input array; used for passing optional data to fast_3dft. The options are as follows:

| Packaged Options for FAST_3DFT |  |  |
| :---: | :--- | :---: |
| Option Prefix $=?$ | Option Name | Option Value |
| C_, $_{-}$ | fast_3dft_scan_for_NaN | 1 |
| C_, $_{-}$ | fast_3dft_near_power_of_2 | 2 |
| C_, $_{-}$ | fast_3dft_scale_forward | 3 |
| C_, $_{-}$ | fast_3dft_scale_inverse | 4 |

iopt(IO) = ?_options(?_fast_3dft_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that isNaN(x(i,j,k)) ==.true.
See the isnan () function, Chapter 10.
Default: Does not scan for NaNs.
iopt (IO) = ?_options(?_fast_3dft_near_power_of_2, ?_dummy)
Nearest powers of $2 \geq \bar{m}, \geq n$, and $\geq \bar{k}$ are returned as an outputs in
iopt (IO 1 ) \%idummy, iopt (IO +2 ) \%idummy and iopt(IO +3 ) \%idummy

```
iopt(IO) = ?_options(?_fast_3dft_scale_forward, real_part_of_scale)
```

iopt(IO+1) = ? options(?_dummy, imaginary_part_of_scale)

Complex number defined by the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is multiplied by the forward transformed array.
Default value is 1 .
iopt(IO) = ? options(? fast_3dft_scale_inverse, real_part_of_scale)
iopt(IO+1) = ?_options(?_dummy, imaginary_part_of_scale)
Complex number defined by the factor
cmplx(real_part_of_scale, imaginary_part_of_scale) is
multiplied by the inverse transformed array.
Default value is 1 .

## FORTRAN 90 Interface

Generic: None

Specific: The specific interface names are S_FAST_3DFT, D_FAST_3DFT, C_FAST_3DFT, and Z_FAST_3DFT.

## Description

The fast_3dft routine is a Fortran 90 version of the FFT suite of IMSL (1994, pp. 772-776).

## Fatal and Terminal Messages

See the messages.gls file for error messages for FAST_3DFT. These error messages are numbered 685-695; 740-750.

## Example: Transforming an Array of Random Complex Numbers

An array of random complex numbers is obtained. The transform of the numbers is inverted and the final results are compared with the input array.

```
    use fast_3dft_int
    implicit none
! This is Example 1 for FAST_3DFT.
    integer i, j, k
    integer, parameter :: n=64
    real(kind(le0)), parameter :: one=1e0, zero=0e0
    real(kind(le0)) r(n,n,n), err
    complex(kind(1e0)) a(n,n,n), b(n,n,n), c(n,n,n)
! Fill in value one for points inside the sphere
! with radius=16.
    a = zero
    do i=1,n
        do j=1,n
            do k=1,n
                r(i,j,k)=(i-n/2)**2+(j-n/2)**2+(k-n/2)**2
            end do
        end do
    end do
    where (r <= (n/4)**2) a = one
    c = a
! Transform the image and then invert it back.
    call c_fast_3dft(forward_in=a, &
        for}ward_out=b
    call c_fast_3dft(inverse_in=b, &
            inverse_out=a)
! Check that inverse(transform(image)) = image.
    err = maxval(abs(c-a))/maxval(abs(c))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for FAST_3DFT is correct.'
    end if
    end
```


## Output

Example 1 for FAST_3DFT is correct.

## FFTRF

Computes the Fourier coefficients of a real periodic sequence.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
SEQ - Array of length N containing the periodic sequence. (Input)
COEF - Array of length N containing the Fourier coefficients. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTRF ( $\mathrm{N}, \mathrm{SEQ}$, COEF)
Specific: The specific interface names are S_FFTRF and D_FFTRF.

## FORTRAN 77 Interface

Single: CALL FFTRF (N, SEQ, COEF)
Double: The double precision name is DFFTRF.

## Description

The routine FFTRF computes the discrete Fourier transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm that is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an $N$-vector $s=$ SEQ, FFTRF returns in $c=$ COEF, if $N$ is even:

$$
\begin{aligned}
c_{2 m-2} & =\sum_{n=1}^{N} s_{n} \cos \left[\frac{(m-1)(n-1) 2 \pi}{N}\right] \quad m=2, \ldots, N / 2+1 \\
c_{2 m-1} & =-\sum_{n=1}^{N} s_{n} \sin \left[\frac{(m-1)(n-1) 2 \pi}{N}\right] \quad m=2, \ldots, N / 2 \\
c_{1} & =\sum_{n=1}^{N} s_{n}
\end{aligned}
$$

If $N$ is odd, $c_{m}$ is defined as above for $m$ from 2 to $(N+1) / 2$.

We now describe a fairly common usage of this routine. Let $f$ be a real valued function of time. Suppose we sample $f$ at $N$ equally spaced time intervals of length $\Delta$ seconds starting at time $t_{0}$. That is, we have

$$
\mathrm{SEQ}_{i}:=f\left(t_{0}+(i-1) \Delta\right) i=1,2, \ldots, N
$$

The routine FFTRF treats this sequence as if it were periodic of period $N$. In particular, it assumes that $f\left(t_{0}\right)=f\left(t_{0}+N \Delta\right)$. Hence, the period of the function is assumed to be $T=N \Delta$.

Now, FFTRF accepts as input SEQ and returns as output coefficients $c=$ COEF that satisfy the following relation when $N$ is odd ( $N$ even is similar):

$$
\mathrm{SEQ}_{i}=\frac{1}{N}\left[c_{1}+2 \sum_{n=2}^{(N+1) / 2} c_{2 n-2} \cos \left[\frac{2 \pi(n-1)(i-1)}{N}\right]-2 \sum_{n=2}^{(N+1) / 2} c_{2 n-1} \sin \left[\frac{2 \pi(n-1)(i-1)}{N}\right]\right]
$$

This formula is very revealing. It can be interpreted in the following manner. The coefficients produced by FFTRF produce an interpolating trigonometric polynomial to the data. That is, if we define

$$
\begin{aligned}
g(t) & :=\frac{1}{N}\left[c_{1}+2 \sum_{n=2}^{(N+1) / 2} c_{2 n-2} \cos \left[\frac{2 \pi(n-1)\left(t-t_{0}\right)}{N \Delta}\right]-2 \sum_{n=2}^{(N+1) / 2} c_{2 n-1} \sin \left[\frac{2 \pi(n-1)\left(t-t_{0}\right)}{N \Delta}\right]\right] \\
& =\frac{1}{N}\left[c_{1}+2 \sum_{n=2}^{(N+1) / 2} c_{2 n-2} \cos \left[\frac{2 \pi(n-1)\left(t-t_{0}\right)}{T}\right]-2 \sum_{n=2}^{(N+1) / 2} c_{2 n-1} \sin \left[\frac{2 \pi(n-1)\left(t-t_{0}\right)}{T}\right]\right]
\end{aligned}
$$

then, we have

$$
f\left(t_{0}+(i-1) \Delta\right)=g\left(t_{0}+(i-1) \Delta\right)
$$

Now, suppose we want to discover the dominant frequencies. One forms the vector $P$ of length $N / 2$ as follows:

$$
\begin{aligned}
P_{1} & :=\left|c_{1}\right| \\
P_{k} & :=\sqrt{c_{2 k-2}^{2}+c_{2 k-1}^{2}} \quad k=2,3, \ldots,(N+1) / 2
\end{aligned}
$$

These numbers correspond to the energy in the spectrum of the signal. In particular, $P_{k}$ corresponds to the energy level at frequency

$$
\frac{k-1}{T}=\frac{k-1}{N \Delta} \quad k=1,2, \ldots, \frac{N+1}{2}
$$

Furthermore, note that there are only $(N+1) / 2 \approx T /(2 \Delta)$ resolvable frequencies when $N$ observations are taken. This is related to the Nyquist phenomenon, which is induced by discrete sampling of a continuous signal.

Similar relations hold for the case when $N$ is even.
Finally, note that the Fourier transform hsas an (unnormalized) inverse that is implemented in FFTRB. The routine FFTRF is based on the real FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2TRF / DF2TRF. The reference is:

CALL F2TRF (N, SEQ, COEF, WFFTR)
The additional argument is
WFFTR - Array of length $2 \mathrm{~N}+15$ initialized by FFTRI. (Input) The initialization depends on N .
2. The routine FFTRF is most efficient when $N$ is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If FFTRF/FFTRB is used repeatedly with the same value of $N$, then call FFTRI followed by repeated calls to F2TRF/F2TRB. This is more efficient than repeated calls to FFTRF/FFTRB.

## Example

In this example, a pure cosine wave is used as a data vector, and its Fourier series is recovered. The Fourier series is a vector with all components zero except at the appropriate frequency where it has an $N$.

```
USE FFTRF INT
USE CONST - INT
USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=7)
    INTEGER I, NOUT
    REAL COEF(N), COS, FLOAT, TWOPI, SEQ(N)
    INTRINSIC COS, FLOAT
    TWOPI = CONST('PI')
    TWOPI = 2.0*TWOPI
    CALL UMACH (2, NOUT)
    DO 10 I=1, N
        SEQ(I) = COS(FLOAT(I-1)*TWOPI/FLOAT(N))
    CONTINUE
    CALL FFTRF (N, SEQ, COEF)
    WRITE (NOUT,99998)
99998 FORMAT (9X, 'INDEX', 5X, 'SEQ', 6X, 'COEF')
```

!

```
        WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99999 FORMAT (1X, I11, 5X, F5.2, 5X, F5.2)
        END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 0.00 |
| 2 | 0.62 | 3.50 |
| 3 | -0.22 | 0.00 |
| 4 | -0.90 | 0.00 |
| 5 | -0.90 | 0.00 |
| 6 | -0.22 | 0.00 |
| 7 | 0.62 | 0.00 |

## FFTRB

Computes the real periodic sequence from its Fourier coefficients.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
COEF - Array of length n containing the Fourier coefficients. (Input)
SEQ - Array of length N containing the periodic sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTRB (N, COEF, SEQ [, ...])

Specific: The specific interface names are S_FFTRB and D_FFTRB.

## FORTRAN 77 Interface

Single: CALL FFTRB ( $\mathrm{N}, \mathrm{COEF}, \mathrm{SEQ}$ )
Double: The double precision name is DFFTRB.

## Description

The routine FFTRB is the unnormalized inverse of the routine FFTRF. This routine computes the discrete inverse Fourier transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an $N$-vector $c=$ COEF, FFTRB returns in $s=$ SEQ, if $N$ is even:

$$
\begin{aligned}
s_{m}=c_{1}+(-1)^{(m-1)} c_{N} & +2 \sum_{n=2}^{N / 2} c_{2 n-2} \cos \frac{[(n-1)(m-1) 2 \pi]}{N} \\
& -2 \sum_{n=2}^{N / 2} c_{2 n-1} \sin \frac{[(n-1)(m-1) 2 \pi]}{N}
\end{aligned}
$$

If $N$ is odd:

$$
\begin{aligned}
s_{m}=c_{1}+ & 2 \sum_{n=2}^{(N+1) / 2} c_{2 n-2} \cos \frac{[(n-1)(m-1) 2 \pi]}{N} \\
& -2 \sum_{n=2}^{(N+1) / 2} c_{2 n-1} \sin \frac{[(n-1)(m-1) 2 \pi]}{N}
\end{aligned}
$$

The routine FFTRB is based on the inverse real FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2TRB/DF2TRB. The reference is:

CALL F2TRB (N, COEF, SEQ, WFFTR)
The additional argument is
WFFTR - Array of length $2 \mathrm{~N}+15$ initialized by FFTRI. (Input)
The initialization depends on N .
2. The routine FFTRB is most efficient when N is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If $\operatorname{FFTRF} /$ FFTRB is used repeatedly with the same value of N , then call FFTRI followed by repeated calls to $\operatorname{F2TRF} / \mathrm{F} 2 \mathrm{TRB}$. This is more efficient than repeated calls to FFTRF/FFTRB.

## Example

We compute the forward real FFT followed by the inverse operation. In this example, we first compute the Fourier transform

$$
\hat{x}=\mathrm{COEF}
$$

of the vector $x$, where $x_{j}=(-1)^{j}$ for $j=1$ to $N$. This vector

$$
\hat{x}
$$

is now input into FFTRB with the resulting output $s=N x$, that is, $s_{j}=(-1)^{j} N$ for $j=1$ to $N$.

```
USE FFTRB INT
```

```
USE CONST_INT
USE FFTRF INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
!
INTEGER I, NOUT
REAL COEF(N), FLOAT, SEQ(N), TWOPI, X(N)
INTRINSIC FLOAT
TWOPI = CONST('PI')
!
TWOPI = TWOPI
! Get output unit number
CALL UMACH (2, NOUT)
DO 10 I=1, N
    X(I) = FLOAT((-1)**I)
    CONTINUE
CALL FFTRF (N, X, COEF)
WRITE (NOUT,99994)
WRITE (NOUT,99995)
99994 FORMAT (9X, 'Result after forward transform')
99995 FORMAT (9X, 'INDEX', 5X, 'X', 8X, 'COEF')
WRITE (NOUT,99996) (I, X(I), COEF(I), I=1,N)
99996 FORMAT (1X, I11, 5X, F5.2, 5X, F5.2)
! Compute the backward transform of
CALL FFTRB (N, COEF, SEQ)
Print results
WRITE (NOUT,99997)
WRITE (NOUT,99998)
99997 FORMAT (/, 9X, 'Result after backward transform')
99998 FORMAT (9X, 'INDEX', 4X, 'COEF', 6X, 'SEQ')
WRITE (NOUT,99999) (I, COEF(I), SEQ(I), I=1,N)
99999 FORMAT (1X, I11, 5X, F5.2, 5X, F5.2)
END
```


## Output

| Result | after | forward transform |
| :---: | ---: | :---: |
| INDEX | X | COEF |
| 1 | -1.00 | -1.00 |
| 2 | 1.00 | -1.00 |
| 3 | -1.00 | -0.48 |
| 4 | 1.00 | -1.00 |
| 5 | -1.00 | -1.25 |
| 6 | 1.00 | -1.00 |
| 7 | -1.00 | -4.38 |

Result after backward transform
INDEX COEF SEQ

| 1 | -1.00 | -7.00 |
| :--- | :--- | ---: |
| 2 | -1.00 | 7.00 |
| 3 | -0.48 | -7.00 |
| 4 | -1.00 | 7.00 |
| 5 | -1.25 | -7.00 |
| 6 | -1.00 | 7.00 |
| 7 | -4.38 | -7.00 |

## FFTRI

Computes parameters needed by FFTRF and FFTRB.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
WFFTR - Array of length $2 \mathrm{~N}+15$ containing parameters needed by FFTRF and FFTRB. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTRI (N, WFFTR)

Specific: The specific interface names are S_FFTRI and D_FFTRI.

## FORTRAN 77 Interface

Single: CALL FFTRI (N, WFFTR)
Double: The double precision name is DFFTRI.

## Description

The routine FFTRI initializes the routines FFTRF and FFTRB. An efficient way to make multiple calls for the same $N$ to routine FFTRF or FFTRB, is to use routine FFTRI for initialization. (In this case, replace FFTRF or FFTRB with F2TRF or F2TRB, respectively.) The routine FFTRI is based on the routine RFFTI in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different WFFTR arrays are needed for different values of N .

## Example

In this example, we compute three distinct real FFTs by calling FFTRI once and then calling F2TRF three times.

```
    USE FFTRI_INT
    USE CONST INT
    USE F2TRF INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=7)
!
    INTEGER I, K, NOUT
    REAL COEF(N), COS, FLOAT, TWOPI, WFFTR(29), SEQ(N)
    INTRINSIC COS, FLOAT
!
    TWOPI = CONST('PI')
    TWOPI = 2* TWOPI
    Set the work vector
    CALL FFTRI (N, WFFTR)
    DO 20 K=1, 3
! This loop fills out the data vector
with a pure exponential signal
    DO 10 I=1, N
        SEQ(I) = COS (FLOAT (K* (I-1))*TWOPI/FLOAT (N))
    CONTINUE
    CALL F2TRF (N, SEQ, COEF, WFFTR)
    WRITE (NOUT, 99998)
    FORMAT (/, 9X, 'INDEX', 5X, 'SEQ', 6X, 'COEF')
    WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99999 FORMAT (1X, I11, 5X, F5.2, 5X, F5.2)
!
    20 CONTINUE
END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 0.00 |
| 2 | 0.62 | 3.50 |
| 3 | -0.22 | 0.00 |
| 4 | -0.90 | 0.00 |
| 5 | -0.90 | 0.00 |
| 6 | -0.22 | 0.00 |
| 7 | 0.62 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 1.00 | 0.00 |
| 2 | -0.22 | 0.00 |
| 3 | -0.90 | 0.00 |
| 4 | 0.62 | 3.50 |
| 5 | 0.62 | 0.00 |


| 6 | -0.90 | 0.00 |
| :--- | :---: | ---: |
| 7 | -0.22 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 1.00 | 0.00 |
| 2 | -0.90 | 0.00 |
| 3 | 0.62 | 0.00 |
| 4 | -0.22 | 0.00 |
| 5 | -0.22 | 0.00 |
| 6 | 0.62 | 3.50 |
| 7 | -0.90 | 0.00 |

## FFTCF

Computes the Fourier coefficients of a complex periodic sequence.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
$\mathbf{S E Q}$ - Complex array of length N containing the periodic sequence. (Input)
COEF - Complex array of length N containing the Fourier coefficients. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTCF (N, SEQ, COEF)

Specific: The specific interface names are S_FFTCF and D_FFTCF.

## FORTRAN 77 Interface

Single: CALL FFTCF (N, SEQ, COEF)
Double: The double precision name is DFFTCF.

## Description

The routine FFTCF computes the discrete complex Fourier transform of a complex vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N$-vector $x$, FFTCF returns in $c=$ COEF

$$
c_{m}=\sum_{n=1}^{N} x_{n} e^{-2 \pi i(n-1)(m-1) / N}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
\sqrt{N} S
$$

Finally, note that we can invert the Fourier transform as follows:

$$
x_{n}=\frac{1}{N} \sum_{m=1}^{N} c_{m} e^{2 \pi i(m-1)(n-1) / N}
$$

This formula reveals the fact that, after properly normalizing the Fourier coefficients, one has the coefficients for a trigonometric interpolating polynomial to the data. An unnormalized inverse is implemented in FFTCB. FFTCF is based on the complex FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2TCF/DF2TCF. The reference is:

CALL F2TCF (N, SEQ, COEF, WFFTC, CPY)
The additional arguments are as follows:
WFFTC - Real array of length $4 * \mathrm{~N}+15$ initialized by FFTCI. The initialization depends on $N$. (Input)
$\boldsymbol{C P Y}$ - Real array of length $2 *$ n. (Workspace)
2. The routine FFTCF is most efficient when N is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If FFTCF/FFTCB is used repeatedly with the same value of N , then call FFTCI followed by repeated calls to $\mathrm{F} 2 \mathrm{TCF} / \mathrm{F} 2 \mathrm{TCB}$. This is more efficient than repeated calls to FFTCF/FFTCB.

## Example

In this example, we input a pure exponential data vector and recover its Fourier series, which is a vector with all components zero except at the appropriate frequency where it has an $N$. Notice that the norm of the input vector is

$$
\sqrt{N}
$$

but the norm of the output vector is $N$.

```
USE FFTCF_INT
USE CONST_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
```

```
    PARAMETER (N=7)
    INTEGER I, NOUT
    REAL TWOPI
    COMPLEX C, CEXP, COEF(N), H, SEQ(N)
    INTRINSIC CEXP
!
    C = (0.,1.)
    TWOPI = CONST('PI')
    TWOPI = 2.0 * TWOPI
    H}=(TWOPI*C/N)*3
        DO 10 I=1, N
        SEQ(I) = CEXP((I-1)*H)
    1 0 ~ C O N T I N U E
! Compute the Fourier transform of SEQ
    CALL FFTCF (N, SEQ, COEF)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
99998 FORMAT (9X, 'INDEX', 8X, 'SEQ', 15X, 'COEF')
    WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99999 FORMAT (1X, I11, 5X,'(',F5.2,',',F5.2,')', &
    5X,'(',F5.2,',',F5.2,')')
    END
```


## Output

| INDEX | SEQ | COEF |
| :---: | :---: | :---: |
| 1 | $(1.00,0.00)$ | $(0.00,0.00)$ |
| 2 | $(-0.90,0.43)$ | $(0.00,0.00)$ |
| 3 | $(0.62,-0.78)$ | $(0.00,0.00)$ |
| 4 | $(-0.22,0.97)$ | $(7.00,0.00)$ |
| 5 | $(-0.22,-0.97)$ | $(0.00,0.00)$ |
| 6 | $(0.62,0.78)$ | $(0.00,0.00)$ |
| 7 | $(-0.90,-0.43)$ | $(0.00,0.00)$ |

## FFTCB

Computes the complex periodic sequence from its Fourier coefficients.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)

COEF - Complex array of length N containing the Fourier coefficients. (Input)
SEQ - Complex array of length N containing the periodic sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTCB ( $\mathrm{N}, \mathrm{COEF}, \mathrm{SEQ}$ )
Specific: The specific interface names are S_FFTCB and D_FFTCB.

## FORTRAN 77 Interface

Single: CALL FFTCB (N, COEF, SEQ)

Double: $\quad$ The double precision name is DFFTCB.

## Description

The routine FFTCB computes the inverse discrete complex Fourier transform of a complex vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N$-vector $c=$ COEF, FFTCB returns in $s=$ SEQ

$$
s_{m}=\sum_{n=1}^{N} c_{n} e^{2 \pi i(n-1)(m-1) / N}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
\sqrt{N} S
$$

Finally, note that we can invert the inverse Fourier transform as follows:

$$
c_{n}=\frac{1}{N} \sum_{m=1}^{N} s_{m} e^{-2 \pi i(n-1)(m-1) / N}
$$

This formula reveals the fact that, after properly normalizing the Fourier coefficients, one has the coefficients for a trigonometric interpolating polynomial to the data. FFTCB is based on the complex inverse FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of $\operatorname{F2TCB} / D F 2 T C B$. The reference is:

CALL F2TCB (N, COEF, SEQ, WFFTC, CPY)
The additional arguments are as follows:
WFFTC - Real array of length $4 * \mathrm{~N}+15$ initialized by FFTCI. The initialization depends on N. (Input)
$\boldsymbol{C P Y}$ - Real array of length $2 *$ N. (Workspace)
2. The routine FFTCB is most efficient when N is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If $\operatorname{FFTCF} / \mathrm{FFTCB}$ is used repeatedly with the same value of N ; then call FFTCI followed by repeated calls to F2TCF/F2TCB. This is more efficient than repeated calls to FFTCF/FFTCB.

## Example

In this example, we first compute the Fourier transform of the vector $x$, where $x_{j}=j$ for $j=1$ to $N$. Note that the norm of $x$ is $(N[N+1][2 N+1] / 6)^{1 / 2}$, and hence, the norm of the transformed vector

$$
\hat{x}=c
$$

is $N([N+1][2 N+1] / 6)^{1 / 2}$. The vector

$$
\hat{x}
$$

is used as input into FFTCB with the resulting output $s=N x$, that is, $s_{j}=j N$, for $j=1$ to $N$.

```
USE FFTCB INT
USE FFTCF }\mp@subsup{}{}{-}\mathrm{ INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
    INTEGER I, NOUT
    COMPLEX CMPLX, SEQ(N), COEF(N), X(N)
    INTRINSIC CMPLX
DO 10 I=1, N
        X(I) = CMPLX(I,0)
    CONTINUE
    CALL FFTCF (N, X, COEF)
    CALL FFTCB (N, COEF, SEQ)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
    WRITE (NOUT,99999) (I, X(I), COEF(I), SEQ(I), I=1,N)
99998 FORMAT (5X, 'INDEX', 9X, 'INPUT', 9X, 'FORWARD TRANSFORM', 3X, &
    'BACKWARD TRANSFORM')
99999 FORMAT (1X, I7, 7X,'(',F5.2,',',F5.2,')', &
    7X,'(',F5.2,',',F5.2,')', &
    7X,'(',F5.2,',',F5.2,')')
```

!

## Output

| INDEX | INPUT | FORWARD TRANSFORM | BACKWARD TRANSFORM |
| :---: | :---: | :---: | :---: |
| 1 | $(1.00,0.00)$ | $(28.00,0.00)$ | $(7.00,0.00)$ |
| 2 | $(2.00,0.00)$ | $(-3.50,7.27)$ | $(14.00,0.00)$ |
| 3 | $(3.00,0.00)$ | $(-3.50,2.79)$ | $(21.00,0.00)$ |
| 4 | $(4.00,0.00)$ | $(-3.50,0.80)$ | $(28.00,0.00)$ |
| 5 | $(5.00,0.00)$ | $(-3.50,-0.80)$ | $(35.00,0.00)$ |
| 6 | $(6.00,0.00)$ | $(-3.50,-2.79)$ | $(42.00,0.00)$ |
| 7 | $(7.00,0.00)$ | $(-3.50,-7.27)$ | $(49.00,0.00)$ |

## FFTCI

Computes parameters needed by FFTCF and FFTCB.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
WFFTC - Array of length $4 \mathrm{~N}+15$ containing parameters needed by FFTCF and FFTCB. (Output)

## FORTRAN 90 Interface

Generic: CALL FFTCI (N, WFFTC)

Specific: The specific interface names are S_FFTCI and D_FFTCI.

## FORTRAN 77 Interface

Single: CALL FFTCI (N, WFFTC)
Double: The double precision name is DFFTCI.

## Description

The routine FFTCI initializes the routines FFTCF and FFTCB. An efficient way to make multiple calls for the same N to IMSL routine FFTCF or FFTCB is to use routine FFTCI for initialization. (In this case, replace FFTCF or FFTCB with F2TCF or F2TCB, respectively.) The routine FFTCI is based on the routine CFFTI in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different WFFTC arrays are needed for different values of N .

## Example

In this example, we compute a two-dimensional complex FFT by making one call to FFTCI followed by 2 N calls to F 2 TCF .

```
    USE FFTCI_INT
    USE CONST_INT
    USE F2TCF_INT
    USE UMACH_INT
    IMPLICIT NONE
    SPECIFICATIONS FOR PARAMETERS
    INTEGER N
    PARAMETER (N=4)
        INTEGER I, IR, IS, J, NOUT
        REAL FLOAT, TWOPI, WFFTC(35), CPY(2*N)
        COMPLEX CEXP, CMPLX, COEF(N,N), H, SEQ(N,N), TEMP
        INTRINSIC CEXP, CMPLX, FLOAT
        TWOPI = CONST('PI')
        TWOPI = 2*TWOPI
        IR = 3
        IS = 1
        Here we compute e**(2*pi*i/N)
        TEMP = CMPLX(0.0,TWOPI/FLOAT (N))
        H = CEXP(TEMP)
        DO 20 I=1, N
        DO 10 J=1, N
            SEQ(I,J) = H**((I-1)*(IR-1)+(J-1)*(IS-1))
    10 CONTInUE
    2 0 ~ C O N T I N U E ~
                                    Print out SEQ
                                    Get output unit number
```

!
CALL UMACH (2, NOUT)
WRITE (NOUT,99997)
DO $30 \mathrm{I}=1$, N
WRITE (NOUT,99998) (SEQ(I, J), J=1,N)
30 CONTINUE
$!$
!
CALL FFTCI (N, WFFTC)
DO 40 I=1, N
$\operatorname{CALL} \operatorname{F2TCF}(N, \operatorname{SEQ}(1:, I), \operatorname{COEF}(1:, I), \operatorname{WFFTC}, \operatorname{CPY})$
40
CONTINUE
DO 60 I=1, N
DO $50 \mathrm{~J}=\mathrm{I}+1, \mathrm{~N}$
TEMP $=\operatorname{COEF}(I, J)$
$\operatorname{COEF}(I, J)=\operatorname{COEF}(J, I)$
$\operatorname{COEF}(\mathrm{J}, \mathrm{I})=\operatorname{TEMP}$
50 CONTINUE
60 CONTINUE
!
Transform the columns of this result

```
DO 70 I=1, N
                CALL F2TCF (N, COEF(1:,I), SEQ(1:,I), WFFTC, CPY)
        7 0
    CONTINUE
!
DO 90 I=1, N
        DO 80 J=I + 1,N
                        TEMP = SEQ(I,J)
                        SEQ(I,J) = SEQ(J,I)
                SEQ(J,I) = TEMP
        80 CONTINUE
        90 CONTINUE
!
    WRITE (NOUT,99999)
        DO 100 I=1, N
        WRITE (NOUT, 99998) (SEQ (I, J) , J=1,N)
    100 CONTINUE
!
9 9 9 9 7 ~ F O R M A T ~ ( 1 X , ~ ' T h e ~ i n p u t ~ m a t r i x ~ i s ~ b e l o w ' )
99998 FORMAT (1X, 4(' (',F5.2,',',F5.2,')'))
99999 FORMAT (/, 1X, 'Result of two-dimensional transform')
    END
```


## Output

```
The input matrix is below
    (1.00, 0.00) ( 1.00, 0.00) ( 1.00, 0.00) ( 1.00, 0.00)
    (-1.00, 0.00) (-1.00, 0.00) (-1.00, 0.00) (-1.00, 0.00)
    (1.00, 0.00) ( 1.00, 0.00) ( 1.00, 0.00) ( 1.00, 0.00)
    (-1.00, 0.00) (-1.00, 0.00) (-1.00, 0.00) (-1.00, 0.00)
Result of two-dimensional transform
    ( 0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00)
    (0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00)
    (16.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00)
    (0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00) ( 0.00, 0.00)
```


## FSINT

Computes the discrete Fourier sine transformation of an odd sequence.

## Required Arguments

$\boldsymbol{N}$ - Length of the sequence to be transformed. It must be greater than 1. (Input)
$\mathbf{S E Q}$ - Array of length N containing the sequence to be transformed. (Input)
COEF - Array of length $N+1$ containing the transformed sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL FSINT (N, SEQ, COEF)

Specific: The specific interface names are S_FSINT and D_FSINT.

## FORTRAN 77 Interface

Single: CALL FSINT (N, SEQ, COEF)
Double: The double precision name is DFSINT.

## Description

The routine FSINT computes the discrete Fourier sine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N+1$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an $N$-vector $s=$ SEQ, FSINT returns in $c=$ COEF

$$
c_{m}=2 \sum_{n=1}^{N} s_{n} \sin \left(\frac{m n \pi}{N+1}\right)
$$

Finally, note that the Fourier sine transform is its own (unnormalized) inverse. The routine FSINT is based on the sine FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2INT/DF2INT. The reference is:

CALL F2INT (N, SEQ, COEF, WFSIN)
The additional argument is:
WFSIN - Array of length INT (2.5 * N + 15) initialized by FSINI. The initialization depends on N . (Input)
2. The routine FSINT is most efficient when $\mathrm{N}+1$ is the product of small primes.
3. The routine FSINT is its own (unnormalized) inverse. Applying FSINT twice will reproduce the original sequence multiplied by $2 *(\mathrm{~N}+1)$.
4. The arrays COEF and SEQ may be the same, if SEQ is also dimensioned at least $\mathrm{N}+1$.
5. $\operatorname{COEF}(\mathrm{N}+1)$ is needed as workspace.
6. If FSINT is used repeatedly with the same value of N , then call FSINI followed by repeated calls to F2INT. This is more efficient than repeated calls to FSINT.

## Example

In this example, we input a pure sine wave as a data vector and recover its Fourier sine series, which is a vector with all components zero except at the appropriate frequency it has an $N$.

```
USE FSINT INT
USE CONST-INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
INTEGER I, NOUT
REAL COEF(N+1), FLOAT, PI, SIN, SEQ(N)
INTRINSIC FLOAT, SIN
CALL UMACH (2, NOUT)
! Fill the data vector SEQ
PI = CONST('PI')
DO 10 I=1, N
        SEQ(I) = SIN(FLOAT(I)*PI/FLOAT (N+1))
    10 CONTINUE
CALL FSINT (N, SEQ, COEF)
WRITE (NOUT,99998)
WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99998 FORMAT (9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```

!
$!$

## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 0.38 | 8.00 |
| 2 | 0.71 | 0.00 |
| 3 | 0.92 | 0.00 |
| 4 | 1.00 | 0.00 |
| 5 | 0.92 | 0.00 |
| 6 | 0.71 | 0.00 |
| 7 | 0.38 | 0.00 |

## FSINI

Computes parameters needed by FSINT.

## Required Arguments

$N$ - Length of the sequence to be transformed. N must be greater than 1. (Input)

WFSIN - Array of length $\operatorname{INT}(2.5 * N+15)$ containing parameters needed by FSINT. (Output)

## FORTRAN 90 Interface

Generic: CALL FSINI (N, WFSIN)

Specific: The specific interface names are S_FSINI and D_FSINI.

## FORTRAN 77 Interface

Single: CALL FSINI (N, WFSIN)
Double: The double precision name is DFSINI.

## Description

The routine FSINI initializes the routine FSINT. An efficient way to make multiple calls for the same $N$ to IMSL routine FSINT, is to use routine FSINI for initialization. (In this case, replace FSINT with F2INT.) The routine FSINI is based on the routine SINTI in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different WFSIN arrays are needed for different values of N .

## Example

In this example, we compute three distinct sine FFTs by calling FSINI once and then calling F2INT three times.

```
USE FSINI_INT
USE UMACH_INT
USE CONST INT
USE F2INT_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
    INTEGER I, K, NOUT
    REAL COEF(N+1), FLOAT, PI, SIN, WFSIN(32), SEQ(N)
    INTRINSIC FLOAT, SIN
    CALL UMACH (2, NOUT)
    CALL FSINI (N, WFSIN)
    DO 20 K=1, 3
```

!

```
!
                                    with a pure sine wave
        PI = CONST('PI')
        DO 10 I=1, N
            SEQ(I) = SIN(FLOAT (K*I)*PI/FLOAT (N+1))
        CONTINUE
! Compute the transform of SEQ
    CALL F2INT (N, SEQ, COEF, WFSIN)
        WRITE (NOUT,99998)
        WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
    20 CONTINUE
99998 FORMAT (/, 9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
        END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 0.38 | 8.00 |
| 2 | 0.71 | 0.00 |
| 3 | 0.92 | 0.00 |
| 4 | 1.00 | 0.00 |
| 5 | 0.92 | 0.00 |
| 6 | 0.71 | 0.00 |
| 7 | 0.38 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 0.71 | 0.00 |
| 2 | 1.00 | 8.00 |
| 3 | 0.71 | 0.00 |
| 4 | 0.00 | 0.00 |
| 5 | -0.71 | 0.00 |
| 6 | -1.00 | 0.00 |
| 7 | -0.71 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 0.92 | 0.00 |
| 2 | 0.71 | 0.00 |
| 3 | -0.38 | 8.00 |
| 4 | -1.00 | 0.00 |
| 5 | -0.38 | 0.00 |
| 6 | 0.71 | 0.00 |
| 7 | 0.92 | 0.00 |

## FCOST

Computes the discrete Fourier cosine transformation of an even sequence.

## Required Arguments

$N$ - Length of the sequence to be transformed. It must be greater than 1 . (Input)
SEQ - Array of length N containing the sequence to be transformed. (Input)

COEF - Array of length N containing the transformed sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL FCOST (N, SEQ, COEF)

Specific: The specific interface names are S_FCOST and D_FCOST.

## FORTRAN 77 Interface

Single: CALL FCOST (N, SEQ, COEF)
Double: The double precision name is DFCOST.

## Description

The routine FCOST computes the discrete Fourier cosine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N-1$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an $N$-vector $s=\mathrm{SEQ}, \mathrm{FCOST}$ returns in $c=\mathrm{COEF}$

$$
c_{m}=2 \sum_{n=2}^{N-1} s_{n} \cos \left[\frac{(m-1)(n-1) \pi}{N-1}\right]+s_{1}+s_{N}(-1)^{(m-1)}
$$

Finally, note that the Fourier cosine transform is its own (unnormalized) inverse. Two applications of FCOST to a vector $s$ produces $(2 N-2) s$. The routine FCOST is based on the cosine FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2OST/DF2OST. The reference is:

CALL F2OST (N, SEQ, COEF, WFCOS)
The additional argument is
WFCOS - Array of length $3 * \mathrm{~N}+15$ initialized by FCOSI. The initialization depends on N . (Input)
2. The routine $\operatorname{FCOST}$ is most efficient when $N-1$ is the product of small primes.
3. The routine FCOST is its own (unnormalized) inverse. Applying FCOST twice will reproduce the original sequence multiplied by $2 *(\mathrm{~N}-1)$.
4. The arrays COEF and SEQ may be the same.
5. If FCOST is used repeatedly with the same value of N , then call $\operatorname{FCOSI}$ followed by repeated calls to F2OST. This is more efficient than repeated calls to FCOST.

## Example

In this example, we input a pure cosine wave as a data vector and recover its Fourier cosine series, which is a vector with all components zero except at the appropriate frequency it has an $N-1$.

```
USE FCOST INT
USE CONST INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
INTEGER I, NOUT
REAL COEF(N), COS, FLOAT, PI, SEQ(N)
INTRINSIC COS, FLOAT
CALL UMACH (2, NOUT)
! Fill the data vector SEQ
PI = CONST('PI')
DO 10 I=1, N
        SEQ(I) = COS(FLOAT (I-1)*PI/FLOAT (N-1))
    10 CONTINUE
CALL FCOST (N, SEQ, COEF)
WRITE (NOUT,99998)
WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99998 FORMAT (9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```

!
!
!

## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 0.00 |
| 2 | 0.87 | 6.00 |
| 3 | 0.50 | 0.00 |
| 4 | 0.00 | 0.00 |
| 5 | -0.50 | 0.00 |
| 6 | -0.87 | 0.00 |
| 7 | -1.00 | 0.00 |

## FCOSI

Computes parameters needed by FCOST.

## Required Arguments

$N$ - Length of the sequence to be transformed. N must be greater than 1. (Input)
WFCOS - Array of length $3 \mathrm{~N}+15$ containing parameters needed by FCOST. (Output)

## FORTRAN 90 Interface

Generic: CALL FCOSI (N, WFCOS)

Specific: The specific interface names are S_FCOSI and D_FCOSI.

## FORTRAN 77 Interface

Single: CALL FCOSI (N, WFCOS)

Double: The double precision name is DFCOSI.

## Description

The routine FCOSI initializes the routine FCOST. An efficient way to make multiple calls for the same $N$ to IMSL routine $\operatorname{FCOST}$ is to use routine FCOSI for initialization. (In this case, replace FCOST with F2OST.) The routine FCOSI is based on the routine COSTI in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different WFCOS arrays are needed for different values of N .

## Example

In this example, we compute three distinct cosine FFTs by calling FCOSI once and then calling F2OST three times.

```
USE FCOSI INT
USE CONST_INT
USE F2OST INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
INTEGER I, K, NOUT
REAL COEF(N), COS, FLOAT, PI, WFCOS(36), SEQ(N)
INTRINSIC COS, FLOAT
CALL UMACH (2, NOUT)
CALL FCOSI (N, WFCOS)
```

!

```
! wave will be transformed
    PI = CONST('PI')
    DO 20 K=1, 3
! Fill the data vector SEQ
        DO 10 I=1, N
                SEQ(I) = COS (FLOAT (K* (I-1))*PI/FLOAT (N-1))
    10 CONTINUE
! Compute the transform of SEQ
        CALL F2OST (N, SEQ, COEF, WFCOS)
        WRITE (NOUT, 99998)
        WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
    20 CONTINUE
99998 FORMAT (/, 9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
        END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 0.00 |
| 2 | 0.87 | 6.00 |
| 3 | 0.50 | 0.00 |
| 4 | 0.00 | 0.00 |
| 5 | -0.50 | 0.00 |
| 6 | -0.87 | 0.00 |
| 7 | -1.00 | 0.00 |


| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 0.00 |
| 2 | 0.50 | 0.00 |
| 3 | -0.50 | 6.00 |
| 4 | -1.00 | 0.00 |
| 5 | -0.50 | 0.00 |
| 6 | 0.50 | 0.00 |
| 7 | 1.00 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 1.00 | 0.00 |
| 2 | 0.00 | 0.00 |
| 3 | -1.00 | 0.00 |
| 4 | 0.00 | 6.00 |
| 5 | 1.00 | 0.00 |
| 6 | 0.00 | 0.00 |
| 7 | -1.00 | 0.00 |

## QSINF

Computes the coefficients of the sine Fourier transform with only odd wave numbers.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
SEQ - Array of length N containing the sequence. (Input)
COEF - Array of length N containing the Fourier coefficients. (Output)

## FORTRAN 90 Interface

Generic: CALL QSINF (N, SEQ, COEF)
Specific: The specific interface names are S_QSINF and D_QSINF.

## FORTRAN 77 Interface

Single: CALL QSINF (N, SEQ, COEF)
Double: The double precision name is DQSINF.

## Description

The routine QSINF computes the discrete Fourier quarter sine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an N -vector $s=$ SEQ, QSINF returns in $c=$ COEF

$$
c_{m}=2 \sum_{n=1}^{N-1} s_{n} \sin \left[\frac{(2 m-1) n \pi}{2 N}\right]+s_{N}(-1)^{m-1}
$$

Finally, note that the Fourier quarter sine transform has an (unnormalized) inverse, which is implemented in the IMSL routine QSINB. The routine QSINF is based on the quarter sine FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of Q2INF/DQ2 INF. The reference is:

CALL Q2INF (N, SEQ, COEF, WQSIN)
The additional argument is:
WQSIN - Array of length $3 * \mathrm{~N}+15$ initialized by QSINI. The initialization depends on $N$. (Input)
2. The routine QSINF is most efficient when $N$ is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If QSINF/QSINB is used repeatedly with the same value of $N$, then call QSINI followed by repeated calls to $Q 2 I N F / Q 2 I N B$. This is more efficient than repeated calls to QSINF/QSINB.

## Example

In this example, we input a pure quarter sine wave as a data vector and recover its Fourier quarter sine series.

```
USE QSINF_INT
USE CONST }\mp@subsup{}{}{-}\mathrm{ INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
!
INTEGER I, NOUT
REAL COEF(N), FLOAT, PI, SIN, SEQ(N)
INTRINSIC FLOAT, SIN
CALL UMACH (2, NOUT)
! Fill the data vector SEQ
! with a pure sine wave
PI = CONST('PI')
DO 10 I=1, N
        SEQ(I) = SIN(FLOAT (I)*(PI/2.0)/FLOAT (N))
    10 CONTINUE
! Compute the transform of SEQ
CALL QSINF (N, SEQ, COEF)
WRITE (NOUT,99998)
WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99998 FORMAT (9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 0.22 | 7.00 |
| 2 | 0.43 | 0.00 |
| 3 | 0.62 | 0.00 |
| 4 | 0.78 | 0.00 |
| 5 | 0.90 | 0.00 |
| 6 | 0.97 | 0.00 |
| 7 | 1.00 | 0.00 |

## QSINB

Computes a sequence from its sine Fourier coefficients with only odd wave numbers.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
COEF - Array of length N containing the Fourier coefficients. (Input)
$\boldsymbol{S E Q}$ - Array of length N containing the sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL QSINB ( $\mathrm{N}, \mathrm{COEF}, \mathrm{SEQ}$ )
Specific: The specific interface names are S_QSINB and D_QSINB.

## FORTRAN 77 Interface

Single: CALL QSINB (N, COEF, SEQ)
Double: The double precision name is DQS INB.

## Description

The routine QSINB computes the discrete (unnormalized) inverse Fourier quarter sine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.
Specifically, given an $N$-vector $c=$ COEF, QSINB returns in $s=$ SEQ

$$
s_{m}=4 \sum_{n=1}^{N} c_{n} \sin \left(\frac{(2 n-1) m \pi}{2 N}\right)
$$

Furthermore, a vector $x$ of length $N$ that is first transformed by QSINF and then by QSINB will be returned by QSINB as $4 N X$. The routine QSINB is based on the inverse quarter sine FFT in FFTPACK which was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of Q2 INB / DQ2 INB. The reference is:

CALL Q2INB (N, SEQ, COEF, WQSIN)
The additional argument is:

WQSIN — ray of length $3 * \mathrm{~N}+15$ initialized by QSINI. The initialization depends on N.(Input)
2. The routine QSINB is most efficient when $N$ is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If QSINF/QSINB is used repeatedly with the same value of $N$, then call QSINI followed by repeated calls to Q2INF/Q2INB. This is more efficient than repeated calls to QSINF/QSINB.

## Example

In this example, we first compute the quarter wave sine Fourier transform $c$ of the vector $x$ where $x_{n}=n$ for $n=1$ to $N$. We then compute the inverse quarter wave Fourier transform of $c$ which is $4 N x=s$.

USE QSINB_INT
USE QSINF_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER ( $\mathrm{N}=7$ )
!
INTEGER I, NOUT
REAL FLOAT, SEQ (N), COEF (N), X(N)
INTRINSIC FLOAT
CALL UMACH (2, NOUT)
$!$ Fill the data vector X
!
DO $10 \quad \mathrm{I}=1$, N
X(I) = FLOAT (I)
10 CONTINUE
$!$
CALL QSINF (N, X, COEF)
CALL QSINB (N, COEF, SEQ)
WRITE (NOUT,99998)
WRITE (NOUT,99999) (X(I), COEF(I), SEQ(I), I=1,N)
99998 FORMAT (5X, 'INPUT', 5X, 'FORWARD TRANSFORM', 3X, 'BACKWARD ', \&
'TRANSFORM')
99999 FORMAT (3X, F6.2, 10X, F6.2, 15X, F6.2)
END

## Output

INPUT FORWARD TRANSFORM BACKWARD TRANSFORM
$1.00 \quad 39.88 \quad 28.00$
$2.00-4.58 \quad 56.00$

| 3.00 | 1.77 | 84.00 |
| ---: | ---: | ---: |
| 4.00 | -1.00 | 112.00 |
| 5.00 | 0.70 | 140.00 |
| 6.00 | -0.56 | 168.00 |
| 7.00 | 0.51 | 196.00 |

## QSINI

Computes parameters needed by QSINF and QSINB.
CALL QSINI (N, WQSIN)

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
WQSIN - Array of length $3 \mathrm{~N}+15$ containing parameters needed by QSINF and QSINB. (Output)

## FORTRAN 90 Interface

Generic: CALL QSINI (N,WQSIN)

Specific: The specific interface names are S_QSINI and D_QSINI.

## FORTRAN 77 Interface

Single: CALL QSINI (N, WQSIN)
Double: The double precision name is DQSINI.

## Description

The routine QSINI initializes the routines QSINF and QSINB. An efficient way to make multiple calls for the same $N$ to IMSL routine QSINF or QSINB is to use routine QSINI for initialization. (In this case, replace QSINF or QSINB with Q2INF or Q2INB, respectively.) The routine QSINI is based on the routine SINQI in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different WQSIN arrays are needed for different values of N .

## Example

In this example, we compute three distinct quarter sine transforms by calling QSINI once and then calling Q2INF three times.

```
USE QSINI INT
USE CONST_INT
```

```
    USE Q2INF_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=7)
!
    INTEGER I, K, NOUT
    REAL COEF(N), FLOAT, PI, SIN, WQSIN(36), SEQ(N)
    INTRINSIC FLOAT, SIN
    CALL UMACH (2, NOUT)
    CALL QSINI (N, WQSIN)
    PI = CONST('PI')
    DO 20 K=1, 3
        Fill the data vector SEQ
                                with a pure sine wave
        DO 10 I=1, N
                SEQ(I) = SIN(FLOAT((2*K-1)*I)*(PI/2.0)/FLOAT (N))
    CONTINUE
                                    Compute the transform of SEQ
    CALL Q2INF (N, SEQ, COEF, WQSIN)
                                    Print results
    WRITE (NOUT,99998)
    WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
    20 CONTINUE
99998 FORMAT (/, 9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```


## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 0.22 | 7.00 |
| 2 | 0.43 | 0.00 |
| 3 | 0.62 | 0.00 |
| 4 | 0.78 | 0.00 |
| 5 | 0.90 | 0.00 |
| 6 | 0.97 | 0.00 |
| 7 | 1.00 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 0.62 | 0.00 |
| 2 | 0.97 | 7.00 |
| 3 | 0.90 | 0.00 |
| 4 | 0.43 | 0.00 |
| 5 | -0.22 | 0.00 |
| 6 | -0.78 | 0.00 |
| 7 | -1.00 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 0.90 | 0.00 |


| 2 | 0.78 | 0.00 |
| ---: | ---: | ---: |
| 3 | -0.22 | 7.00 |
| 4 | -0.97 | 0.00 |
| 5 | -0.62 | 0.00 |
| 6 | 0.43 | 0.00 |
| 7 | 1.00 | 0.00 |

## QCOSF

Computes the coefficients of the cosine Fourier transform with only odd wave numbers.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
$\mathbf{S E Q}$ - Array of length N containing the sequence. (Input)
COEF - Array of length $n$ containing the Fourier coefficients. (Output)

## FORTRAN 90 Interface

Generic: CALL QCOSF (N, SEQ, COEF [, ...])
Specific: The specific interface names are S_QCOSF and D_QCOSF.

## FORTRAN 77 Interface

Single: CALL QCOSF (N, SEQ, COEF)
Double: The double precision name is DQCOSF.

## Description

The routine QCOSF computes the discrete Fourier quarter cosine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$.

Specifically, given an $N$-vector $s=\mathrm{SEQ}$, QCOSF returns in $c=\mathrm{COEF}$

$$
c_{m}=s_{1}+2 \sum_{n=2}^{N} s_{n} \cos \left(\frac{(2 m-1)(n-1) \pi}{2 N}\right)
$$

Finally, note that the Fourier quarter cosine transform has an (unnormalized) inverse which is implemented in QCOSB. The routine QCOSF is based on the quarter cosine FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of Q2OSF/DQ2OSF. The reference is:

CALL Q2OSF (N, SEQ, COEF, WQCOS)
The additional argument is:
WQCOS - Array of length $3 * \mathrm{~N}+15$ initialized by QCOSI. The initialization depends on N . (Input)
2. The routine QCOSF is most efficient when N is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If $Q C O S F / Q C O S B$ is used repeatedly with the same value of $N$, then call QCOSI followed by repeated calls to $Q 20 S F / Q 20 S B$. This is more efficient than repeated calls to QCOSF/QCOSB.

## Example

In this example, we input a pure quarter cosine wave as a data vector and recover its Fourier quarter cosine series.

```
USE QCOSF_INT
USE CONST INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
INTEGER I, NOUT
REAL COEF(N), COS, FLOAT, PI, SEQ(N)
INTRINSIC COS, FLOAT
CALL UMACH (2, NOUT)
PI = CONST('PI')
DO 10 I=1, N
        SEQ(I) = COS(FLOAT(I-1)*(PI/2.0)/FLOAT(N))
    CONTINUE
        Call QCOSF (N, SEQ, COEF)
        WRITE (NOUT,99998)
        WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
99998 FORMAT (9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
99999 FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```

$!$
10
$!$

## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 7.00 |
| 2 | 0.97 | 0.00 |
| 3 | 0.90 | 0.00 |
| 4 | 0.78 | 0.00 |
| 5 | 0.62 | 0.00 |
| 6 | 0.43 | 0.00 |
| 7 | 0.22 | 0.00 |

## QCOSB

Computes a sequence from its cosine Fourier coefficients with only odd wave numbers.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
COEF - Array of length n containing the Fourier coefficients. (Input)
SEQ - Array of length n containing the sequence. (Output)

## FORTRAN 90 Interface

Generic: CALL QCOSB (N, COEF, SEQ)

Specific: The specific interface names are S_QCOSB and D_QCOSB.

## FORTRAN 77 Interface

Single: CALL QCOSB (N, COEF, SEQ)
Double: The double precision name is DQCOSB.

## Description

The routine $Q \operatorname{COSB}$ computes the discrete (unnormalized) inverse Fourier quarter cosine transform of a real vector of size $N$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ is a product of small prime factors. If $N$ satisfies this condition, then the computational effort is proportional to $N \log N$. Specifically, given an $N$-vector $c=$ COEF, QCOSB returns in $s=\mathrm{SEQ}$

$$
s_{m}=4 \sum_{n=1}^{N} c_{n} \cos \left(\frac{(2 n-1)(m-1) \pi}{2 N}\right)
$$

Furthermore, a vector $x$ of length $N$ that is first transformed by QCOSF and then by QCOSB will be returned by QCOSB as $4 N x$. The routine QCOSB is based on the inverse quarter cosine FFT in

FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of Q2OSB/DQ2OSB. The reference is:

CALL Q2OSB (N, COEF, SEQ, WQCOS)
The additional argument is:
WQCOS - Array of length $3 * N+15$ initialized by QCOSI . The initialization depends on N. (Input)
2. The routine $\ell \operatorname{CosB}$ is most efficient when $N$ is the product of small primes.
3. The arrays COEF and SEQ may be the same.
4. If $Q \operatorname{COSF} / Q \operatorname{COSB}$ is used repeatedly with the same value of N , then call QCOSI followed by repeated calls to $Q 20 S F / Q 2 O S B$. This is more efficient than repeated calls to QCOSF/QCOSB.

## Example

In this example, we first compute the quarter wave cosine Fourier transform $c$ of the vector $x$, where $x_{n}=n$ for $n=1$ to $N$. We then compute the inverse quarter wave Fourier transform of $c$ which is $4 N x=s$.

```
USE QCOSB_INT
USE QCOSF_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
INTEGER I, NOUT
REAL FLOAT, SEQ(N), COEF(N), X(N)
INTRINSIC FLOAT
! CALL UMACH (2, NOUT)
DO 10 I=1, N
        X(I) = FLOAT(I)
    1 0 ~ C O N T I N U E
CALL QCOSF (N, X, COEF)
CALL QCOSB (N, COEF, SEQ)
```

!
$!$
!

```
! Print results
    WRITE (NOUT,99998)
    DO 20 I=1, N
        WRITE (NOUT,99999) X(I), COEF(I), SEQ(I)
    20 CONTINUE
99998 FORMAT (5X, 'INPUT', 5X, 'FORWARD TRANSFORM', 3X, 'BACKWARD ', &
            'TRANSFORM')
99999 FORMAT (3X, F6.2, 10X, F6.2, 15X, F6.2)
    END
```


## Output

| INPUT | FORWARD | TRANSFORM |
| :--- | :---: | :---: |
| 1.00 | 31.12 | BACKWARD |
| 2.00 | -27.45 | 28.00 |
| 3.00 | 10.97 | 56.00 |
| 4.00 | -9.00 | 84.00 |
| 5.00 | 4.33 | 112.00 |
| 6.00 | -3.36 | 140.00 |
| 7.00 | 0.40 | 168.00 |
|  |  | 196.00 |

## QCOSI

Computes parameters needed by QCOSF and QCOSB.

## Required Arguments

$N$ - Length of the sequence to be transformed. (Input)
WQCOS - Array of length $3 \mathrm{~N}+15$ containing parameters needed by QCOSF and QCOSB. (Output)

## FORTRAN 90 Interface

Generic: CALL QCOSI (N,WQCOS)

Specific: The specific interface names are S_QCOSI and D_QCOSI.

## FORTRAN 77 Interface

Single: CALL QCOSI (N, WQCOS)
Double: The double precision name is DQCOSI.

## Description

The routine QCOSI initializes the routines $Q C O S F$ and $Q C O S B$. An efficient way to make multiple calls for the same $N$ to IMSL routine $\varrho C O S F$ or $\varrho C O S B$ is to use routine QCOSI for initialization. (In this case, replace $Q C O S F$ or $Q C O S B$ with $Q 2 O S F$ or $Q 2 O S B$, respectively.) The routine $2 C O S I$ is
based on the routine COSQI in FFTPACK, which was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

Different $W Q C O S$ arrays are needed for different values of N .

## Example

In this example, we compute three distinct quarter cosine transforms by calling QCOSI once and then calling Q2OSF three times.

```
USE QCOSI INT
USE CONST INT
USE Q2OSF INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=7)
    INTEGER I, K, NOUT
    REAL COEF(N), COS, FLOAT, PI, WQCOS(36), SEQ(N)
    INTRINSIC COS, FLOAT
    CALL UMACH (2, NOUT)
    CALL QCOSI (N, WQCOS)
    PI = CONST('PI')
    DO 20 K=1, 3
! Fill the data vector SEQ
        DO 10 I=1, N
            SEQ(I) = COS(FLOAT((2*K-1)*(I-1))*(PI/2.0)/FLOAT (N))
        CONTINUE
        CALL Q2OSF (N, SEQ, COEF, WQCOS)
        Print results
        WRITE (NOUT,99998)
        WRITE (NOUT,99999) (I, SEQ(I), COEF(I), I=1,N)
    CONTINUE
FORMAT (/, 9X, 'INDEX', 6X, 'SEQ', 7X, 'COEF')
FORMAT (1X, I11, 5X, F6.2, 5X, F6.2)
END
```

!

## Output

| INDEX | SEQ | COEF |
| :---: | ---: | ---: |
| 1 | 1.00 | 7.00 |
| 2 | 0.97 | 0.00 |
| 3 | 0.90 | 0.00 |
| 4 | 0.78 | 0.00 |


| 5 | 0.62 | 0.00 |
| :---: | ---: | ---: |
| 6 | 0.43 | 0.00 |
| 7 | 0.22 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 1.00 | 0.00 |
| 2 | 0.78 | 7.00 |
| 3 | 0.22 | 0.00 |
| 4 | -0.43 | 0.00 |
| 5 | -0.90 | 0.00 |
| 6 | -0.97 | 0.00 |
| 7 | -0.62 | 0.00 |
|  |  |  |
| INDEX | SEQ | COEF |
| 1 | 1.00 | 0.00 |
| 2 | 0.43 | 0.00 |
| 3 | -0.62 | 7.00 |
| 4 | -0.97 | 0.00 |
| 5 | -0.22 | 0.00 |
| 6 | 0.78 | 0.00 |
| 7 | 0.90 | 0.00 |

## FFT2D

Computes Fourier coefficients of a complex periodic two-dimensional array.

## Required Arguments

$\boldsymbol{A}$ — NRA by NCA complex matrix containing the periodic data to be transformed. (Input)
COEF - NRA by NCA complex matrix containing the Fourier coefficients of A. (Output)

## Optional Arguments

NRA - The number of rows of A. (Input)
Default: NRA $=\operatorname{size}(A, 1)$.
$N C A$ - The number of columns of A. (Input)
Default: NCA $=\operatorname{size}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
LDCOEF - Leading dimension of COEF exactly as specified in the dimension statement of the calling program. (Input)
Default: LDCOEF $=\operatorname{size}(C O E F, 1)$.

## FORTRAN 90 Interface

Generic: CALL FFT2D (A, COEF [,...])

Specific: The specific interface names are S_FFT2D and D_FFT2D.

## FORTRAN 77 Interface

Single: CALL FFT2D (NRA, NCA, A, LDA, COEF, LDCOEF)

Double: $\quad$ The double precision name is DFFT2D.

## Description

The routine FFT2D computes the discrete complex Fourier transform of a complex two dimensional array of size $(\mathrm{NRA}=N) \times(\mathrm{NCA}=M)$. The method used is a variant of the CooleyTukey algorithm, which is most efficient when $N$ and $M$ are each products of small prime factors. If $N$ and $M$ satisfy this condition, then the computational effort is proportional to $N M \log N M$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N \times M$ array $a$, FFT2D returns in $c=\mathrm{COEF}$

$$
c_{j k}=\sum_{n=1}^{N} \sum_{m=1}^{M} a_{n m} e^{-2 \pi i(j-1)(n-1) / N} e^{-2 \pi i(k-1)(m-1) / M}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
\sqrt{N M} S
$$

Finally, note that an unnormalized inverse is implemented in FFT2B. The routine FFT2D is based on the complex FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2T2D/DF2T2D. The reference is:

CALL F2T2D (NRA, NCA, A, LDA, COEF, LDCOEF, WFF1, WFF2, CWK, CPY)

The additional arguments are as follows:
WFF1 - Real array of length 4 *NRA +15 initialized by FFTCI. The initialization depends on NRA. (Input)

WFF2 - Real array of length 4 * NCA +15 initialized by FFTCI. The initialization depends on NCA. (Input)

CWK — Complex array of length 1. (Workspace)
$\boldsymbol{C P Y}$ - Real array of length 2 * MAX (NRA, NCA). (Workspace)
2. The routine FFT2D is most efficient when NRA and NCA are the product of small primes.
3. The arrays COEF and A may be the same.
4. If FFT2D/FFT2B is used repeatedly, with the same values for NRA and NCA, then use $\operatorname{FFTCI}$ to fill $\operatorname{WFF} 1(\mathrm{~N}=\mathrm{NRA})$ and $\operatorname{WFF} 2(\mathrm{~N}=\mathrm{NCA})$. Follow this with repeated calls to F2T2D/F2T2B. This is more efficient than repeated calls to FFT2D/FFT2B.

## Example

In this example, we compute the Fourier transform of the pure frequency input for a $5 \times 4$ array

$$
a_{n m}=e^{2 \pi i(n-1) 2 / N} e^{2 \pi i(m-1) 3 / M}
$$

for $1 \leq n \leq 5$ and $1 \leq m \leq 4$ using the IMSL routine FFT2D. The result

$$
\hat{a}=c
$$

has all zeros except in the $(3,4)$ position.

```
USE FFT2D_INT
USE CONST INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER I, IR, IS, J, NCA, NRA
REAL FLOAT, TWOPI
COMPLEX A(5,4), C, CEXP, CMPLX, COEF (5,4), H
CHARACTER TITLE1*26, TITLE2*26
INTRINSIC CEXP, CMPLX, FLOAT
TITLE1 = 'The input matrix is below '
TITLE2 = 'The output matrix is below'
NRA = 5
NCA = 4
IR = 3
IS = 4
! Fill A with initial data
TWOPI = CONST('PI')
TWOPI = 2.0*TWOPI
C = CMPLX(0.0,1.0)
H = CEXP(TWOPI*C)
DO 10 I=1, NRA
        DO 10 J=1, NCA
        A(I,J) = CEXP(TWOPI*C*((FLOAT((I-1)*(IR-1))/FLOAT(NRA) + &
                            FLOAT((J-1)*(IS-1))/FLOAT (NCA))))
    1 0 ~ C O N T I N U E ~
CALL WRCRN (TITLE1, A)
```

!

```
!
    CALL FFT2D (A, COEF)
    CALL WRCRN (TITLE2, COEF)
!
END
```


## Output



## FFT2B

Computes the inverse Fourier transform of a complex periodic two-dimensional array.

## Required Arguments

COEF - NRCOEF by NCCOEF complex array containing the Fourier coefficients to be transformed. (Input)

A - nRCOEF by NCCOEf complex array containing the Inverse Fourier coefficients of coef. (Output)

## Optional Arguments

NRCOEF - The number of rows of COEF. (Input)
Default: NRCOEF $=\operatorname{size}($ COEF, 1$)$.

NCCOEF - The number of columns of COEF. (Input)
Default: $\mathrm{NCCOEF}=\operatorname{size}(\mathrm{COEF}, 2)$.
LDCOEF - Leading dimension of COEF exactly as specified in the dimension statement of the calling program. (Input)
Default: LDCOEF = size (COEF, 1).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL FFT2B (COEF,A [, ...])

Specific: The specific interface names are S_FFT2B and D_FFT2B.

## FORTRAN 77 Interface

Single: CALL FFT2B (NRCOEF, NCCOEF, COEF, LDCOEF, A, LDA)
Double: The double precision name is DFFT2B.

## Description

The routine FFT2B computes the inverse discrete complex Fourier transform of a complex twodimensional array of size $($ NRCOEF $=N) \times(\mathrm{NCCOEF}=M)$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N$ and $M$ are both products of small prime factors. If $N$ and $M$ satisfy this condition, then the computational effort is proportional to $N M \log$ $N M$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N \times M$ array $c=\operatorname{COEF}$, FFT2B returns in $a$

$$
a_{j k}=\sum_{n=1}^{N} \sum_{m=1}^{M} c_{n m} e^{2 \pi i(j-1)(n-1) / N} e^{2 \pi i(k-1)(m-1) / M}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
S \sqrt{N M}
$$

Finally, note that an unnormalized inverse is implemented in FFT2D. The routine FFT2B is based on the complex FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2T2B/DF2T2B. The reference is:

CALL F2T2B (NRCOEF, NCCOEF, A, LDA, COEF, LDCOEF, WFF1, WFF2, CWK, CPY)
The additional arguments are as follows:
WFF1 - Real array of length 4 * NRCOEF +15 initialized by FFTCI. The initialization depends on NRCOEF. (Input)

WFF2 - Real array of length 4 * NCCOEF +15 initialized by FFTCI. The initialization depends on NCCOEF. (Input)

CWK - Complex array of length 1. (Workspace)
$\boldsymbol{C P Y}$ - Real array of length 2 * MAX(NRCOEF, NCCOEF). (Workspace)
2. The routine FFT2B is most efficient when NRCOEF and NCCOEF are the product of small primes.
3. The arrays COEF and A may be the same.
4. If $\mathrm{FFT} 2 \mathrm{D} / \mathrm{FFT} 2 \mathrm{~B}$ is used repeatedly, with the same values for NRCOEF and NCCOEF, then use FFTCI to fill WFF1 $(\mathrm{N}=\mathrm{NRCOEF})$ and $\operatorname{WFF} 2(\mathrm{~N}=\mathrm{NCCOEF})$. Follow this with repeated calls to $\mathrm{F} 2 \mathrm{~T} 2 \mathrm{D} / \mathrm{F} 2 \mathrm{~T} 2 \mathrm{~B}$. This is more efficient than repeated calls to FFT2D/FFT2B.

## Example

In this example, we first compute the Fourier transform of the $5 \times 4$ array

$$
x_{n m}=n+5(m-1)
$$

for $1 \leq n \leq 5$ and $1 \leq m \leq 4$ using the IMSL routine FFT2D. The result

$$
\hat{x}=c
$$

is then inverted by a call to FFT2B. Note that the result is an array a satisfying $a=(5)(4) x=20 x$. In general, FFT2B is an unnormalized inverse with expansion factor $N M$.

```
USE FFT2B INT
USE FFT2D_INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER M, N, NCA, NRA
COMPLEX CMPLX, X(5,4), A(5,4), COEF(5,4)
CHARACTER TITLE1*26, TITLE2*26, TITLE3*26
INTRINSIC CMPLX
    TITLE1 = 'The input matrix is below '
    TITLE2 = 'After FFT2D
    TITLE3 = 'After FFT2B '
    NRA = 5
    NCA = 4
! Fill X with initial data
    DO 20 N=1, NRA
        DO 10 M=1, NCA
        X(N,M) = CMPLX(FLOAT (N+5*M-5),0.0)
        CONTINUE
    CONTINUE
```

!
!

```
    CALL WRCRN (TITLE1, X)
    CALL FFT2D (X, COEF)
    CALL WRCRN (TITLE2, COEF)
    CALL FFT2B (COEF, A)
    CALL WRCRN (TITLE3, A)
    END
```


## Output

|  |  |  | input ma | $\begin{gathered} \text { rix is } \\ 2 \end{gathered}$ | below | 3 |  | 4 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1.00, | $0.00)$ | ( 6.00, | $0.00)$ | ( 11.00, | 0.00) | ( 16.00, | $0.00)$ |
| 2 | ( 2.00, | $0.00)$ | ( 7.00, | $0.00)$ | ( 12.00, | 0.00) | ( 17.00, | $0.00)$ |
| 3 | ( 3.00, | $0.00)$ | ( 8.00, | $0.00)$ | ( 13.00, | 0.00) | ( 18.00, | $0.00)$ |
| 4 | ( 4.00, | $0.00)$ | ( 9.00, | $0.00)$ | ( 14.00, | 0.00) | ( 19.00, | $0.00)$ |
| 5 | 5.00, | $0.00)$ | ( 10.00, | $0.00)$ | ( 15.00, | 0.00) | ( 20.00, | $0.00)$ |
| After FFT2D |  |  |  |  |  |  |  |  |
|  |  | 1 |  | 2 |  | 3 |  | 4 |
| 1 | ( 210.0, | 0.0) | ( -50.0, | $50.0)$ | ( -50.0, | 0.0) | ( -50.0, | -50.0) |
| 2 | ( -10.0, | 13.8) | ( 0.0, | 0.0) | ( 0.0, | 0.0) | ( 0.0, | 0.0) |
| 3 | ( -10.0, | 3.2) | 0.0, | 0.0) | 0.0, | 0.0) | ( 0.0, | 0.0) |
| 4 | ( -10.0, | -3.2) | 0.0, | 0.0) | ( 0.0, | 0.0) | ( 0.0, | 0.0) |
| 5 | ( -10.0, | -13.8) | ( 0.0, | 0.0) | 0.0, | $0.0)$ | 0.0, | $0.0)$ |
| After FFT2B |  |  |  |  |  |  |  |  |
|  |  | 1 |  | 2 |  | 3 |  | 4 |
| 1 | ( 20.0, | 0.0) | ( 120.0, | 0.0) | ( 220.0, | 0.0) | ( 320.0, | 0.0) |
| 2 | 40.0, | 0.0) | ( 140.0, | 0.0) | ( 240.0, | 0.0) | ( 340.0, | 0.0) |
| 3 | 60.0, | 0.0) | ( 160.0, | 0.0) | ( 260.0, | 0.0) | ( 360.0, | 0.0) |
| 4 | ( 80.0, | 0.0) | ( 180.0, | 0.0) | ( 280.0, | 0.0) | ( 380.0, | 0.0) |
| 5 | ( 100.0, | $0.0)$ | ( 200.0, | 0.0) | ( 300.0, | 0.0) | ( 400.0, | 0.0) |

## FFT3F

Computes Fourier coefficients of a complex periodic three-dimensional array.

## Required Arguments

A - Three-dimensional complex matrix containing the data to be transformed. (Input)
$\boldsymbol{B}$ - Three-dimensional complex matrix containing the Fourier coefficients of A. (Output) The matrices A and B may be the same.

## Optional Arguments

N1 - Limit on the first subscript of matrices A and B. (Input)
Default: N1 $=\operatorname{size}(\mathrm{A}, 1)$
N2 - Limit on the second subscript of matrices A and B. (Input)
Default: $\mathrm{N} 2=\operatorname{size}(\mathrm{A}, 2)$

N3 - Limit on the third subscript of matrices A and B. (Input)
Default: N3 $=\operatorname{size}(\mathrm{A}, 3)$
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.
MDA - Middle dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: MDA $=\operatorname{size}(\mathrm{A}, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{size}(\mathrm{B}, 1)$.
MDB - Middle dimension of B exactly as specified in the dimension statement of the calling program. (Input)
Default: $\mathrm{MDB}=\operatorname{size}(\mathrm{B}, 2)$.

## FORTRAN 90 Interface

Generic: CALL FFT3F (A, B $[, \ldots]$ )
Specific: The specific interface names are S_FFT3F and D_FFT3F.

## FORTRAN 77 Interface

Single: CALL FFT3F (N1, N2, N3, A, LDA, MDA, B, LDB, MDB)
Double: $\quad$ The double precision name is DFFT3F.

## Description

The routine FFT3F computes the forward discrete complex Fourier transform of a complex threedimensional array of size $(\mathrm{N} 1=N) \times(\mathrm{N} 2=M) \times(\mathrm{N} 3=L)$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N, M$, and $L$ are each products of small prime factors. If $N, M$, and $L$ satisfy this condition, then the computational effort is proportional to $N M L \log N M L$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N \times M \times L$ array $a$, FFT3F returns in $c=\mathrm{COEF}$

$$
c_{j k l}=\sum_{n=1}^{N} \sum_{m=1}^{M} \sum_{l=1}^{L} a_{n m l} e^{-2 \pi i(j-1)(n-1) / N} e^{-2 \pi i(k-1)(m-1) / M} e^{-2 \pi i(k-1)(l-1) / L}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
\sqrt{N M L} S
$$

Finally, note that an unnormalized inverse is implemented in FFT3B. The routine FFT3F is based on the complex FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2T3F/DF2T3F. The reference is:

CALL F2T3F (N1, N2, N3, A, LDA, MDA, B, LDB, MDB, WFF1, WFF2, WFF3, CPY)
The additional arguments are as follows:

WFF1 - Real array of length 4 * N1 +15 initialized by FFTCI. The initialization depends on N1. (Input)

WFF2 - Real array of length 4 * N2 +15 initialized by FFTCI. The initialization depends on N 2 . (Input)

WFF3 - Real array of length $4 *$ N3 +15 initialized by FFTCI. The initialization depends on N3. (Input)
$\boldsymbol{C P Y}$ - Real array of size 2 * MAX(N1, N2, N3). (Workspace)
2. The routine FFT3F is most efficient when N1, N2, and N3 are the product of small primes.
3. If $\operatorname{FFT} 3 \mathrm{~F} / \mathrm{FFT} 3 \mathrm{~B}$ is used repeatedly with the same values for $\mathrm{N} 1, \mathrm{~N} 2$ and N 3 , then use $\operatorname{FFTCI}$ to fill WFF1 $(\mathrm{N}=\mathrm{N} 1)$, $\operatorname{WFF} 2(\mathrm{~N}=\mathrm{N} 2)$, and $\operatorname{WFF} 3(\mathrm{~N}=\mathrm{N} 3)$. Follow this with repeated calls to $\mathrm{F} 2 \mathrm{~T} 3 \mathrm{~F} / \mathrm{F} 2 \mathrm{~T} 3 \mathrm{~B}$. This is more efficient than repeated calls to FFT3F/FFT3B.

## Example

In this example, we compute the Fourier transform of the pure frequency input for a $2 \times 3 \times 4$ array

$$
a_{n m l}=e^{2 \pi i(n-1) 1 / 2} e^{2 \pi i(m-1) 2 / 3} e^{2 \pi i(l-1) 2 / 4}
$$

for $1 \leq n \leq 2,1 \leq m \leq 3$, and $1 \leq l \leq 4$ using the IMSL routine $\operatorname{FFT} 3 F$. The result

$$
\hat{a}=c
$$

has all zeros except in the $(2,3,3)$ position.

```
    USE FFT3F INT
    USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
    USE CONST_INT
    IMPLICIT NONE
    INTEGER LDA, LDB, MDA, MDB, NDA, NDB
    PARAMETER (LDA=2, LDB=2, MDA=3, MDB=3, NDA=4, NDB=4)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
    INTEGER I, J, K, L, M, N, N1, N2, N3, NOUT
    REAL PI
    COMPLEX A(LDA,MDA,NDA), B(LDB,MDB,NDB), C, H
        SPECIFICATIONS FOR INTRINSICS
    INTRINSIC CEXP, CMPLX
    COMPLEX CEXP, CMPLX
        SPECIFICATIONS FOR SUBROUTINES
        SPECIFICATIONS FOR FUNCTIONS
        Get output unit number
    CALL UMACH (2, NOUT)
    PI = CONST('PI')
    C = CMPLX(0.0,2.0*PI)
!
    DO 30 N=1, 2
        DO 20 M=1, 3
            DO 10 L=1, 4
                        H}=C*(N-1)*1/2+C*(M-1)*2/3 + C* (L-1)*2/4
                A(N,M,L) = CEXP(H)
        CONTINUE
        CONTINUE
    CONTINUE
    CALL FFT3F (A, B)
    WRITE (NOUT,99996)
    DO 50 I=1, 2
        WRITE (NOUT,99998) I
        DO 40 J=1, 3
            WRITE (NOUT, 99999) (A (I, J,K),K=1,4)
        CONTINUE
    CONTINUE
    WRITE (NOUT,99997)
    DO 70 I=1, 2
        WRITE (NOUT,99998) I
        DO 60 J=1, 3
        WRITE (NOUT, 99999) (B (I, J,K),K=1, 4)
        CONTINUE
    CONTINUE
!
9 9 9 9 6 ~ F O R M A T ~ ( 1 3 X , ~ ' T h e ~ i n p u t ~ f o r ~ F F T 3 F ~ i s ' )
9 9 9 9 7 ~ F O R M A T ~ ( / , ~ 1 3 X , ~ ' T h e ~ r e s u l t s ~ f r o m ~ F F T 3 F ~ a r e ' )
9 9 9 9 8 ~ F O R M A T ~ ( / , ~ ' ~ F a c e ~ n o . ~ ' , ~ I 1 ) ~
99999 FORMAT (1X, 4('(',F6.2,',',F6.2,')',3X))
    END
```


## Output



## FFT3B

Computes the inverse Fourier transform of a complex periodic three-dimensional array.

## Required Arguments

A- Three-dimensional complex matrix containing the data to be transformed. (Input)
$\boldsymbol{B}$ - Three-dimensional complex matrix containing the inverse Fourier coefficients of A. (Output)
The matrices A and B may be the same.

## Optional Arguments

N1 - Limit on the first subscript of matrices A and B. (Input)
Default: N1 $=\operatorname{size}(\mathrm{A}, 1)$.
N2 - Limit on the second subscript of matrices A and B. (Input)
Default: N2 $=\operatorname{size}(\mathrm{A}, 2)$.
N3 - Limit on the third subscript of matrices A and B. (Input)
Default: N3 $=\operatorname{size}(A, 3)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{size}(\mathrm{A}, 1)$.

MDA - Middle dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: MDA $=\operatorname{size}(A, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{size}(\mathrm{B}, 1)$.
MDB - Middle dimension of B exactly as specified in the dimension statement of the calling program. (Input)
Default: MDB $=\operatorname{size}(B, 2)$.

## FORTRAN 90 Interface

Generic: CALL $\operatorname{FFT} 3 B$ (A, B $[, \ldots]$ )

Specific: The specific interface names are S_FFT3B and D_FFT3B.

## FORTRAN 77 Interface

Single: CALL FFT3B (N1, N2, N3, A, LDA, MDA, B, LDB, MDB)
Double: The double precision name is DFFT3B.

## Description

The routine FFT3B computes the inverse discrete complex Fourier transform of a complex threedimensional array of size $(\mathrm{N} 1=N) \times(\mathrm{N} 2=M) \times(\mathrm{N} 3=L)$. The method used is a variant of the Cooley-Tukey algorithm, which is most efficient when $N, M$, and $L$ are each products of small prime factors. If $N, M$, and $L$ satisfy this condition, then the computational effort is proportional to $N M L \log N M L$. This considerable savings has historically led people to refer to this algorithm as the "fast Fourier transform" or FFT.

Specifically, given an $N \times M \times L$ array $a$, FFT3B returns in $b$

$$
b_{j k l} \sum_{n=1}^{N} \sum_{m=1}^{M} \sum_{l=1}^{L} a_{n m l} e^{2 \pi i(j-1)(n-1) / N} e^{2 \pi i(k-1)(m-1) / M} e^{2 \pi i(k-1)(l-1) / L}
$$

Furthermore, a vector of Euclidean norm $S$ is mapped into a vector of norm

$$
\sqrt{N M L} S
$$

Finally, note that an unnormalized inverse is implemented in FFT3F. The routine FFT3B is based on the complex FFT in FFTPACK. The package FFTPACK was developed by Paul Swarztrauber at the National Center for Atmospheric Research.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F 2 т 3 B/ DF 2 т 3 B. The reference is:

CALL F2T3B (N1, N2, N3, A, LDA, MDA, B, LDB, MDB, WFF1, WFF2, WFF3, CPY)
The additional arguments are as follows:
WFF1 - Real array of length 4 * N1 + 15 initialized by FFTCI. The initialization depends on N1. (Input)

WFF2 - Real array of length 4 * N2 +15 initialized by FFTCI. The initialization depends on N2. (Input)

WFF3 - Real array of length 4 * N3 +15 initialized by FFTCI. The initialization depends on N3. (Input)
$\boldsymbol{C P Y}$ - Real array of size 2 * MAX(N1, N2, N3). (Workspace)
2. The routine FFT T3B is most efficient when $\mathrm{N} 1, \mathrm{~N} 2$, and N 3 are the product of small primes.
3. If FFT3F/FFT3B is used repeatedly with the same values for N 1 , N 2 and N 3 , then use
 repeated calls to $\mathrm{F} 2 \mathrm{~T} 3 \mathrm{~F} / \mathrm{F} 2 \mathrm{~T} 3 \mathrm{~B}$. This is more efficient than repeated calls to fFT3F/FFT3B.

## Example

In this example, we compute the Fourier transform of the $2 \times 3 \times 4$ array

$$
x_{n m l}=n+2(m-1)+2(3)(l-1)
$$

for $1 \leq n \leq 2,1 \leq m \leq 3$, and $1 \leq l \leq 4$ using the IMSL routine FFT3F. The result

$$
a=\hat{x}
$$

is then inverted using FFT3B. Note that the result is an array $b$ satisfying $b=2(3)(4) x=24 x$. In general, FFT 3 B is an unnormalized inverse with expansion factor $N M L$.

```
USE FFT3B INT
USE FFT3F - INT
USE UMACH_INT
IMPLICIT NONE
INTEGER LDA, LDB, MDA, MDB, NDA, NDB
PARAMETER (LDA=2, LDB=2, MDA=3, MDB=3, NDA=4, NDB=4)
                                    SPECIFICATIONS FOR LOCAL VARIABLES
INTEGER I, J, K, L, M, N, N1, N2, N3, NOUT
COMPLEX A(LDA,MDA,NDA), B(LDB,MDB,NDB), X(LDB,MDB,NDB)
```

```
! SPECIFICATIONS FOR INTRINSICS
    INTRINSIC CEXP, CMPLX
    COMPLEX CEXP, CMPLX
! SPECIFICATIONS FOR SUBROUTINES
    CALL UMACH (2, NOUT)
    N1 = 2
    N2 = 3
    N3 = 4
!
    DO 30 N=1, 2
        DO 20 M=1, 3
            DO 10 L=1, 4
                    X(N,M,L) = N + 2* (M-1) + 2* 3* (L-1)
            CONTINUE
        CONTINUE
    CONTINUE
!
    CALL FFT3F (X, A)
    CALL FFT3B (A, B)
!
    WRITE (NOUT,99996)
    DO 50 I=1, 2
        WRITE (NOUT,99998) I
        DO 40 J=1, 3
            WRITE (NOUT,99999) (X (I, J,K),K=1,4)
        CONTINUE
    CONTINUE
!
    WRITE (NOUT,99997)
    DO 70 I=1, 2
        WRITE (NOUT,99998) I
        DO 60 J=1, 3
            WRITE (NOUT, 99999) (A (I, J,K),K=1,4)
        CONTINUE
    CONTINUE
!
    WRITE (NOUT, 99995)
    DO 90 I=1, 2
        WRITE (NOUT,99998) I
        DO 80 J=1, 3
            WRITE (NOUT, 99999) (B (I, J,K),K=1,4)
        CONTINUE
    CONTINUE
    FORMAT (13X, 'The unnormalized inverse is')
    FORMAT (13X, 'The input for FFT3F is')
    FORMAT (/, 13X, 'The results from FFT3F are')
    FORMAT (/, ' Face no. ', Il)
    FORMAT (1X, 4('(',F6.2,',',F6.2,')',3X))
    END
```


## Output

```
The input for FFT3F is
```

| Face no. 1 |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ( 1.00, 0.00) | ( 7.00, | $0.00)$ | ( 13.00, | $0.00)$ | ( 19.00, | $0.00)$ |
| ( 3.00, 0.00) | ( 9.00, | $0.00)$ | ( 15.00, | $0.00)$ | ( 21.00, | $0.00)$ |
| ( 5.00, 0.00) | ( 11.00, | $0.00)$ | ( 17.00, | $0.00)$ | ( 23.00, | $0.00)$ |
| Face no. 2 |  |  |  |  |  |  |
| ( 2.00, 0.00) | ( 8.00, | $0.00)$ | ( 14.00, | $0.00)$ | ( 20.00, | $0.00)$ |
| ( 4.00, 0.00) | ( 10.00, | $0.00)$ | ( 16.00, | $0.00)$ | ( 22.00, | $0.00)$ |
| ( 6.00, 0.00) | ( 12.00, | $0.00)$ | ( 18.00, | $0.00)$ | ( 24.00, | $0.00)$ |
| The results from FFT3F are |  |  |  |  |  |  |
| Face no. 1 |  |  |  |  |  |  |
| (300.00, 0.00) | (-72.00, | 72.00) | (-72.00, | $0.00)$ | (-72.00, | $72.00)$ |
| $(-24.00,13.86)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ |
| $(-24.00,-13.86)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ |
| Face no. 2 |  |  |  |  |  |  |
| $(-12.00,0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ |
| ( 0.00, 0.00) | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ |
| $(0.00,0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ | ( 0.00, | $0.00)$ |
| The unnormalized inverse is |  |  |  |  |  |  |
| Face no. 1 |  |  |  |  |  |  |
| ( 24.00, 0.00) | (168.00, | $0.00)$ | (312.00, | $0.00)$ | (456.00, | $0.00)$ |
| ( 72.00, 0.00) | (216.00, | $0.00)$ | (360.00, | $0.00)$ | (504.00, | $0.00)$ |
| (120.00, 0.00) | (264.00, | $0.00)$ | (408.00, | $0.00)$ | (552.00, | $0.00)$ |
| Face no. 2 |  |  |  |  |  |  |
| ( 48.00, 0.00) | (192.00, | $0.00)$ | (336.00, | $0.00)$ | (480.00, | $0.00)$ |
| ( 96.00, 0.00) | (240.00, | $0.00)$ | (384.00, | $0.00)$ | (528.00, | $0.00)$ |
| (144.00, 0.00) | (288.00, | $0.00)$ | (432.00, | $0.00)$ | (576.00, | $0.00)$ |

## RCONV

Computes the convolution of two real vectors.

## Required Arguments

$\boldsymbol{X}$ - Real vector of length NX. (Input)
$\boldsymbol{Y}$ - Real vector of length NY. (Input)
$Z$ - Real vector of length NZ ontaining the convolution of X and Y . (Output)
ZHAT - Real vector of length NZ containing the discrete Fourier transform of z. (Output)

## Optional Arguments

IDO - Flag indicating the usage of RCONV. (Input)
Default: IDO $=0$.

0 If this is the only call to RCONV.

If RCONV is called multiple times in sequence with the same NX, NY, and IPAD, IDO should be set to

1 on the first call
2 on the intermediate calls

3 on the final call.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
Default: NX $=\operatorname{size}(X, 1)$.
$N Y$ - Length of the vector Y. (Input)
Default: NY $=\operatorname{size}(Y, 1)$.
IPAD - IPAD should be set to zero for periodic data or to one for nonperiodic data. (Input) Default: $\operatorname{IPAD}=0$.
$N Z$ - Length of the vector Z. (Input/Output)
Upon input: When IPAD is zero, NZ must be at least MAX(NX, NY). When IPAD is one, $N Z$ must be greater than or equal to the smallest integer greater than or equal to $(N X+N Y-1)$ of the form $\left(2^{\alpha}\right)$ * $\left(3^{\beta}\right)$ * $\left(5^{\gamma}\right)$ where alpha, beta, and gamma are nonnegative integers. Upon output, the value for NZ that was used by RCONV.
Default: NZ = size $(Z, 1)$.

## FORTRAN 90 Interface

Generic: CALL RCONV (X, Y, Z, ZHAT [,...])

Specific: The specific interface names are S_RCONV and D_RCONV.

## FORTRAN 77 Interface

Single: CALL RCONV (IDO, NX, X, NY, Y, IPAD, NZ, Z, ZHAT)
Double: The double precision name is DRCONV.

## Description

The routine RCONV computes the discrete convolution of two sequences $x$ and $y$. More precisely, let $n_{x}$ be the length of $x$ and $n_{y}$ denote the length of $y$. If a circular convolution is desired, then IPAD must be set to zero. We set

$$
n_{z}:=\max \left\{n_{x}, n_{y}\right\}
$$

and we pad out the shorter vector with zeroes. Then, we compute

$$
z_{i}=\sum_{j=1}^{n_{z}} x_{i-j+1} y_{j}
$$

where the index on $x$ is interpreted as a positive number between 1 and $n_{z}$, modulo $n_{z}$.
The technique used to compute the $z_{i}$ 's is based on the fact that the (complex discrete) Fourier transform maps convolution into multiplication. Thus, the Fourier transform of $z$ is given by

$$
\hat{z}(n)=\hat{x}(n) \hat{y}(n)
$$

where

$$
\hat{z}(n)=\sum_{m=1}^{n_{z}} z_{m} e^{-2 \pi i(m-1)(n-1) / n_{z}}
$$

The technique used here to compute the convolution is to take the discrete Fourier transform of $x$ and $y$, multiply the results together component-wise, and then take the inverse transform of this product. It is very important to make sure that $n_{z}$ is a product of small primes if IPAD is set to zero. If $n_{z}$ is a product of small primes, then the computational effort will be proportional to $n_{z} \log \left(n_{z}\right)$. If IPAD is one, then a good value is chosen for $n_{z}$ so that the Fourier transforms are efficient and $n_{z} \geq n_{x}+n_{y}-1$. This will mean that both vectors will be padded with zeroes.
We point out that no complex transforms of $x$ or $y$ are taken since both sequences are real, we can take real transforms and simulate the complex transform above. This can produce a savings of a factor of six in time as well as save space over using the complex transform.

## Comments

1. Workspace may be explicitly provided, if desired, by use of R2ONV/DR2ONV. The reference is:

CALL R2ONV (IDO, NX, X, NY, Y, IPAD, NZ, Z, ZHAT, XWK, YWK, WK)
The additional arguments are as follows:
$\boldsymbol{X W K}$ - Real work array of length Nz.
YWK - Real work array of length NZ.
$W K-$ Real work arrary of length $2 * N Z+15$.
2. Informational error

Type Code
41 The length of the vector $z$ must be large enough to hold the results. An acceptable length is returned in NZ.

## Example

In this example, we compute both a periodic and a non-periodic convolution. The idea here is that one can compute a moving average of the type found in digital filtering using this routine. The averaging operator in this case is especially simple and is given by averaging five consecutive points in the sequence. The periodic case tries to recover a noisy sin function by averaging five nearby values. The nonperiodic case tries to recover the values of an exponential function contaminated by noise. The large error for the last value printed has to do with the fact that the convolution is averaging the zeroes in the "pad" rather than function values. Notice that the signal size is 100 , but we only report the errors at ten points.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER NFLTR, NY, A
    PARAMETER (NFLTR=5, NY=100)
!
    INTEGER I, IPAD, K, MOD, NOUT, NZ
    REAL ABS, EXP, F1, F2, FLOAT, FLTR(NFLTR), &
        FLTRER, ORIGER, SIN, TOTAL1, TOTAL2, TWOPI, X, &
        Y(NY), Z(2* (NFLTR+NY-1)), ZHAT(2* (NFLTR+NY-1))
    INTRINSIC ABS, EXP, FLOAT, MOD, SIN
    F1(X) = SIN(X)
    F2(X) = EXP(X)
!
    CALL RNSET (1234579)
    CALL UMACH (2, NOUT)
    TWOPI = CONST('PI')
    TWOPI = 2.0*TWOPI
    DO 10 I=1, 5
        FLTR(I) = 0.2
    1 0 \text { CONTINUE}
        SET UP Y-VECTOR FOR THE PERIODIC
                                CASE.
    DO 20 I=1, NY
        X = TWOPI*FLOAT (I-1)/FLOAT (NY-1)
        Y(I) = RNUNF()
        Y(I) = F1(X) + 0.5*Y(I) - 0.25
    20 CONTINUE
!
!
    NZ = 2*(NFLTR+NY-1)
    CALL RCONV (FLTR, Y, Z, ZHAT, IPAD=0, NZ=NZ)
                            PRINT RESULTS
    WRITE (NOUT,99993)
    WRITE (NOUT,99995)
    TOTAL1 = 0.0
    TOTAL2 = 0.0
    DO 30 I=1, NY
        IF (I .GE. NY-1) THEN
        K=I - NY + 2
    ELSE
```

```
            K = I + 2
            END IF
!
        X = TWOPI*FLOAT(I-1)/FLOAT (NY-1)
        ORIGER = ABS(Y(I)-F1(X))
        FLTRER = ABS (Z (K)-F1 (X))
        IF (MOD(I,11) .EQ. 1) WRITE (NOUT,99997) X, F1(X), ORIGER, &
            FLTRER
        TOTAL1 = TOTAL1 + ORIGER
        TOTAL2 = TOTAL2 + FLTRER
    3 0 ~ C O N T I N U E ~
    WRITE (NOUT,99998) TOTAL1/FLOAT (NY)
    WRITE (NOUT,99999) TOTAL2/FLOAT (NY)
                                    SET UP Y-VECTOR FOR THE NONPERIODIC
                                    CASE.
    DO 40 I=1, NY
        A = FLOAT (I-1)/FLOAT (NY-1)
        Y(I) = RNUNF()
        Y(I) = F2(A) + 0.5*Y(I) - 0.25
    4 0 ~ C O N T I N U E ~
!
! - < <
    NZ = 2*(NFLTR+NY-1)
    CALL RCONV (FLTR, Y, Z, ZHAT, IPAD=1)
                                    PRINT RESULTS
    WRITE (NOUT,99994)
    WRITE (NOUT,99996)
    TOTAL1 = 0.0
    TOTAL2 = 0.0
    DO 50 I=1, NY
        X = FLOAT (I-1)/FLOAT (NY-1)
        ORIGER = ABS(Y(I)-F2(X))
        FLTRER = ABS (Z(I+2)-F2(X))
        IF (MOD(I,11) .EQ. 1) WRITE (NOUT,99997) X, F2(X), ORIGER, &
            FLTRER
        TOTAL1 = TOTAL1 + ORIGER
        TOTAL2 = TOTAL2 + FLTRER
    5 0 ~ C O N T I N U E ~
    WRITE (NOUT,99998) TOTAL1/FLOAT (NY)
    WRITE (NOUT,99999) TOTAL2/FLOAT (NY)
9 9 9 9 3 ~ F O R M A T ~ ( ' ~ P e r i o d i c ~ C a s e ' )
99994 FORMAT (/,' Nonperiodic Case')
99995 FORMAT (8X, 'x', 9X, 'sin(x)', 6X, 'Original Error', 5X, &
    'Filtered Error')
99996 FORMAT (8X, 'x', 9X, 'exp(x)', 6X, 'Original Error', 5X, &
            'Filtered Error')
99997 FORMAT (1X, F10.4, F13.4, 2F18.4)
99998 FORMAT (' Average absolute error before filter:', F10.5)
99999 FORMAT (' Average absolute error after filter:', F11.5)
    END
```


## Output

Periodic Case
$x \quad$ sin $(x) \quad$ Original Error Filtered Error

| 0.0000 | 0.0000 | 0.0811 |  | 0.0587 |
| :---: | :---: | :---: | :---: | :---: |
| 0.6981 | 0.6428 | 0.0226 |  | 0.0781 |
| 1.3963 | 0.9848 | 0.1526 |  | 0.0529 |
| 2.0944 | 0.8660 | 0.0959 |  | 0.0125 |
| 2.7925 | 0.3420 | 0.1747 |  | 0.0292 |
| 3.4907 | -0.3420 | 0.1035 |  | 0.0238 |
| 4.1888 | -0.8660 | 0.0402 |  | 0.0595 |
| 4.8869 | -0.9848 | 0.0673 |  | 0.0798 |
| 5.5851 | -0.6428 | 0.1044 |  | 0.0074 |
| 6.2832 | 0.0000 | 0.0154 |  | 0.0018 |
| Average absolute error before filter: Average absolute error after filter: |  |  | 0.12481 |  |
|  |  |  | 0.04778 |  |
| Nonperiodic Case |  |  |  |  |
| x | exp (x) | Original Error | Filt | ered Error |
| 0.0000 | 1.0000 | 0.1476 |  | 0.3915 |
| 0.1111 | 1.1175 | 0.0537 |  | 0.0326 |
| 0.2222 | 1.2488 | 0.1278 |  | 0.0932 |
| 0.3333 | 1.3956 | 0.1136 |  | 0.0987 |
| 0.4444 | 1.5596 | 0.1617 |  | 0.0964 |
| 0.5556 | 1.7429 | 0.0071 |  | 0.0662 |
| 0.6667 | 1.9477 | 0.1248 |  | 0.0713 |
| 0.7778 | 2.1766 | 0.1556 |  | 0.0158 |
| 0.8889 | 2.4324 | 0.1529 |  | 0.0696 |
| 1.0000 | 2.7183 | 0.2124 |  | 1.0562 |
| Average | solute er | fore filter: | 0.12538 |  |
| Average | solute er | fer filter: | 0.07764 |  |

## CCONV

Computes the convolution of two complex vectors.

## Required Arguments

$\boldsymbol{X}$ - Complex vector of length NX . (Input)
$\boldsymbol{Y}$ - Complex vector of length NY. (Input)
$Z$ - Complex vector of length NZ containing the convolution of X and Y . (Output)
ZHAT - Complex vector of length nZ containing the discrete complex Fourier transform of Z. (Output)

## Optional Arguments

IDO - Flag indicating the usage of CCONV. (Input) Default: IDO $=0$.

IDO Usage
0 If this is the only call to CCONV.

If CCONV is called multiple times in sequence with the same NX, NY, and IPAD, IDO should be set to:

1 on the first call

2 on the intermediate calls
3 on the final call.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
Default: $N X=\operatorname{size}(X, 1)$.
$N Y$ - Length of the vector Y. (Input)
Default: NY $=\operatorname{size}(Y, 1)$.

IPAD - IPAD should be set to zero for periodic data or to one for nonperiodic data. (Input) Default: $I P A D=0$.
$N Z$ — Length of the vector Z. (Input/Output)
Upon input: When IPAD is zero, NZ must be at least MAX(NX, NY). When IPAD is one, NZ must be greater than or equal to the smallest integer greater than or equal to $(N X+N Y-1)$ of the form $\left(2^{\alpha}\right) *\left(3^{\beta}\right) *\left(5^{\gamma}\right)$ where alpha, beta, and gamma are nonnegative integers. Upon output, the value for NZ that was used by CCONV. Default: NZ $=\operatorname{size}(\mathrm{Z}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CCONV (X,Y, Z, ZHAT [,...])
Specific: The specific interface names are S_CCONV and D_CCONV.

## FORTRAN 77 Interface

Single: CALL CCONV (IDO, NX, X, NY, Y, IPAD, NZ, Z, ZHAT)
Double: The double precision name is DCCONV.

## Description

The subroutine CCONV computes the discrete convolution of two complex sequences $x$ and $y$. More precisely, let $n_{x}$ be the length of $x$ and $n_{y}$ denote the length of $y$. If a circular convolution is desired, then IPAD must be set to zero. We set

$$
n_{z}:=\max \left\{n_{x}, n_{y}\right\}
$$

and we pad out the shorter vector with zeroes. Then, we compute

$$
z_{i}=\sum_{j=1}^{n_{z}} x_{i-j+1} y_{j}
$$

where the index on $x$ is interpreted as a positive number between 1 and $n_{z}$, modulo $n_{z}$.
The technique used to compute the $z_{i}$ 's is based on the fact that the (complex discrete) Fourier transform maps convolution into multiplication. Thus, the Fourier transform of $z$ is given by

$$
\hat{z}(n)=\hat{x}(n) \hat{y}(n)
$$

where

$$
\hat{z}(n)=\sum_{m=1}^{n_{z}} z_{m} e^{-2 \pi i(m-1)(n-1) / n_{z}}
$$

The technique used here to compute the convolution is to take the discrete Fourier transform of $x$ and $y$, multiply the results together component-wise, and then take the inverse transform of this product. It is very important to make sure that $n_{z}$ is a product of small primes if IPAD is set to zero. If $n_{z}$ is a product of small primes, then the computational effort will be proportional to $n_{z} \log \left(n_{z}\right)$. If IPAD is one, then a a good value is chosen for $n_{z}$ so that the Fourier transforms are efficient and $n_{z} \geq n_{x}+n_{y}-1$. This will mean that both vectors will be padded with zeroes.

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2ONV/DC2ONV. The reference is:

CALL C2ONV (IDO, NX, X, NY, Y, IPAD, NZ, Z, ZHAT, XWK, YWK, WK)
The additional arguments are as follows:
$X W K$ - Complex work array of length NZ.
YWK - Complex work array of length NZ.
$\boldsymbol{W K}$ - Real work array of length 6 * nz +15 .
2. Informational error

Type Code
41 The length of the vector z must be large enough to hold the results. An acceptable length is returned in NZ.

## Example

In this example, we compute both a periodic and a non-periodic convolution. The idea here is that one can compute a moving average of the type found in digital filtering using this routine. The averaging operator in this case is especially simple and is given by averaging five consecutive points in the sequence. The periodic case tries to recover a noisy function $f_{1}(x)=\cos (x)+i \sin (x)$ by averaging five nearby values. The nonperiodic case tries to recover the values of the function
$f_{2}(x)=e^{x} f_{1}(x)$ contaminated by noise. The large error for the first and last value printed has to do with the fact that the convolution is averaging the zeroes in the "pad" rather than function values. Notice that the signal size is 100 , but we only report the errors at ten points.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER NFLTR, NY
PARAMETER (NFLTR=5, NY=100)
INTEGER I, IPAD, K, MOD, NOUT, NZ
REAL CABS, COS, EXP, FLOAT, FLTRER, ORIGER, &
    SIN, TOTAL1, TOTAL2, TWOPI, X, T1, T2
COMPLEX CMPLX, F1, F2, FLTR(NFLTR), Y(NY), Z(2*(NFLTR+NY-1)), &
            ZHAT (2* (NFLTR+NY-1))
INTRINSIC CABS, CMPLX, COS, EXP, FLOAT, MOD, SIN
                    DEFINE FUNCTIONS
F1(X) = CMPLX(COS (X),SIN (X))
F2(X) = EXP(X)*CMPLX(COS (X),SIN (X))
CALL RNSET (1234579)
CALL UMACH (2, NOUT)
TWOPI = CONST('PI')
TWOPI = 2.0*TWOPI
                                    SET UP THE FILTER
CALL CSET(NFLTR,(0.2,0.0),FLTR,1)
                                    SET UP Y-VECTOR FOR THE PERIODIC
                                    CASE.
DO 20 I=1, NY
        X = TWOPI*FLOAT (I-1)/FLOAT (NY-1)
        T1 = RNUNF ()
        T2 = RNUNF()
        Y(I) = F1(X) + CMPLX(0.5*T1-0.25,0.5*T2-0.25)
CONTINUE
                                    CALL THE CONVOLUTION ROUTINE FOR THE
                                    PERIODIC CASE.
NZ = 2*(NFLTR+NY-1)
CALL CCONV (FLTR, Y, Z, ZHAT)
                                    PRINT RESULTS
WRITE (NOUT, 99993)
WRITE (NOUT,99995)
TOTAL1 = 0.0
TOTAL2 = 0.0
DO 30 I=1, NY
                                    COMPUTE THE OFFSET FOR THE Z-VECTOR
    IF (I .GE. NY-1) THEN
        K = I - NY + 2
    ELSE
        K = I + 2
    END IF
    X = TWOPI*FLOAT (I-1)/FLOAT (NY-1)
    ORIGER = CABS (Y(I) -F1 (X))
    FLTRER = CABS (Z (K) - F1 (X))
    IF (MOD(I,11) .EQ. 1) WRITE (NOUT,99997) X, F1(X), ORIGER, &
```

    20
    ```
                FLTRER
            TOTAL1 = TOTAL1 + ORIGER
            TOTAL2 = TOTAL2 + FLTRER
    3 0 ~ C O N T I N U E
    WRITE (NOUT,99998) TOTAL1/FLOAT (NY)
    WRITE (NOUT,99999) TOTAL2/FLOAT (NY)
                                    SET UP Y-VECTOR FOR THE NONPERIODIC
                                    CASE.
    DO 40 I=1, NY
        X = FLOAT (I-1)/FLOAT (NY-1)
        T1 = RNUNF ()
        T2 = RNUNF()
        Y(I) = F2(X) + CMPLX(0.5*T1-0.25,0.5*T2-0.25)
    4 0 ~ C O N T I N U E ~
!
!
                                    CALL THE CONVOLUTION ROUTINE FOR THE
                                    NONPERIODIC CASE
    NZ = 2*(NFLTR+NY-1)
    CALL CCONV (FLTR, Y, Z, ZHAT, IPAD=1)
                                    PRINT RESULTS
    WRITE (NOUT,99994)
    WRITE (NOUT,99996)
    TOTAL1 = 0.0
    TOTAL2 = 0.0
    DO 50 I=1, NY
        X = FLOAT (I-1)/FLOAT (NY-1)
        ORIGER = CABS (Y(I) -F2(X))
        FLTRER = CABS (Z (I+2)-F2 (X))
        IF (MOD(I,11) .EQ. 1) WRITE (NOUT,99997) X, F2(X), ORIGER, &
        FLTRER
        TOTAL1 = TOTAL1 + ORIGER
        TOTAL2 = TOTAL2 + FLTRER
    5 0 ~ C O N T I N U E ~
    WRITE (NOUT,99998) TOTAL1/FLOAT (NY)
    WRITE (NOUT,99999) TOTAL2/FLOAT (NY)
9 9 9 9 3 ~ F O R M A T ~ ( ' ~ P e r i o d i c ~ C a s e ' )
9 9 9 9 4 ~ F O R M A T ~ ( / , ~ ' ~ N o n p e r i o d i c ~ C a s e ' ) ,
99995 FORMAT (8X, 'x', 15X, 'f1(x)', 8X, 'Original Error', 5X, &
            'Filtered Error')
99996 FORMAT (8X, 'x', 15X, 'f2(x)', 8X, 'Original Error', 5X, &
            'Filtered Error')
99997 FORMAT (1X, F10.4, 5X, '(', F7.4, ',', F8.4, ' )', 5X, F8.4, &
                10X, F8.4)
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ A v e r a g e ~ a b s o l u t e ~ e r r o r ~ b e f o r e ~ f i l t e r : ' , ~ F 1 1 . 5 )
99999 FORMAT (' Average absolute error after filter:', F12.5)
    END
```


## Output

| Periodic Case <br> x | fl (x) |  |  |  |
| :---: | :---: | :---: | ---: | ---: |
| 0.0000 | $(1.0000$, | $0.0000)$ | Original Error | Filtered Error |
| 0.6981 | $(0.7660$, | $0.6428)$ | 0.1666 | 0.0773 |
| 1.3963 | $(0.1736$, | $0.9848)$ | 0.1685 | 0.1399 |
| 2.0944 | $(-0.5000$, | $0.8660)$ | 0.2171 | 0.0368 |
| 2.7925 | $(-0.9397$, | $0.3420)$ | 0.1147 | 0.0142 |
|  |  |  |  | 0.0200 |



## RCORL

Computes the correlation of two real vectors.

## Required Arguments

$\boldsymbol{X}$ - Real vector of length N. (Input)
$\boldsymbol{Y}$ - Real vector of length N. (Input)
$Z$ - Real vector of length NZ containing the correlation of X and Y . (Output)
ZHAT - Real vector of length NZ containing the discrete Fourier transform of z. (Output)

## Optional Arguments

IDO - Flag indicating the usage of RCORL. (Input) Default: $I D O=0$.

## IDO Usage

0 If this is the only call to RCORL.
If RCORL is called multiple times in sequence with the same NX, NY, and IPAD, IDO should be set to:

1 on the first call
$N$ - Length of the X and Y vectors. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)$.

IPAD - I PAD should be set as follows. (Input)
Default: $\operatorname{IPAD}=0$.
IPAD Value

IPAD 0 for periodic data with X and Y different.

IPAD 1 for nonperiodic data with X and Y different.
IPAD 2 for periodic data with X and Y identical.
IPAD 3 for nonperiodic data with X and Y identical.
$N Z$ - Length of the vector Z. (Input/Output)
Upon input: When IPAD is zero or two, NZ must be at least $(2 * N-1)$. When IPAD is one or three, NZ must be greater than or equal to the smallest integer greater than or equal to $(2 * N-1)$ of the form $\left(2^{\alpha}\right) *\left(3^{\beta}\right) *\left(5^{\gamma}\right)$ where alpha, beta, and gamma are nonnegative integers. Upon output, the value for NZ that was used by RCORL.
Default: NZ $=\operatorname{size}(\mathrm{Z}, 1)$.

## FORTRAN 90 Interface

Generic: CALL RCORL (X, Y, Z, ZHAT [,...])
Specific: The specific interface names are S_RCORL and D_RCORL.

## FORTRAN 77 Interface

Single: CALL RCORL (IDO, N, X, Y, IPAD, NZ, Z, ZHAT)
Double: The double precision name is DRCORL.

## Description

The subroutine RCORL computes the discrete correlation of two sequences $x$ and $y$. More precisely, let $n$ be the length of $x$ and $y$. If a circular correlation is desired, then IPAD must be set to zero (for $x$ and $y$ distinct) and two (for $x=y$ ). We set (on output)

$$
\begin{aligned}
n_{z}=n & \text { if IPAD }=0,2 \\
n_{z}=2^{\alpha} 3^{\beta} 5^{\gamma} \geq 2 n-1 & \text { if IPAD }=1,3
\end{aligned}
$$

where $\alpha, \beta, \gamma$ are nonnegative integers yielding the smallest number of the type $2^{\alpha} 3^{\beta} 5^{\gamma}$ satisfying the inequality. Once $n_{z}$ is determined, we pad out the vectors with zeroes. Then, we compute

$$
z_{i}=\sum_{j=1}^{n_{z}} x_{i+j-1} y_{j}
$$

where the index on $x$ is interpreted as a positive number between one and $n_{z}$, modulo $n_{z}$. Note that this means that

$$
z_{n_{z}-k}
$$

contains the correlation of $x(\cdot-\mathrm{k}-1)$ with $y$ as $k=0,1, \ldots, n_{z} / 2$. Thus, if $x(k-1)=y(k)$ for all $k$, then we would expect

$$
z_{n_{z}}
$$

to be the largest component of $z$.
The technique used to compute the $z_{i}$ 's is based on the fact that the (complex discrete) Fourier transform maps correlation into multiplication. Thus, the Fourier transform of $z$ is given by

$$
\hat{z}_{j}=\hat{x}_{j} \overline{\hat{y}}_{j}
$$

where

$$
\hat{z}_{j}=\sum_{m=1}^{n_{z}} z_{m} e^{-2 \pi i(m-1)(j-1) / n_{z}}
$$

Thus, the technique used here to compute the correlation is to take the discrete Fourier transform of $x$ and the conjugate of the discrete Fourier transform of $y$, multiply the results together component-wise, and then take the inverse transform of this product. It is very important to make sure that $n_{z}$ is a product of small primes if IPAD is set to zero or two. If $n_{z}$ is a product of small primes, then the computational effort will be proportional to $n_{z} \log \left(n_{z}\right)$. If IPAD is one or three, then a good value is chosen for $n_{z}$ so that the Fourier transforms are efficient and $n_{z} \geq 2 n-1$. This will mean that both vectors will be padded with zeroes.
We point out that no complex transforms of $x$ or $y$ are taken since both sequences are real, and we can take real transforms and simulate the complex transform above. This can produce a savings of a factor of six in time as well as save space over using the complex transform.

## Comments

1. Workspace may be explicitly provided, if desired, by use of R2ORL/DR2ORL. The reference is:

CALL R2ORL (IDO, N, X, Y, IPAD, NZ, Z, ZHAT, XWK, YWK, WK)
The additional arguments are as follows:
$\boldsymbol{X W K}$ - Real work array of length NZ.
$\mathbf{Y W K}$ - Real work array of length NZ.
$\boldsymbol{W} \boldsymbol{K}$ - Real work arrary of length $2 * \mathrm{NZ}+15$.
2. Informational error

Type Code
41 The length of the vector z must be large enough to hold the results. An acceptable length is returned in NZ.

## Example

In this example, we compute both a periodic and a non-periodic correlation between two distinct signals $x$ and $y$. In the first case we have 100 equally spaced points on the interval $[0,2 \pi]$ and $f_{1}(x)=\sin (x)$. We define $x$ and $y$ as follows

$$
\begin{array}{ll}
x_{i}=f_{1}\left(2 \pi \frac{i-1}{n-1}\right) & i=1, \ldots, n \\
y_{i} & =f_{1}\left(2 \pi \frac{i-1}{n-1}+\frac{\pi}{2}\right)
\end{array} \quad i=1, \ldots, n
$$

Note that the maximum value of $z$ (the correlation of $x$ with $y$ ) occurs at $i=26$, which corresponds to the offset.
The nonperiodic case uses the function $f_{2}(x)=\sin \left(x^{2}\right)$. The two input signals are on the interval $[0,4 \pi]$.

$$
\begin{array}{lll}
x_{i} & =f_{2}\left(4 \pi \frac{i-1}{n-1}\right) & i=1, \ldots, n \\
y_{i} & =f_{2}\left(4 \pi \frac{i-1}{n-1}+\pi\right) & i=1, \ldots, n
\end{array}
$$

The offset of $x$ to $y$ is again (roughly) 26 and this is where $z$ has its maximum value.

```
USE IMSL_LIBRARIES
IMPLICIT NONE
INTEGER N
PARAMETER (N=100)
INTEGER I, IPAD, K, NOUT, NZ
REAL A, F1, F2, FLOAT, PI, SIN, X(N), XNORM, &
    Y(N), YNORM, Z(4*N), ZHAT (4*N)
INTRINSIC FLOAT, SIN
F1(A) = SIN(A)
F2(A) = SIN (A*A)
CALL UMACH (2, NOUT)
PI = CONST('pi')
DO 10 I=1, N
    X(I) = F1(2.0*PI*FLOAT(I-1)/FLOAT (N-1))
    Y(I) = F1(2.0*PI*FLOAT(I-1)/FLOAT(N-1) +PI/2.0)
```

!

10 CONTINUE

```
!
!
    NZ = 2*N
    CALL RCORL (X, Y, Z, ZHAT)
!
!
    XNORM = SNRM2 (N, X,1)
    YNORM = SNRM2 (N,Y,1)
    DO 20 I=1, N
        Z(I) = Z(I)/(XNORM*YNORM)
    20 CONTINUE
    K = ISMAX (N, Z,1)
        Print results for the periodic
        case.
    WRITE (NOUT,99995)
    WRITE (NOUT,99994)
    WRITE (NOUT,99997)
    WRITE (NOUT,99998) K
    WRITE (NOUT,99999) K, Z(K)
!
    DO 30 I=1, N
        X(I) = F2(4.0*PI*FLOAT (I-1)/FLOAT (N-1))
        Y(I) = F2(4.0*PI*FLOAT (I-1)/FLOAT (N-1) +PI)
    3 0 ~ C O N T I N U E ~
!
!
    Call the correlation routine for the
    nonperiodic case.
    NZ = 4*N
    CALL RCORL (X, Y, Z, ZHAT, IPAD=1)
                                    Find the element of Z with the
                                    largest normalized value.
    XNORM = SNRM2 (N, X,1)
    YNORM = SNRM2 (N,Y,1)
    DO 40 I=1, N
        Z(I) = Z(I)/(XNORM*YNORM)
    40 CONTINUE
    K = ISMAX (N,Z,1)
        Print results for the nonperiodic
        case.
    WRITE (NOUT,99996)
    WRITE (NOUT,99994)
    WRITE (NOUT,99997)
    WRITE (NOUT, 99998) K
    WRITE (NOUT,99999) K, Z(K)
99994 FORMAT (1X, 28('-'))
9 9 9 9 5 ~ F O R M A T ~ ( ' ~ C a s e ~ \# 1 : ~ P e r i o d i c ~ d a t a ' )
9 9 9 9 6 ~ F O R M A T ~ ( / , ~ ' ~ C a s e ~ \# 2 : ~ N o n p e r i o d i c ~ d a t a ' )
9 9 9 9 7 ~ F O R M A T ~ ( ' ~ T h e ~ e l e m e n t ~ o f ~ Z ~ w i t h ~ t h e ~ l a r g e s t ~ n o r m a l i z e d ~ ' ) ,
99998 FORMAT (' value is Z(', I2, ').')
99999 FORMAT (' The normalized value of Z(', I2, ') is', F6.3)
    END
```


## Output

```
Example #1: Periodic case
----------------------------
The element of Z with the largest normalized value is Z(26).
The normalized value of Z(26) is 1.000
Example #2: Nonperiodic case
-----------------------------
The element of Z with the largest normalized value is Z(26).
The normalized value of Z(26) is 0.661
```


## CCORL

Computes the correlation of two complex vectors.

## Required Arguments

$\boldsymbol{X}$ - Complex vector of length N. (Input)
$\boldsymbol{Y}$ - Complex vector of length N. (Input)
$Z$ - Complex vector of length NZ containing the correlation of X and Y . (Output)
ZHAT - Complex vector of length NZ containing the inverse discrete complex Fourier transform of z. (Output)

## Optional Arguments

IDO - Flag indicating the usage of CCORL. (Input)
Default: $I D O=0$.

IDO Usage
0 If this is the only call to CCORL.

If CCORL is called multiple times in sequence with the same NX, NY, and IPAD, IDO should be set to:

1 on the first call
2 on the intermediate calls
3 on the final call.
$N$ - Length of the X and Y vectors. (Input)
Default: $N=\operatorname{size}(X, 1)$.

IPAD - IPAD should be set as follows. (Input)
IPAD $=0$ for periodic data with X and Y different. $\operatorname{IPAD}=1$ for nonperiodic data with X
and $Y$ different. $I P A D=2$ for periodic data with $X$ and $Y$ identical. $I P A D=3$ for nonperiodic data with X and Y identical.
Default: $\operatorname{IPAD}=0$.
$N Z$ - Length of the vector Z. (Input/Output)
Upon input: When IPAD is zero or two, NZ must be at least ( $2 * N-1$ ). When IPAD is one or three, NZ must be greater than or equal to the smallest integer greater than or equal to $(2 * N-1)$ of the form $\left(2^{\alpha}\right) *\left(3^{\beta}\right) *\left(5^{\gamma}\right)$ where alpha, beta, and gamma are nonnegative integers. Upon output, the value for NZ that was used by CCORL.
Default: NZ $=\operatorname{size}(\mathrm{Z}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CCORL (X, Y, Z, ZHAT [,...])
Specific: The specific interface names are S_CCORL and D_CCORL.

## FORTRAN 77 Interface

Single: CALL CCORL (IDO, N, X, Y, IPAD, NZ, Z, ZHAT)
Double: The double precision name is DCCORL.

## Description

The subroutine CCORL computes the discrete correlation of two complex sequences $x$ and $y$. More precisely, let $n$ be the length of $x$ and $y$. If a circular correlation is desired, then IPAD must be set to zero (for $x$ and $y$ distinct) and two (for $x=y$ ). We set (on output)

$$
\begin{aligned}
n_{z}=n & \text { if IPAD }=0,2 \\
n_{z}=2^{\alpha} 3^{\beta} 5^{\gamma} \geq 2 n-1 & \text { if } \operatorname{IPAD}=1,3
\end{aligned}
$$

where $\alpha, \beta, \gamma$ are nonnegative integers yielding the smallest number of the type $2^{\alpha} 3^{\beta} 5^{\gamma}$ satisfying the inequality. Once $n_{z}$ is determined, we pad out the vectors with zeroes. Then, we compute

$$
z_{i}=\sum_{j=1}^{n_{z}} x_{i+j-1} \bar{y}_{j}
$$

where the index on $x$ is interpreted as a positive number between one and $n_{z}$, modulo $n_{z}$. Note that this means that

$$
z_{n_{z}-k}
$$

contains the correlation of $x(\cdot-\mathrm{k}-1)$ with $y$ as $k=0,1, \ldots, n_{z} / 2$. Thus, if $x(k-1)=y(k)$ for all $k$, then we would expect

$$
\mathfrak{R} z_{n_{z}}
$$

to be the largest component of $\mathfrak{R z}$.

The technique used to compute the $z_{i}$ 's is based on the fact that the (complex discrete) Fourier transform maps correlation into multiplication. Thus, the Fourier transform of $z$ is given by

$$
\hat{z}_{j}=\hat{x}_{j} \overline{\hat{y}}_{j}
$$

where

$$
\hat{z}_{j}=\sum_{m=1}^{n_{z}} z_{m} e^{-2 \pi i(m-1)(j-1) / n_{z}}
$$

Thus, the technique used here to compute the correlation is to take the discrete Fourier transform of $x$ and the conjugate of the discrete Fourier transform of $y$, multiply the results together component-wise, and then take the inverse transform of this product. It is very important to make sure that $n_{z}$ is a product of small primes if IPAD is set to zero or two. If $n_{z}$ is a product of small primes, then the computational effort will be proportional to $n_{z} \log \left(n_{z}\right)$. If IPAD is one or three, then a good value is chosen for $n_{z}$ so that the Fourier transforms are efficient and $n_{z} \geq 2 n-1$. This will mean that both vectors will be padded with zeroes.

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2ORL/DC2ORL. The reference is:

CALL C2ORL (IDO, N, X, Y, IPAD, NZ, Z, ZHAT, XWK, YWK, WK)
The additional arguments are as follows:
$X W K$ — Complex work array of length NZ.
YWK - Complex work array of length nz.
$W K-$ Real work arrary of length $6 *$ NZ +15.
2. Informational error

Type Code
41 The length of the vector z must be large enough to hold the results. An acceptable length is returned in NZ.

## Example

In this example, we compute both a periodic and a non-periodic correlation between two distinct signals $x$ and $y$. In the first case, we have 100 equally spaced points on the interval $[0,2 \pi]$ and $f_{1}(x)=\cos (x)+i \sin (x)$. We define $x$ and $y$ as follows

$$
\begin{array}{ll}
x_{i} & =f_{1}\left(2 \pi \frac{i-1}{n-1}\right) \\
y_{i} & =f_{1}\left(2 \pi \frac{i-1}{n-1}+\frac{\pi}{2}\right)
\end{array} \quad i=1, \ldots, n, n, n, ~ l i=1, n
$$

Note that the maximum value of $z$ (the correlation of $x$ with $y$ ) occurs at $i=26$, which corresponds to the offset.

The nonperiodic case uses the function $f_{2}(x)=\cos \left(x^{2}\right)+i \sin \left(x^{2}\right)$. The two input signals are on the interval $[0,4 \pi]$.

$$
\begin{array}{lll}
x_{i} & =f_{2}\left(4 \pi \frac{i-1}{n-1}\right) & i=1, \ldots, n \\
y_{i} & =f_{2}\left(4 \pi \frac{i-1}{n-1}+\pi\right) & i=1, \ldots, n
\end{array}
$$

The offset of $x$ to $y$ is again (roughly) 26 and this is where $z$ has its maximum value.

```
    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=100)
    INTEGER I, IPAD, K, NOUT, NZ
    REAL A, COS, F1, F2, FLOAT, PI, SIN, &
        XNORM, YNORM, ZREAL1(4*N)
    COMPLEX CMPLX, X(N), Y(N), Z(4*N), ZHAT(4*N)
    INTRINSIC CMPLX, COS, FLOAT, SIN
        Define functions
    F1(A) = CMPLX(COS (A),SIN (A))
    F2(A) = CMPLX(COS (A*A),SIN (A*A))
    CALL RNSET (1234579)
    CALL UMACH (2, NOUT)
    PI = CONST('pi')
                Set up the vectors for the
                periodic case.
    DO 10 I=1, N
        X(I) = F1 (2.0*PI*FLOAT(I-1)/FLOAT (N-1))
        Y(I) = F1(2.0*PI*FLOAT(I-1)/FLOAT (N-1) +PI/2.0)
    1 0 ~ C O N T I N U E ~
    NZ = 2*N
    CALL CCORL (X, Y, Z, ZHAT, IPAD=0, NZ=NZ)
                        Find the element of Z with the
                                largest normalized real part.
    XNORM = SCNRM2(N,X,1)
    YNORM = SCNRM2(N,Y,1)
    DO 20 I=1, N
        ZREAL1(I) = REAL(Z(I))/(XNORM*YNORM)
    20 CONTINUE
    K = ISMAX(N,ZREAL1,1)
        Print results for the periodic
                case.
    WRITE (NOUT,99995)
    WRITE (NOUT,99994)
    WRITE (NOUT,99997)
```

```
    WRITE (NOUT,99998) K
    WRITE (NOUT,99999) K, ZREAL1(K)
        Set up the vectors for the
        nonperioddic case.
    DO 30 I=1, N
        X(I) = F2(4.0*PI*FLOAT (I-1)/FLOAT (N-1))
        Y(I) = F2(4.0*PI*FLOAT(I-1)/FLOAT (N-1) +PI)
    3 0 ~ C O N T I N U E
!
!
    NZ = 4*N
    CALL CCORL (X, Y, Z, ZHAT, IPAD=1, NZ=NZ)
                        Find the element of z with the
                        largest normalized real part.
    XNORM = SCNRM2 (N, X,1)
    YNORM = SCNRM2 (N,Y,1)
    DO 40 I=1, N
        ZREAL1(I) = REAL(Z(I))/(XNORM*YNORM)
    40 CONTINUE
    K = ISMAX (N,ZREAL1,1)
!
!
        Print results for the nonperiodic
        case.
    WRITE (NOUT,99996)
    WRITE (NOUT,99994)
    WRITE (NOUT,99997)
    WRITE (NOUT,99998) K
    WRITE (NOUT,99999) K, ZREAL1(K)
99994 FORMAT (1X, 28('-'))
9 9 9 9 5 ~ F O R M A T ~ ( ' ~ C a s e ~ \# 1 : ~ p e r i o d i c ~ d a t a ' )
9 9 9 9 6 ~ F O R M A T ~ ( / , ~ ' ~ C a s e ~ \# 2 : ~ n o n p e r i o d i c ~ d a t a ' )
99997 FORMAT (' The element of Z with the largest normalized ')
99998 FORMAT (' real part is Z(', I2, ').')
99999 FORMAT (' The normalized value of real(Z(', I2, ')) is', F6.3)
    END
```


## Output

```
Example #1: periodic case
```

---------------------------------
The element of $Z$ with the largest normalized real part is $Z(26)$.
The normalized value of real(Z(26)) is 1.000
Example \#2: nonperiodic case
The element of $Z$ with the largest normalized real part is $Z(26)$.
The normalized value of real(Z(26)) is 0.638

## INLAP

Computes the inverse Laplace transform of a complex function.

## Required Arguments

$\boldsymbol{F}$ - User-supplied FUNCTION to which the inverse Laplace transform will be computed. The form is $F(Z)$, where

Z - Complex argument. (Input) F - The complex function value. (Output)

F must be declared EXTERNAL in the calling program. F should also be declared COMPLEX.
$\boldsymbol{T}$ - Array of length N containing the points at which the inverse Laplace transform is desired. (Input)
$T(I)$ must be greater than zero for all $I$.
FINV - Array of length N whose I-th component contains the approximate value of the Laplace transform at the point $\mathrm{T}(\mathrm{I})$. (Output)

## Optional Arguments

$N$ - Number of points at which the inverse Laplace transform is desired. (Input) Default: $\mathrm{N}=\operatorname{size}(\mathrm{T}, 1)$.

ALPHA - An estimate for the maximum of the real parts of the singularities of F . If unknown, set ALPHA $=0$. (Input)
Default: ALPHA $=0.0$.
$\boldsymbol{K M A X}$ - The number of function evaluations allowed for each $T$ (I). (Input) Default: KMAX $=500$.

RELERR - The relative accuracy desired. (Input)
Default: RELERR $=1.1920929 \mathrm{e}-5$ for single precision and $2.22 \mathrm{~d}-10$ for double precision.

## FORTRAN 90 Interface

Generic: CALL INLAP (F, T, FINV [,...])
Specific: The specific interface names are S_INLAP and D_INLAP.

## FORTRAN 77 Interface

Single: CALL INLAP (F, N, T, ALPHA, RELERR, KMAX, FINV)
Double: The double precision name is DINLAP.

## Description

The routine INLAP computes the inverse Laplace transform of a complex-valued function. Recall that if $f$ is a function that vanishes on the negative real axis, then we can define the Laplace transform of $f$ by

$$
L[f](s):=\int_{0}^{\infty} e^{-s x} f(x) d x
$$

It is assumed that for some value of $s$ the integrand is absolutely integrable.
The computation of the inverse Laplace transform is based on applying the epsilon algorithm to the complex Fourier series obtained as a discrete approximation to the inversion integral. The initial algorithm was proposed by K.S. Crump (1976) but was significantly improved by de Hoog et al. (1982). Given a complex-valued transform $F(s)=L[f](s)$, the trapezoidal rule gives the approximation to the inverse transform

$$
g(t)=\left(e^{\alpha t} / T\right) \Re\left\{\frac{1}{2} F(\alpha)+\sum_{k=1}^{\infty} F\left(\alpha+\frac{i k \pi}{T}\right) \exp \left(\frac{i k \pi t}{T}\right)\right\}
$$

This is the real part of the sum of a complex power series in $z=\exp (i \pi t / T)$, and the algorithm accelerates the convergence of the partial sums of this power series by using the epsilon algorithm to compute the corresponding diagonal Pade approximants. The algorithm attempts to choose the order of the Pade approximant to obtain the specified relative accuracy while not exceeding the maximum number of function evaluations allowed. The parameter $\alpha$ is an estimate for the maximum of the real parts of the singularities of $F$, and an incorrect choice of $\alpha$ may give false convergence. Even in cases where the correct value of $\alpha$ is unknown, the algorithm will attempt to estimate an acceptable value. Assuming satisfactory convergence, the discretization error $E:=g-f$ satisfies

$$
E=\sum_{n=1}^{\infty} e^{-2 n \alpha T} f(2 n T+t)
$$

It follows that if $|f(t)| \leq M e^{\beta t}$, then we can estimate the expression above to obtain (for $0 \leq t \leq 2 T$ )

$$
E \leq M e^{\alpha t} /\left(e^{2 T(\alpha-\beta)}-1\right)
$$

## Comments

Informational errors
Type Code
41 The algorithm was not able to achieve the accuracy requested within KMAX function evaluations for some $T(I)$.
42 Overflow is occurring for a particular value of $T$.

## Example

We invert the Laplace transform of the simple function $(s-1)^{-2}$ and print the computed answer, the true solution and the difference at five different points. The correct inverse transform is $x e^{x}$.

```
USE INLAP_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER I, KMAX, N, NOUT
REAL ALPHA, DIF(5), EXP, FINV(5), FLOAT, RELERR, T(5), &
TRUE (5)
COMPLEX F
INTRINSIC EXP, FLOAT
EXTERNAL F
! Get output unit number
CALL UMACH (2, NOUT)
DO 10 I=1, 5
        T(I) = FLOAT(I) - 0.5
    10 CONTINUE
    N = 5
    ALPHA = 1.0E0
    RELERR = 5.0E-4
    CALL INLAP (F, T, FINV, ALPHA=ALPHA, RELERR=RELERR)
                                    Evaluate the true solution and the
                                    difference
    DO 20 I=1, 5
        TRUE(I) = T(I)*EXP(T(I))
        DIF(I) = TRUE(I) - FINV(I)
    CONTINUE
    WRITE (NOUT,99999) (T (I),FINV (I),TRUE (I), DIF (I), I=1, 5)
99999 FORMAT (7X, 'T', 8X, 'FINV', 9X, 'TRUE', 9X, 'DIFF', /, &
        5(1X,E9.1,3X,1PE10.3,3X,1PE10.3,3X,1PE10.3,/))
        END
!
    COMPLEX FUNCTION F (S)
    COMPLEX S
    F = 1./(S-1.)**2
    RETURN
    END
```


## Output

| T | FINV | TRUE | DIFF |
| :---: | :---: | :---: | :---: |
| $0.5 \mathrm{E}+00$ | $8.244 \mathrm{E}-01$ | $8.244 \mathrm{E}-01$ | $-4.768 \mathrm{E}-06$ |
| $1.5 \mathrm{E}+00$ | $6.723 \mathrm{E}+00$ | $6.723 \mathrm{E}+00$ | $-3.481 \mathrm{E}-05$ |
| $2.5 \mathrm{E}+00$ | $3.046 \mathrm{E}+01$ | $3.046 \mathrm{E}+01$ | $-1.678 \mathrm{E}-04$ |
| $3.5 \mathrm{E}+00$ | $1.159 \mathrm{E}+02$ | $1.159 \mathrm{E}+02$ | $-6.027 \mathrm{E}-04$ |
| $4.5 \mathrm{E}+00$ | $4.051 \mathrm{E}+02$ | $4.051 \mathrm{E}+02$ | $-2.106 \mathrm{E}-03$ |

## SINLP

Computes the inverse Laplace transform of a complex function.

## Required Arguments

$\boldsymbol{F}$ - User-supplied FUNCTION to which the inverse Laplace transform will be computed. The form is $F(Z)$, where
z - Complex argument. (Input)
F - The complex function value. (Output)
F must be declared EXTERNAL in the calling program. F must also be declared COMPLEX.
$\boldsymbol{T}$ - Vector of length N containing points at which the inverse Laplace transform is desired. (Input)
$T(I)$ must be greater than zero for all $I$.
FINV - Vector of length $N$ whose I-th component contains the approximate value of the inverse Laplace transform at the point $\mathrm{T}(\mathrm{I})$. (Output)

## Optional Arguments

$\boldsymbol{N}$ - The number of points at which the inverse Laplace transform is desired. (Input) Default: $\mathrm{N}=\operatorname{size}(\mathrm{T}, 1)$.

SIGMAO - An estimate for the maximum of the real parts of the singularities of F . (Input) If unknown, set SIGMA0 $=0.0$.
Default: SIGMA0 $=0 . e 0$.
EPSTOL - The required absolute uniform pseudo accuracy for the coefficients and inverse Laplace transform values. (Input)
Default: $\operatorname{EPSTOL}=1.1920929 \mathrm{e}-5$ for single precision and $2.22 \mathrm{~d}-10$ for double precision.
$\boldsymbol{E R R V E C}$ - Vector of length eight containing diagnostic information. (Output) All components depend on the intermediately generated Laguerre coefficients. See Comments.

## FORTRAN 90 Interface

Generic: CALL SINLP (F, T, FINV [, ...])
Specific: The specific interface names are S_SINLP and D_SINLP.

## FORTRAN 77 Interface

Single: CALL SINLP (F, N, T, SIGMA0, EPSTOL, ERRVEC, FINV)
Double: The double precision name is DSINLP.

## Description

The routine SINLP computes the inverse Laplace transform of a complex-valued function. Recall that if $f$ is a function that vanishes on the negative real axis, then we can define the Laplace transform of $f$ by

$$
L[f](s):=\int_{0}^{\infty} e^{-s x} f(x) d x
$$

It is assumed that for some value of $s$ the integrand is absolutely integrable.
The computation of the inverse Laplace transform is based on a modification of Weeks' method (see W.T. Weeks (1966)) due to B.S. Garbow et. al. (1988). This method is suitable when $f$ has continuous derivatives of all orders on $[0, \infty)$. In this situation, this routine should be used in place of the IMSL routine INLAP. It is especially efficient when multiple function values are desired. In particular, given a complex-valued function $F(s)=L[f](s)$, we can expand $f$ in a Laguerre series whose coefficients are determined by $F$. This is fully described in B.S. Garbow et. al. (1988) and Lyness and Giunta (1986).

The algorithm attempts to return approximations $g(t)$ to $f(t)$ satisfying

$$
\left|\frac{g(t)-f(t)}{e^{\sigma t}}\right|<\varepsilon
$$

where $\varepsilon:=$ EPSTOL and $\sigma:=$ SIGMA $>$ SIGMA0. The expression on the left is called the pseudo error. An estimate of the pseudo error is available in ERRVEC(1).

The first step in the method is to transform $F$ to $\phi$ where

$$
\phi(z)=\frac{b}{1-z} F\left(\frac{b}{1-z}-\frac{b}{2}+\sigma\right)
$$

Then, if $f$ is smooth, it is known that $\phi$ is analytic in the unit disc of the complex plane and hence has a Taylor series expansion

$$
\phi(z)=\sum_{s=0}^{\infty} a_{s} z^{s}
$$

which converges for all $z$ whose absolute value is less than the radius of convergence $R_{c}$. This number is estimated in ERRVEC(6). In ERRVEC(5), we estimate the smallest number $K$ which satisfies

$$
\left|a_{s}\right|<\frac{K}{R^{s}}
$$

for all $R<R_{c}$.
The coefficients of the Taylor series for $\phi$ can be used to expand $f$ in a Laguerre series

$$
f(t)=e^{\sigma t} \sum_{s=0}^{\infty} a_{s} e^{-b t / 2} L_{s}(b t)
$$

## Comments

1. Workspace may be explicitly provided, if desired, by use of S2NLP/DS2NLP. The reference is:
```
CALL S2NLP (F,N,T,SIGMA0, EPSTOL, ERRVEC, FINV, SIGMA, BVALUE, MTOP, WK,
IFLOVC)
```

The additional arguments are as follows:
SIGMA - The first parameter of the Laguerre expansion. If SIGMA is not greater than SIGMA 0 , it is reset to SIGMA $0+0.7$. (Input)

BVALUE - The second parameter of the Laguerre expansion. If BVALUE is less than 2.0 * (SIGMA - SIGMA $)$, it is reset to 2.5 * (SIGMA - SIGMA0). (Input)

MTOP - An upper limit on the number of coefficients to be computed in the Laguerre expansion. MTOP must be a multiple of four. Note that the maximum number of Laplace transform evaluations is мтор/2 + 2. (Default: 1024.) (Input)

WK - Real work vector of length 9 * MTOP/4.
IFLOVC - Integer vector of length N , the I-th component of which contains the overflow/underflow indicator for the computed value of FINV(I). (Output) See Comment 3.
2. Informational errors

Type Code
11 Normal termination, but with estimated error bounds slightly larger than EPSTOL. Note, however, that the actual errors on the final results may be smaller than EPSTOL as bounds independent of T are pessimistic.
32 Normal calculation, terminated early at the roundoff error level estimate because this estimate exceeds the required accuracy (usually due to overly optimistic expectation by the user about attainable accuracy).
43 The decay rate of the coefficients is too small. It may improve results to use S2NLP and increase MTOP.
44 The decay rate of the coefficients is too small. In addition, the roundoff error level is such that required accuracy cannot be reached.
45 No error bounds are returned as the behavior of the coefficients does not enable reasonable prediction. Results are probably wrong. Check the value of SIGMA0. In this case, each of ERRVEC( $J$ ), $J=1, \ldots, 5$, is set to - 1.0.
3. The following are descriptions of the vectors ERRVEC and IFLOVC.

ERRVEC - Real vector of length eight.
$\operatorname{ERRVEC}(1)=$ Overall estimate of the pseudo error, $\operatorname{ERRVEC}(2)+\operatorname{ERRVEC}(3)+$ $\operatorname{ERRVEC}(4)$. Pseudo error $=$ absolute error $/ \exp ($ sigma $*$ tvalue $)$.
$\operatorname{ERRVEC}(2)=$ Estimate of the pseudo discretization error.
$\operatorname{ERRVEC}(3)=$ Estimate of the pseudo truncation error.
$\operatorname{ERRVEC}(4)=$ Estimate of the pseudo condition error on the basis of minimal noise levels in the function values.
$\operatorname{ERRVEC}(5)=\mathrm{K}$, the coefficient of the decay function for ACOEF, the coefficients of the Laguerre expansion.
$\operatorname{ERRVEC}(6)=R$, the base of the decay function for ACOEF. Here $\operatorname{abs}(\operatorname{ACOEF}(J+1)$ ).LE.K/R**J for J.GE.MACT/2, where MACT is the number of Laguerre coefficients actually computed.
$\operatorname{ERRVEC}(7)=A L P H A$, the logarithm of the largest ACOEF.
$\operatorname{ERRVEC}(8)=\mathrm{BETA}$, the logarithm of the smallest nonzero ACOEF.
IFLOVC - Integer vector of length N containing the overflow/underflow indicators for FINV. For each I, the value of IFLOVC(I) signifies the following.
$0=$ Normal termination.
$1=$ The value of the inverse Laplace transform is found to be too large to be representable; $\operatorname{FINV}(\mathrm{I})$ is set to $\operatorname{AMACH}(6)$.
$-1=$ The value of the inverse Laplace transform is found to be too small to be representable; $\operatorname{FINV}(\mathrm{I})$ is set to 0.0 .
$2=$ The value of the inverse Laplace transform is estimated to be too large, even before the series expansion, to be representable; $\operatorname{FINV}(\mathrm{I})$ is set to $\operatorname{AMACH}(6)$.
$-2=$ The value of the inverse Laplace transform is estimated to be too small, even before the series expansion, to be representable; $\operatorname{FINV}(I)$ is set to 0.0 .

## Example

We invert the Laplace transform of the simple function $(s-1)^{-2}$ and print the computed answer, the true solution, and the difference at five different points. The correct inverse transform is $x e^{x}$.

```
USE SINLP_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER I, NOUT
REAL DIF(5), ERRVEC(8), EXP, FINV(5), FLOAT, RELERR, &
    SIGMA0, T(5), TRUE(5), EPSTOL
COMPLEX
    F
```

```
    INTRINSIC EXP, FLOAT
    EXTERNAL F
! Get output unit number
    CALL UMACH (2, NOUT)
!
    DO 10 I=1, 5
        T(I) = FLOAT(I) - 0.5
    1 0 ~ C O N T I N U E
    SIGMAO = 1.OEO
    RELERR = 5.0E-4
    EPSTOL = 1.OE-4
    CALL SINLP (F, T, FINV, SIGMA0=SIGMA0, EPSTOL=RELERR)
                                    Evaluate the true solution and the
                                    difference
    DO 20 I=1, 5
        TRUE(I) = T(I)*EXP(T(I))
        DIF(I) = TRUE(I) - FINV(I)
    20 CONTINUE
!
    WRITE (NOUT,99999) (T (I),FINV (I),TRUE (I), DIF (I), I=1, 5)
99999 FORMAT (7X, 'T', 8X, 'FINV', 9X, 'TRUE', 9X, 'DIFF', /, &
            5 (1X, E9.1,3X,1PE10.3,3X,1PE10.3,3X,1PE10.3, /))
        END
!
    COMPLEX FUNCTION F (S)
    COMPLEX S
l
    F=1./(S-1.)**2
    RETURN
    END
```


## Output

| T | FINV | TRUE | DIFF |
| :---: | :---: | :---: | ---: |
| $0.5 \mathrm{E}+00$ | $8.244 \mathrm{E}-01$ | $8.244 \mathrm{E}-01$ | $-2.086 \mathrm{E}-06$ |
| $1.5 \mathrm{E}+00$ | $6.723 \mathrm{E}+00$ | $6.723 \mathrm{E}+00$ | $-8.583 \mathrm{E}-06$ |
| $2.5 \mathrm{E}+00$ | $3.046 \mathrm{E}+01$ | $3.046 \mathrm{E}+01$ | $0.000 \mathrm{E}+00$ |
| $3.5 \mathrm{E}+00$ | $1.159 \mathrm{E}+02$ | $1.159 \mathrm{E}+02$ | $2.289 \mathrm{E}-05$ |
| $4.5 \mathrm{E}+00$ | $4.051 \mathrm{E}+02$ | $4.051 \mathrm{E}+02$ | $-2.136 \mathrm{E}-04$ |

## Chapter 7: Nonlinear Equations

## Routines

7.1. Zeros of a Polynomial
Real coefficients using Laguerre method ............................ ZPLRC ..... 1184
Real coefficients using Jenkins-Traub method ZPORC ..... 1186
Complex coefficients ZPOCC ..... 1188
7.2. Zero(s) of a Function
Zeros of a complex analytic function ..... ZANLY1189
Zero of a real function with sign changes ..... ZBREN ..... 1192
Zeros of a real function ..... 1195
7.3. Root of a System of Equations
Finite-difference Jacobian NEQNF ..... 1198
Analytic Jacobian NEQNJ ..... 1201
Broyden's update and Finite-difference Jacobian .NEQBF ..... 1204
Broyden's update and Analytic Jacobian. NEQBJ ..... 1210

## Usage Notes

## Zeros of a Polynomial

A polynomial function of degree $n$ can be expressed as follows:

$$
p(z)=a_{n} z^{n}+a_{n-1} z^{n-1}+\ldots+a_{1} z+a_{0}
$$

where $a_{n} \neq 0$.
There are three routines for zeros of a polynomial. The routines ZPLRC and ZPORC find zeros of the polynomial with real coefficients while the routine ZPOCC finds zeros of the polynomial with complex coefficients.

The Jenkins-Traub method is used for the routines ZPORC and ZPOCC; whereas ZPLRC uses the Laguerre method. Both methods perform well in comparison with other methods. The JenkinsTraub algorithm usually runs faster than the Laguerre method. Furthermore, the routine ZANLY in the next section can also be used for the complex polynomial.

## Zero(s) of a Function

The routines ZANLY and ZREAL use Müller's method to find the zeros of a complex analytic function and real zeros of a real function, respectively. The routine ZBREN finds a zero of a real function, using an algorithm that is a combination of interpolation and bisection. This algorithm requires the user to supply two points such that the function values at these two points have opposite sign. For functions where it is difficult to obtain two such points, ZREAL can be used.

## Root of System of Equations

A system of equations can be stated as follows:

$$
f_{i}(x)=0, \text { for } i=1,2, \ldots, n
$$

where $x \varepsilon \mathbf{R}^{n}$.
The routines NEQNF and NEQNJ use a modified Powell hybrid method to find a zero of a system of nonlinear equations. The difference between these two routines is that the Jacobian is estimated by a finite-difference method in NEQNF, whereas the user has to provide the Jacobian for NEQNJ. It is advised that the Jacobian-checking routine, CHJAC (see Chapter 8, Optimization), be used to ensure the accuracy of the user-supplied Jacobian.

The routines NEQBF and NEQBJ use a secant method with Broyden's update to find a zero of a system of nonlinear equations. The difference between these two routines is that the Jacobian is estimated by a finite-difference method in NEQBF; whereas the user has to provide the Jacobian for NEQBJ. For more details, see Dennis and Schnabel (1983, Chapter 8).

## ZPLRC

Finds the zeros of a polynomial with real coefficients using Laguerre's method.

## Required Arguments

COEFF - Vector of length NDEG +1 containing the coefficients of the polynomial in increasing order by degree. (Input)
The polynomial is
$\operatorname{COEFF}(\operatorname{NDEG}+1) * Z * * \operatorname{NDEG}+\operatorname{COEFF}(\operatorname{NDEG}) * Z * *(\operatorname{NDEG}-1)+\ldots+\operatorname{COEFF}(1)$.
ROOT - Complex vector of length NDEG containing the zeros of the polynomial. (Output)

## Optional Arguments

NDEG - Degree of the polynomial. $1 \leq$ NDEG $\leq 100$ (Input)
Default: NDEG $=\operatorname{size}(\operatorname{COEFF}, 1)-1$.

## FORTRAN 90 Interface

Generic: CALL ZPLRC (COEFF, ROOT [,...])

Specific: The specific interface names are $S \_$ZPLRC and D_ZPLRC.

## FORTRAN 77 Interface

Single: CALL ZPLRC (NDEG, COEFF, ROOT)
Double: The double precision name is DZPLRC.

## Description

Routine ZPLRC computes the $n$ zeros of the polynomial

$$
p(z)=a_{n} z^{n}+a_{n-1} z^{n-1}+\ldots+a_{1} z+a_{0}
$$

where the coefficients $a_{i}$ for $i=0,1, \ldots, n$ are real and $n$ is the degree of the polynomial.
The routine ZPLRC is a modification of B.T. Smith's routine ZERPOL (Smith 1967) that uses Laguerre's method. Laguerre's method is cubically convergent for isolated zeros and linearly convergent for multiple zeros. The maximum length of the step between successive iterates is restricted so that each new iterate lies inside a region about the previous iterate known to contain a zero of the polynomial. An iterate is accepted as a zero when the polynomial value at that iterate is smaller than a computed bound for the rounding error in the polynomial value at that iterate. The original polynomial is deflated after each real zero or pair of complex zeros is found. Subsequent zeros are found using the deflated polynomial.

## Comments

Informational errors

| Type | Code |
| :---: | :--- |
| 3 |  | | 1 | The first several coefficients of the polynomial are equal to zero. Several of <br> the last roots will be set to machine infinity to compensate for this problem. |
| :--- | :--- |
| 3 | 2Fewer than NDEG zeros were found. The ROOT vector will contain the value <br> for machine infinity in the locations that do not contain zeros. |

## Example

This example finds the zeros of the third-degree polynomial

$$
p(z)=z^{3}-3 z^{2}+4 z-2
$$

where $z$ is a complex variable.

```
USE ZPLRC INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER NDEG
PARAMETER (NDEG=3)
REAL COEFF (NDEG+1)
COMPLEX ZERO (NDEG)
```

```
! Set values of COEFF
    COEFF = (-2.0 4.0 -3.0 1.0)
    DATA COEFF/-2.0, 4.0, -3.0, 1.0/
!
    CALL ZPLRC (COEFF, ZERO, NDEG)
    CALL WRCRN ('The zeros found are', ZERO, 1, NDEG, 1)
!
END
```


## Output

The zeros found are
$\left(\begin{array}{cc}\stackrel{1}{2} & { }^{2} \\ (1.000,1.000) & 3 \\ (1.000,-1.000) & (1.000,0.000)\end{array}\right.$

## ZPORC

Finds the zeros of a polynomial with real coefficients using the Jenkins-Traub three-stage algorithm.

## Required Arguments

COEFF - Vector of length NDEG +1 containing the coefficients of the polynomial in increasing order by degree. (Input)
The polynomial is $\operatorname{COEFF}(\operatorname{NDEG}+1) * Z * * \operatorname{NDEG}+\operatorname{COEFF}(\operatorname{NDEG}) * Z * *(\operatorname{NDEG}-1)+\ldots+\operatorname{COEFF}(1)$.

ROOT - Complex vector of length NDEG containing the zeros of the polynomial. (Output)

## Optional Arguments

NDEG - Degree of the polynomial. $1 \leq$ NDEG $\leq 100$ (Input)
Default: NDEG $=\operatorname{size}(\operatorname{COEFF}, 1)-1$.

## FORTRAN 90 Interface

Generic: CALL ZPORC (COEFF, ROOT [,...])

Specific: The specific interface names are S_ZPORC and D_ZPORC.

## FORTRAN 77 Interface

Single: CALL ZPORC (NDEG, COEFF, ROOT)
Double: The double precision name is DZPORC.

## Description

Routine ZPORC computes the $n$ zeros of the polynomial

$$
p(z)=a_{n} z^{n}+a_{n-1} z^{n-1}+\ldots+a_{1} z+a_{0}
$$

where the coefficients $a_{i}$ for $i=0,1, \ldots, n$ are real and $n$ is the degree of the polynomial.
The routine ZPORC uses the Jenkins-Traub three-stage algorithm (Jenkins and Traub 1970; Jenkins 1975). The zeros are computed one at a time for real zeros or two at a time for complex conjugate pairs. As the zeros are found, the real zero or quadratic factor is removed by polynomial deflation.

## Comments

Informational errors

| Type | Code |  |
| :---: | :--- | :--- |
| 3 | 1 | The first several coefficients of the polynomial are equal to zero. Several of <br> the last roots will be set to machine infinity to compensate for this problem. |
| 3 | 2 | Fewer than NDEG zeros were found. The ROOT vector will contain the value <br> for machine infinity in the locations that do not contain zeros. |

## Example

This example finds the zeros of the third-degree polynomial

$$
p(z)=z^{3}-3 z^{2}+4 z-2
$$

where $z$ is a complex variable.

```
USE ZPORC_INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER NDEG
PARAMETER (NDEG=3)
REAL COEFF (NDEG+1)
COMPLEX ZERO(NDEG)
    DATA COEFF/-2.0, 4.0, -3.0, 1.0/
    CALL ZPORC (COEFF, ZERO)
    CALL WRCRN ('The zeros found are', ZERO, 1, NDEG, 1)
    END
```

$!$
!

## Output

The zeros found are
$(1.000,0.000)^{1}(1.000,1.000) \quad(1.000,-1.000)$

## ZPOCC

Finds the zeros of a polynomial with complex coefficients.

## Required Arguments

COEFF - Complex vector of length NDEG +1 containing the coefficients of the polynomial in increasing order by degree. (Input)
The polynomial is
$\operatorname{COEFF}(\operatorname{NDEG}+1) * Z * * \operatorname{NDEG}+\operatorname{COEFF}(\operatorname{NDEG}) * Z * *(\operatorname{NDEG}-1)+\ldots+\operatorname{COEFF}(1)$.
ROOT - Complex vector of length NDEG containing the zeros of the polynomial. (Output)

## Optional Arguments

NDEG - Degree of the polynomial. $1 \leq$ NDEG $<50$ (Input)
Default: NDEG $=\operatorname{size}(\operatorname{COEFF}, 1)-1$.

## FORTRAN 90 Interface

Generic: CALL ZPOCC (COEFF, ROOT [,...])

Specific: The specific interface names are S_ZPOCC and D_ZPOCC.

## FORTRAN 77 Interface

Single: CALL ZPOCC (NDEG, COEFF, ROOT)
Double: The double precision name is DZPOCC.

## Description

Routine ZPOCC computes the $n$ zeros of the polynomial

$$
p(z)=a_{n} z^{n}+a_{n-1} z^{n-1}+\ldots+a_{1} z+a_{0}
$$

where the coefficients $a_{i}$ for $i=0,1, \ldots, n$ are real and $n$ is the degree of the polynomial.
The routine zPOCC uses the Jenkins-Traub three-stage complex algorithm (Jenkins and Traub 1970, 1972). The zeros are computed one at a time in roughly increasing order of modulus. As each zero is found, the polynomial is deflated to one of lower degree.

## Comments

Informational errors

| Type | Code |
| :---: | :--- | :--- |
| 3 |  | | 1 | The first several coefficients of the polynomial are equal to zero. Several of <br> the last roots will be set to machine infinity to compensate for this problem. |
| :--- | :--- |
| 3 | 2 | | Fewer than NDEG zeros were found. The ROOT vector will contain the value |
| :--- |
| for machine infinity in the locations that do not contain zeros. |

## Example

This example finds the zeros of the third-degree polynomial

$$
p(z)=z^{3}-(3+6 i) z^{2}-(8-12 i) z+10
$$

where $z$ is a complex variable.

```
    USE ZPOCC_INT
    USE WRCRN_INT
    IMPLICIT NONE
! Declare variables
    INTEGER NDEG
    PARAMETER (NDEG=3)
    COMPLEX COEFF(NDEG+1), ZERO(NDEG)
        Set values of COEFF
        COEFF = ( 10.0 + 0.0i )
                                (-8.0 + 12.0i )
                                ( -3.0 - 6.0i )
        ( 1.0 + 0.0i )
    DATA COEFF/(10.0,0.0), (-8.0,12.0), (-3.0,-6.0), (1.0,0.0)/
    CALL ZPOCC (COEFF, ZERO)
    CALL WRCRN ('The zeros found are', ZERO, 1, NDEG, 1)
    END
```

!

## Output

The zeros found are
1 2 3
$(1.000,1.000)(1.000,2.000)(1.000,3.000)$

## ZANLY

Finds the zeros of a univariate complex function using Müller's method.

## Required Arguments

$\boldsymbol{F}$ - User-supplied COMPLEX FUNCTION to compute the value of the function of which the zeros will be found. The form is $F(z)$, where

Z - The complex value at which the function is evaluated. (Input) $Z$ should not be changed by F.

F - The computed complex function value at the point Z . (Output) F must be declared EXTERNAL in the calling program.
$\mathbf{Z}$ - A complex vector of length NKNOWN + nNEW. (Output) Z(1), $\ldots, \mathrm{Z}$ (NKNOWN) contain the known zeros. Z (NKNOWN +1 ), $\ldots, \mathrm{z}$ (NKNOWN + NNEW) contain the new zeros found by ZANLY. If ZINIT is not needed, ZINIT and z can share the same storage locations.

## Optional Arguments

$\boldsymbol{E R R A B S}$ - First stopping criterion. (Input)
Let $\operatorname{FP}(Z)=F(Z) / P$ where $P=(Z-Z(1)) *(Z-Z(2)) * \ldots *(Z-Z(K-1))$ and $z(1), \ldots, Z(K-1)$ are previously found zeros.
 then z is accepted as a zero. Default: ERRABS = 1.e-4 for single precision and 1.d-8 for double precision.

ERRREL - Second stopping criterion is the relative error. (Input)
A zero is accepted if the difference in two successive approximations to this zero is within ERRREL. ERRREL must be less than 0.01 ; otherwise, 0.01 will be used. Default: ERRREL $=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.

NKNOWN - The number of previously known zeros, if any, that must be stored in ZINIT(1), $\ldots, \operatorname{ZINIT}(N K N O W N)$ prior to entry to ZANLY. (Input) NKNOWN must be set equal to zero if no zeros are known.
Default: NKNOWN $=0$.

NNEW - The number of new zeros to be found by ZANLY. (Input) Default: NNEW $=1$.

NGUESS - The number of initial guesses provided. (Input) These guesses must be stored in ZINIT(NKNOWN +1 ), $\ldots$, ZINIT(NKNOWN + NGUESS). NGUESS must be set equal to zero if no guesses are provided.
Default: NGUESS $=0$.

ITMAX - The maximum allowable number of iterations per zero. (Input)
Default: $\operatorname{ITMAX}=100$.

ZINIT - A complex vector of length NKNOWN + NNEW. (Input)
$\operatorname{ZINIT}(1), \ldots, \operatorname{ZINIT}(N K N O W N)$ must contain the known zeros. ZINIT(NKNOWN +1 ), $\ldots$, ZINIT(NKNOWN + NNEW) may, on user option, contain initial guesses for the NNEW new zeros that are to be computed. If the user does not provide an initial guess, zero is used.

INFO - An integer vector of length NKNOWN + NNEW. (Output)
$\operatorname{INFO}(\mathrm{J})$ contains the number of iterations used in finding the $J$-th zero when convergence was achieved. If convergence was not obtained in ITMAX iterations, INFO(J) will be greater than ITMAX.

## FORTRAN 90 Interface

Generic: CALL ZANLY (F, Z [,...])
Specific: The specific interface names are S_ZANLY and D_ZANLY.

## FORTRAN 77 Interface

Single: CALL ZANLY (F, ERRABS, ERRREL, NKNOWN, NNEW, NGUESS, ZINIT, ITMAX, Z, INFO)

Double: The double precision name is DZANLY.

## Example

This example finds the zeros of the equation $f(z)=z^{3}+5 z^{2}+9 z+45$, where $z$ is a complex variable.

```
USE ZANLY_INT
USE WRCRN_INT
IMPLICIT NONE
Declare variables
    INTEGER INFO(3), NGUESS, NNEW
    COMPLEX F, Z(3), ZINIT(3)
    EXTERNAL F
                                    Set the guessed zero values in ZINIT
                                    ZINIT = (1.0+1.0i 1.0+1.0i 1.0+1.0i)
    DATA ZINIT/3*(1.0,1.0)/
        Set values for all input parameters
    NNEW = 3
    NGUESS = 3
    CALL ZANLY (F, Z, NNEW=NNEW, NGUESS=NGUESS, &
        ZINIT=ZINIT, INFO=INFO)
            Print results
    CALL WRCRN ('The zeros are', Z)
    END
    External complex function
```

$!$
$\mathrm{F}=\mathrm{Z} * * 3+5.0 * Z * * 2+9.0 * Z+45.0$
RETURN
END

## Output

The zeros are
$(0.000,3.000)(0.000,-3.000) \quad(-5.000,0.000)$

## ZBREN

Finds a zero of a real function that changes sign in a given interval.

## Required Arguments

$\boldsymbol{F}$ - User-supplied FUNCTION to compute the value of the function of which a zero will be found. The form is $F(X)$, where

X - The point at which the function is evaluated. (Input) $X$ should not be changed by $F$.

F - The computed function value at the point X . (Output) F must be declared EXTERNAL in the calling program.

A-See B. (Input/Output)
$B$ - On input, the user must supply two points, $A$ and $B$, such that $F(A)$ and $F(B)$ are opposite in sign. (Input/Output)
On output, both $A$ and $B$ are altered. $B$ will contain the best approximation to the zero of F.

## Optional Arguments

ERRABS - First stopping criterion. (Input)
A zero, $B$, is accepted if $\operatorname{ABS}(F(B))$ is less than or equal to ERRABS. ERRABS may be set to zero.
Default: ERRABS $=1 . e-4$ for single precision and 1.d- 8 for double precision.
ERRREL - Second stopping criterion is the relative error. (Input)
A zero is accepted if the change between two successive approximations to this zero is within ERRREL.
Default: ERRREL $=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.

MAXFN - On input, MAXFN specifies an upper bound on the number of function evaluations required for convergence. (Input/Output)
On output, MAXFN will contain the actual number of function evaluations used.
Default: MAXFN $=100$.

## FORTRAN 90 Interface

Generic: CALL ZBREN (F, A, B [,...])
Specific: The specific interface names are $S_{\text {_ ZBREN }}$ and $D_{-}$ZBREN.

## FORTRAN 77 Interface

Single: CALL ZBREN (F, ERRABS, ERRREL, A, B, MAXFN)
Double: The double precision name is DZBREN.

## Description

The algorithm used by ZBREN is a combination of linear interpolation, inverse quadratic interpolation, and bisection. Convergence is usually superlinear and is never much slower than the rate for the bisection method. See Brent (1971) for a more detailed account of this algorithm.

## Comments

1. Informational error

$$
\begin{array}{ccc}
\text { Type } & \text { Code } \\
4 & 1 & \text { Failure to converge in MAXFN function evaluations. }
\end{array}
$$

2. On exit from ZBREN without any error message, A and B satisfy the following:

$$
\begin{aligned}
& \mathrm{F}(\mathrm{~A}) \mathrm{F}(\mathrm{~B}) \leq 0.0 \\
& |\mathrm{~F}(\mathrm{~B})| \leq|\mathrm{F}(\mathrm{~A})| \text {, and } \\
& \text { either }|\mathrm{F}(\mathrm{~B})| \leq \text { ERRABS or } \\
& |\mathrm{A}-\mathrm{B}| \leq \max (|\mathrm{B}|, 0.1) * \text { ERRREL. }
\end{aligned}
$$

The presence of 0.1 in the stopping criterion causes leading zeros to the right of the decimal point to be counted as significant digits. Scaling may be required in order to accurately determine a zero of small magnitude.
3. ZBREN is guaranteed to convergence within K function evaluations, where
$K=(\ln ((B-A) / D)+1.0)^{2}$, and

$$
\left(\mathrm{D}=\min _{x \in(\mathrm{~A}, \mathrm{~B})}(\max (|\mathrm{x}|, 0.1) \star \operatorname{ERRREL})\right)
$$

This is an upper bound on the number of evaluations. Rarely does the actual number of evaluations used by ZBREN exceed

```
                                    \sqrt{}{K}
D can be computed as follows:
P = AMAX1 (0.1, AMIN1 (|A|, |B|))
IF((A - 0.1) * (B - 0.1) < 0.0) P = 0.1,
D = P * ERRREL
```


## Example

This example finds a zero of the function

$$
f(x)=x^{2}+x-2
$$

in the interval $(-10.0,0.0)$.

```
USE ZBREN_INT
USE UMACH_INT
IMPLICIT NONE
REAL ERRABS, ERRREL
INTEGER NOUT, MAXFN
REAL A, B, F
EXTERNAL F
Set values of A, B, ERRABS,
ERRREL, MAXFN
A}=-10.
B = 0.0
ERRABS = 0.0
ERRREL = 0.001
MAXFN = 100
CALL UMACH (2, NOUT)
CALL ZBREN (F, A, B, ERRABS=ERRABS, ERRREL=ERRREL, MAXFN=MAXFN)
WRITE (NOUT,99999) B, MAXFN
99999 FORMAT (' The best approximation to the zero of F is equal to', &
        F5.1, '.', /, ' The number of function evaluations', &
        ' required was ', I2, '.', //)
    END
    REAL FUNCTION F (X)
    REAL X
F = X**2 + X - 2.0
RETURN
END
```

!
$!$
$!$

## Output

The best approximation to the zero of $F$ is equal to -2.0 . The number of function evaluations required was 12.

## ZREAL

Finds the real zeros of a real function using Müller's method.

## Required Arguments

$\boldsymbol{F}$ - User-supplied FUNCTION to compute the value of the function of which a zero will be found. The form is $F(X)$, where

X - The point at which the function is evaluated. (Input) $X$ should not be changed by $F$.

F - The computed function value at the point X . (Output) F must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - A vector of length nroot. (Output)
X contains the computed zeros.

## Optional Arguments

ERRABS - First stopping criterion. (Input)
A zero $X(I)$ is accepted if $\operatorname{ABS}(F(X(I))$. LT. ERRABS.
Default: ERRABS $=1 . e-4$ for single precision and 1.d-8 for double precision.
ERRREL - Second stopping criterion is the relative error. (Input)
A zero $\mathrm{X}(\mathrm{I})$ is accepted if the relative change of two successive approximations to $\mathrm{X}(\mathrm{I})$ is less than errret.
Default: ERRREL = 1.e-4 for single precision and 1.d-8 for double precision.
EPS - See EtA. (Input)
Default: EPS $=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.
$\boldsymbol{E T A}$ - Spread criteria for multiple zeros. (Input)
If the zero $X(I)$ has been computed and $A B S(X(I)-X(J))$.LT.EPS, where $X(J)$ is a previously computed zero, then the computation is restarted with a guess equal to $X(I)+$ ETA.
Default: ETA $=.01$.
NROOT - The number of zeros to be found by ZREAL. (Input)
Default: NROOT $=1$.

ITMAX - The maximum allowable number of iterations per zero. (Input)
Default: ITMAX $=100$.

XGUESS - A vector of length NROOT. (Input)
XGUESS contains the initial guesses for the zeros.
Default: XGUESS $=0.0$.

INFO - An integer vector of length nROOT. (Output)
INFO(J) contains the number of iterations used in finding the $J$-th zero when convergence was achieved. If convergence was not obtained in ITMAX iterations, INFO(J) will be greater than ITMAX.

## FORTRAN 90 Interface

Generic: CALL ZREAL (F, X [, ...])
Specific: The specific interface names are S_ZREAL and D_ZREAL.

## FORTRAN 77 Interface

Single: CALL ZREAL (F, ERRABS, ERRREL, EPS, ETA, NROOT, ITMAX, XGUESS, X, INFO)

Double: $\quad$ The double precision name is DZREAL.

## Description

Routine ZREAL computes $n$ real zeros of a real function $f$. Given a user-supplied function $f(x)$ and an $n$-vector of initial guesses $x_{1}, x_{2}, \ldots, x_{n}$, the routine uses Müller's method to locate $n$ real zeros of $f$, that is, $n$ real values of $x$ for which $f(x)=0$. The routine has two convergence criteria: the first requires that

$$
\left|f\left(x_{i}^{m}\right)\right|
$$

be less than ERRABS; the second requires that the relative change of any two successive approximations to an $x_{i}$ be less than ERRREL. Here,

$$
x_{i}^{m}
$$

is the $m$-th approximation to $x_{i}$. Let ERRABS be $\varepsilon_{1}$, and ERRREL be $\varepsilon_{2}$. The criteria may be stated mathematically as follows:
Criterion 1:

$$
\left|f\left(x_{i}^{m}\right)\right|<\varepsilon_{1}
$$

Criterion 2:

$$
\left|\frac{x_{i}^{m+1}-x_{i}^{m}}{x_{i}^{m}}\right|<\varepsilon_{2}
$$

"Convergence" is the satisfaction of either criterion.

## Comments

1. Informational error

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Failure to converge within ITMAX iterations for at least one of the <br> NROOT roots. |

2. Routine ZREAL always returns the last approximation for zero $J$ in $X(J)$. If the convergence criterion is satisfied, then $\operatorname{INFO}(J)$ is less than or equal to ITMAX. If the convergence criterion is not satisfied, then INFO( J ) is set to ITMAX +1 .
3. The routine ZREAL assumes that there exist NROOT distinct real zeros for the function $F$ and that they can be reached from the initial guesses supplied. The routine is designed so that convergence to any single zero cannot be obtained from two different initial guesses.
4. Scaling the X vector in the function F may be required, if any of the zeros are known to be less than one.

## Example

This example finds the real zeros of the second-degree polynomial

$$
f(x)=x^{2}+2 x-6
$$

with the initial guess $(4.6,-193.3)$.

```
USE ZREAL INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER NROOT
REAL EPS, ERRABS, ERRREL
PARAMETER (NROOT=2)
INTEGER INFO(NROOT)
REAL F, X(NROOT), XGUESS (NROOT)
EXTERNAL F
Set values of initial guess
XGUESS = ( 4.6 -193.3)
DATA XGUESS/4.6, -193.3/
EPS = 1.0E-5
ERRABS = 1.0E-5
ERRREL = 1.0E-5
Find the zeros
    CALL ZREAL (F, X, errabs=errabs, errrel=errrel, eps=eps, &
        nroot=nroot, xguess=xguess)
    CALL WRRRN ('The zeros are', X, 1, NROOT, 1)
    END
```

!
!

```
REAL FUNCTION F (X)
REAL X
!
F = X*X + 2.0*X - 6.0
RETURN
END
```


## Output

```
The zeros are
    1 2
1.646 -3.646
```


## NEQNF

Solves a system of nonlinear equations using a modified Powell hybrid algorithm and a finitedifference approximation to the Jacobian.

## Required Arguments

$\boldsymbol{F C N}$ - User-supplied SUBROUTINE to evaluate the system of equations to be solved. The usage is CALL FCN (X, F,N), where

X - The point at which the functions are evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function values at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
$N$ - Length of X and F. (Input)
$\boldsymbol{X}$ - A vector of length n. (Output)
$X$ contains the best estimate of the root found by NEQNF.

## Optional Arguments

ERRREL - Stopping criterion. (Input)
The root is accepted if the relative error between two successive approximations to this root is less than ERRREL.
Default: ERRREL = 1.e-4 for single precision and 1.d-8 for double precision.
$N$ - The number of equations to be solved and the number of unknowns. (Input) Default: $\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)$.

ITMAX - The maximum allowable number of iterations. (Input)
The maximum number of calls to FCN is ITMAX * ( $N+1$ ). Suggested value ITMAX $=200$.
Default: $\operatorname{ITMAX}=200$.

XGUESS - A vector of length N. (Input)
XGUESS contains the initial estimate of the root.
Default: XGUESS $=0.0$.

FNORM - A scalar that has the value $\mathrm{F}(1)^{2}+\ldots+\mathrm{F}(\mathrm{N})^{2}$ at the point X . (Output)

## FORTRAN 90 Interface

Generic: CALL NEQNF (FCN, X [,...])
Specific: The specific interface names are S_NEQNF and D_NEQNF.

## FORTRAN 77 Interface

Single: CALL NEQNF (FCN, ERRREL, N, ITMAX, XGUESS, X, FNORM)
Double: The double precision name is DNEQNF.

## Description

Routine NEQNF is based on the MINPACK subroutine HYBRD1, which uses a modification of M.J.D. Powell's hybrid algorithm. This algorithm is a variation of Newton's method, which uses a finite-difference approximation to the Jacobian and takes precautions to avoid large step sizes or increasing residuals. For further description, see More et al. (1980).

Since a finite-difference method is used to estimate the Jacobian, for single precision calculation, the Jacobian may be so incorrect that the algorithm terminates far from a root. In such cases, high precision arithmetic is recommended. Also, whenever the exact Jacobian can be easily provided, IMSL routine NEQNJ should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of N2QNF / DN 2 QNF. The reference is:

CALL N2QNF (FCN, ERRREL, N, ITMAX, XGUESS, X, FNORM, FVEC, FJAC, R, QTF, WK)
The additional arguments are as follows:
$\boldsymbol{F V E C}$ - A vector of length N. FVEC contains the functions evaluated at the point X.
$\boldsymbol{F} \boldsymbol{J} \boldsymbol{A C}$ - An N by N matrix. FJAC contains the orthogonal matrix $Q$ produced by the $Q R$ factorization of the final approximate Jacobian.
$\boldsymbol{R}$ - A vector of length $\mathrm{N}^{*}(\mathrm{~N}+1) / 2$. R contains the upper triangular matrix produced by the $Q R$ factorization of the final approximate Jacobian. $R$ is stored row-wise.
$\boldsymbol{Q T F}$ - A vector of length N. QTF contains the vector TRANS(Q) * FVEC.
$W K-A$ work vector of length $5 * \mathrm{~N}$.
2. Informational errors


41 The number of calls to FCN has exceeded ITMAX * $(\mathrm{N}+1)$. A new initial guess may be tried.
42 ERRREL is too small. No further improvement in the approximate solution is possible.
43 The iteration has not made good progress. A new initial guess may be tried.

## Example

The following $3 \times 3$ system of nonlinear equations

$$
\begin{aligned}
& f_{1}(x)=x_{1}+e^{x_{1}-1}+\left(x_{2}+x_{3}\right)^{2}-27=0 \\
& f_{2}(x)=e^{x_{2}-2} / x_{1}+x_{3}^{2}-10=0 \\
& f_{3}(x)=x_{3}+\sin \left(x_{2}-2\right)+x_{2}^{2}-7=0
\end{aligned}
$$

is solved with the initial guess (4.0, 4.0, 4.0).
USE NEQNF_INT
USE UMACH_INT
IMPLICIT NONE
! $\quad$ Declare variables
INTEGER N
PARAMETER ( $\mathrm{N}=3$ )
!
INTEGER K, NOUT
REAL FNORM, X(N), XGUESS (N)
EXTERNAL FCN
Set values of initial guess XGUESS $=\left(\begin{array}{lll}4.0 & 4.0 & 4.0\end{array}\right)$

DATA XGUESS/4.0, 4.0, 4.0/
$!$
CALL UMACH (2, NOUT)
CALI NEQNF (FCN $X$ xquess= Find the solution
WRITE (NOUT, 99999) ( $\mathrm{X}(\mathrm{K}), \mathrm{K}=1, \mathrm{~N}$ ), FNORM
99999 FORMAT (' The solution to the system is', /, ' $\mathrm{X}=(', 3 F 5.1, ~ \&$ ')', /, ' with FNORM =', F5.4, //)
!
END
! SUBROUTINE FCN (X, F, N)
INTEGER N
REAL $\quad \mathrm{X}(\mathrm{N}), \mathrm{F}(\mathrm{N})$
!

```
    REAL EXP, SIN
    INTRINSIC EXP, SIN
!
    F(1) = X(1) + EXP(X(1)-1.0) + (X(2)+X(3))*(X(2)+X(3)) - 27.0
    F(2) = EXP(X(2)-2.0)/X(1) + X(3)*X(3) - 10.0
    F(3) = X(3) + SIN (X (2)-2.0) + X(2)*X(2) - 7.0
    RETURN
    END
```


## Output

```
The solution to the system is
X = ( 1.0 2.0 3.0)
with FNORM =.0000
```


## NEQNJ

Solves a system of nonlinear equations using a modified Powell hybrid algorithm with a usersupplied Jacobian.

## Required Arguments

FCN - User-supplied SUBROUTINE to evaluate the system of equations to be solved. The usage is CALL FCN (X, F,N), where
x - The point at which the functions are evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function values at the point X . (Output)
N - Length of X, F. (Input)

FCN must be declared EXTERNAL in the calling program.

LSJAC - User-supplied SUbroutine to evaluate the Jacobian at a point X . The usage is CALL LSJAC ( $\mathrm{N}, \mathrm{X}, \mathrm{FJAC}$ ), where

N - Length of X . (Input)
x - The point at which the function is evaluated. (Input) x should not be changed by LSJAC.

FJAC - The computed $n$ by n Jacobian at the point X . (Output)
LSJAC must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - A vector of length n. (Output)
$X$ contains the best estimate of the root found by NEQNJ.

## Optional Arguments

ERRREL - Stopping criterion. (Input)
The root is accepted if the relative error between two successive approximations to this root is less than ERRREL.
Default: ERRREL $=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.
$N$ - The number of equations to be solved and the number of unknowns. (Input) Default: $\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)$.

ITMAX - The maximum allowable number of iterations. (Input)
Suggested value $=200$.
Default: ITMAX $=200$.

XGUESS - A vector of length N. (Input)
XGUESS contains the initial estimate of the root.
Default: XGUESS $=0.0$.
FNORM - A scalar that has the value $\mathrm{F}(1)^{2}+\ldots+\mathrm{F}(\mathrm{N})^{2}$ at the point X . (Output)

## FORTRAN 90 Interface

Generic: CALL NEQNJ (FCN, LSJAC, X [,...])
Specific: The specific interface names are S_NEQNJ and D_NEQNJ.

## FORTRAN 77 Interface

Single: CALL NEQNJ (FCN, LSJAC, ERRREL, N, ITMAX, XGUESS, X, FNORM)

Double: $\quad$ The double precision name is DNEQNJ.

## Description

Routine NEQNJ is based on the MINPACK subroutine HYBRDJ, which uses a modification of M.J.D. Powell's hybrid algorithm. This algorithm is a variation of Newton's method, which takes precautions to avoid large step sizes or increasing residuals. For further description, see More et al. (1980).

## Comments

1. Workspace may be explicitly provided, if desired, by use of 2 QNJ/DN2QNJ. The reference is:

CALL N2QNJ (FCN, LSJAC, ERRREL, N, ITMAX, XGUESS, X, FNORM, FVEC, FJAC, R, QTF, WK)

The additional arguments are as follows:
$\boldsymbol{F V E C}$ - A vector of length N. FVEC contains the functions evaluated at the point X.
$\boldsymbol{F} \boldsymbol{J} \boldsymbol{A C}$ - An N by N matrix. FJAC contains the orthogonal matrix Q produced by the $Q R$ factorization of the final approximate Jacobian.
$\boldsymbol{R}$ - A vector of length $N *(N+1) / 2$. R contains the upper triangular matrix produced by the $Q R$ factorization of the final approximate Jacobian. $R$ is stored row-wise.

QTF - A vector of length N. QTF contains the vector TRANS (Q) * FVEC.
WK - A work vector of length 5 * N .
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 1 | The number of calls to FCN has exceeded ITMAX. A new initial guess <br> may be tried. |
| 4 | 2 | ERRREL is too small. No further improvement in the approximate <br> solution is possible. |
| 4 | 3 | The iteration has not made good progress. A new initial guess may <br> be tried. |

## Example

The following $3 \times 3$ system of nonlinear equations

$$
\begin{aligned}
& f_{1}(x)=x_{1}+e^{x_{1}-1}+\left(x_{2}+x_{3}\right)^{2}-27=0 \\
& f_{2}(x)=e^{x_{2}-2} / x_{1}+x_{3}^{2}-10=0 \\
& f_{3}(x)=x_{3}+\sin \left(x_{2}-2\right)+x_{2}^{2}-7=0
\end{aligned}
$$

is solved with the initial guess (4.0, 4.0, 4.0).

```
USE NEQNJ_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=3)
INTEGER K, NOUT
REAL FNORM, X(N), XGUESS(N)
EXTERNAL FCN, LSJAC
    set values of initial guess
                                    XGUESS = ( 4.0 4.0 4.0 )
DATA XGUESS/4.0, 4.0, 4.0/
CALL UMACH (2, NOUT)
```

!

```
! Find the solution
    CALL NEQNJ (FCN, LSJAC, X, XGUESS=XGUESS, FNORM=FNORM)
                            Output
    WRITE (NOUT,99999) (X (K), K=1,N), FNORM
99999 FORMAT (' The roots found are', /, ' X = (', 3F5.1, &
        ')', /, ' with FNORM = ',F5.4, //)
!
    END
! User-supplied subroutine
    SUBROUTINE FCN (X, F, N)
    INTEGER N
    REAL X(N), F(N)
    REAL EXP, SIN
    INTRINSIC EXP, SIN
!
    F(1) = X(1) + EXP(X(1)-1.0) + (X(2)+X(3))*(X(2)+X(3)) - 27.0
    F(2) = EXP(X(2)-2.0)/X(1) + X(3)*X(3) - 10.0
    F(3) = X(3) + SIN(X(2)-2.0) + X(2)*X(2) - 7.0
    RETURN
    END
    User-supplied subroutine to
    compute Jacobian
    SUBROUTINE LSJAC (N, X, FJAC)
    INTEGER N
    REAL X(N), FJAC (N,N)
    REAL COS, EXP
    INTRINSIC COS, EXP
!
    FJAC(1,1)=1.0 + EXP(X(1)-1.0)
    FJAC(1,2) = 2.0*(X(2)+X(3))
    FJAC}(1,3)=2.0*(X(2)+X(3)
    FJAC (2,1) = - EXP (X (2) - 2.0)* (1.0/X(1)**2)
    FJAC (2,2) = EXP (X (2)-2.0)* (1.0/X(1))
    FJAC (2,3) = 2.0*X(3)
    FJAC (3,1) = 0.0
    FJAC (3,2) = COS (X (2)-2.0) + 2.0*X(2)
    FJAC}(3,3)=1.
    RETURN
    END
```


## Output

The roots found are
$X=\left(\begin{array}{lll}1.0 & 2.0 & 3.0\end{array}\right)$
with $\mathrm{FNORM}=.0000$
NEQBF
Solves a system of nonlinear equations using factored secant update with a finite-difference approximation to the Jacobian.

## Required Arguments

FCN - User-supplied SUBROUTINE to evaluate the system of equations to be solved. The usage is CALL $\operatorname{FCN}(\mathrm{N}, \mathrm{X}, \mathrm{F})$, where

N - Length of X and F . (Input)
x - The point at which the functions are evaluated. (Input)
X should not be changed by FCN.
F - The computed function values at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{x}, 1)$.

XGUESS - Vector of length N containing initial guess of the root. (Input)
Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the distance between two points. In the absence of other information, set all entries to 1.0. If internal scaling is desired for XSCALE, set IPARAM (6) to 1.
Default: $\mathrm{XSCALE}=1.0$.

FSCALE - Vector of length N containing the diagonal scaling matrix for the functions.
(Input)
FSCALE is used mainly in scaling the function residuals. In the absence of other information, set all entries to 1.0 .
Default: $\operatorname{FSCALE}=1.0$.

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM (1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: IPARAM $=0$.

RPARAM - Parameter vector of length 5. (Input/Output)
See Comment 4.
FVEC - Vector of length $N$ containing the values of the functions at the approximate solution. (Output)

## FORTRAN 90 Interface

Generic: CALL NEQBF (FCN, X [,...])

Specific: The specific interface names are S_NEQBF and D_NEQBF.

## FORTRAN 77 Interface

```
Single: CALL NEQBF (FCN, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC)
```

Double: The double precision name is DNEQBF.

## Description

Routine $N E Q B F$ uses a secant algorithm to solve a system of nonlinear equations, i.e.,

$$
F(x)=0
$$

where $F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{n}$, and $x \in \mathbf{R}^{n}$.
From a current point, the algorithm uses a double dogleg method to solve the following subproblem approximately:

$$
\min _{s \in \mathbf{R}^{\mathbf{R}}}\left\|F\left(x_{c}\right)+J\left(x_{c}\right) s\right\|_{2}
$$

$$
\text { subject to }\|s\|_{2} \leq \delta_{c}
$$

to get a direction $s_{c}$, where $F\left(x_{c}\right)$ and $J\left(x_{c}\right)$ are the function values and the approximate Jacobian respectively evaluated at the current point $x_{c}$. Then, the function values at the point $x_{n}=x_{c}+s_{c}$ are evaluated and used to decide whether the new point $x_{n}$ should be accepted.

When the point $x_{n}$ is rejected, this routine reduces the trust region $\delta_{c}$ and goes back to solve the subproblem again. This procedure is repeated until a better point is found.

The algorithm terminates if the new point satisfies the stopping criterion. Otherwise, $\delta_{c}$ is adjusted, and the approximate Jacobian is updated by Broyden's formula,

$$
J_{n}=J_{c}+\frac{\left(y-J_{c} s_{c}\right) s_{c}^{T}}{s_{c}^{T} s_{c}}
$$

where $J_{n}=J\left(x_{n}\right), J_{c}=J\left(x_{c}\right)$, and $y=F\left(x_{n}\right)-F\left(x_{c}\right)$. The algorithm then continues using the new point as the current point, i.e. $x_{c} \leftarrow x_{n}$.

For more details, see Dennis and Schnabel (1983, Chapter 8).
Since a finite-difference method is used to estimate the initial Jacobian, for single precision calculation, the Jacobian may be so incorrect that the algorithm terminates far from a root. In such cases, high precision arithmetic is recommended. Also, whenever the exact Jacobian can be easily provided, IMSL routine NEQBJ should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of N 2 QBF / DN 2 QBF. The reference is:

CALL N2QBF (FCN, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, WK, LWK)

The additional arguments are as follows:
WK - A work vector of length LWK. On output WK contains the following information:
The third N locations contain the last step taken.

The fourth N locations contain the last Newton step.
The final $N^{2}$ locations contain an estimate of the Jacobian at the solution.
$\boldsymbol{L W K}$ - Length of wK, which must be at least 2 * $\mathrm{N}^{2}+11$ * N. (Input)
2. Informational errors

Type Code
31 The last global step failed to decrease the 2-norm of $F(X)$ sufficiently; either the current point is close to a root of $\mathrm{F}(\mathrm{X})$ and no more accuracy is possible, or the secant approximation to the Jacobian is inaccurate, or the step tolerance is too large.
33 The scaled distance between the last two steps is less than the step tolerance; the current point is probably an approximate root of $\mathrm{F}(\mathrm{X})$ (unless STEPTL is too large).
34 Maximum number of iterations exceeded.
35 Maximum number of function evaluations exceeded.
37 Five consecutive steps of length STEPMX have been taken; either the 2-norm of $F(X)$ asymptotes from above to a finite value in some direction or the maximum allowable step size STEPMX is too small.
3. The stopping criterion for NEQBF occurs when the scaled norm of the functions is less than the scaled function tolerance (RPARAM(1)).
4. If the default parameters are desired for $\operatorname{NEQBF}$, then set IPARAM(1) to zero and call routine NEQBF. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling NEQBF:

CALL N4QBJ (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to N4QBJ will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above.

The following is a list of the parameters and the default values:

IPARAM - Integer vector of length 6.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: not used in NEQBF.
$\operatorname{IPARAM}(6)=$ Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values of XSCALE are set internally.
Default: 0.
RPARAM - Real vector of length 5.
$\operatorname{RPARAM}(1)=$ Scaled function tolerance.
The scaled norm of the functions is computed as

$$
\max _{i}\left(\left|f_{i}\right| * f s_{i}\right)
$$

where $f_{i}$ is the $i$-th component of the function vector F , and $f s_{i}$ is the $i$-th component of FSCALE.
Default:

$$
\sqrt{\varepsilon}
$$

where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $($ STEPTL $)$
The scaled norm of the step between two points $x$ and $y$ is computed as

$$
\max _{i}\left\{\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}\right\}
$$

where $s_{i}$ is the $i$-th component of XSCALE.
Default: $\varepsilon^{2 / 3}$, where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ False convergence tolerance.
Default: not used in NEQBF.
$\operatorname{RPARAM}(4)=$ Maximum allowable step size. (STEPMX)

Default: $\quad 1000 * \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$, where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=$ XSCALE, and $t=$ XGUESS.
$\operatorname{RPARAM}(5)=$ Size of initial trust region.
Default: based on the initial scaled Cauchy step.
If double precision is desired, then DN $4 Q B J$ is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The following $3 \times 3$ system of nonlinear equations:

$$
\begin{aligned}
& f_{1}(x)=x_{1}+e^{x_{1}-1}+\left(x_{2}+x_{3}\right)^{2}-27=0 \\
& f_{2}(x)=e^{x_{2}-2} / x_{1}+x_{3}^{2}-10=0 \\
& f_{3}(x)=x_{3}+\sin \left(x_{2}-2\right)+x_{2}^{2}-7=0
\end{aligned}
$$

is solved with the initial guess (4.0, 4.0, 4.0).

```
USE NEQBF INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=3)
INTEGER K, NOUT
REAL X(N), XGUESS (N)
EXTERNAL FCN
DATA XGUESS/3*4.0/
CALL NEQBF (FCN, X, XGUESS=XGUESS)
Output
```

```
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) (X (K),K=1,N)
99999 FORMAT (' The solution to the system is', /, ' X = (', 3F8.3, &
    ')')
!
    END
! User-defined subroutine
    SUBROUTINE FCN (N, X, F)
    INTEGER N
    REAL X(N), F(N)
!
    REAL EXP, SIN
    INTRINSIC EXP, SIN
!
    F(1) = X(1) + EXP(X(1)-1.0) + (X(2)+X(3))*(X(2)+X(3)) - 27.0
    F(2) = EXP(X(2)-2.0)/X(1) + X(3)*X(3) - 10.0
    F(3) = X(3) + SIN (X(2)-2.0) + X(2)*X(2) - 7.0
    RETURN
    END
```


## Output

The solution to the system is
$X=\left(\begin{array}{lll} \\ X & 1.000 & 2.000\end{array}\right.$

## NEQBJ

Solves a system of nonlinear equations using factored secant update with a user-supplied Jacobian.

## Required Arguments

FCN - User-supplied SUBROUTINE to evaluate the system of equations to be solved. The usage is CALL FCN ( $\mathrm{N}, \mathrm{X}, \mathrm{F}$ ), where

N - Length of X and F . (Input)
X - The point at which the functions are evaluated. (Input)
$X$ should not be changed by FCN.
F - The computed function values at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
JAC - User-supplied subroutine to evaluate the Jacobian at a point x . The usage is CALL JAC (N, X, FJAC, LDFJAC), where

N - Length of X . (Input)
$X-$ Vector of length $N$ at which point the Jacobian is evaluated. (Input)
$x$ should not be changed by JAC.
FJAC - The computed n by N Jacobian at the point X. (Output)
LDF JAC - Leading dimension of FJAC. (Input)
JAC must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{size}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing initial guess of the root. (Input)
Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the distance between two points. In the absence of other information, set all entries to 1.0. If internal scaling is desired for XSCALE, set $\operatorname{IPARAM}(6)$ to 1.
Default: $\mathrm{XSCALE}=1.0$.
FSCALE - Vector of length N containing the diagonal scaling matrix for the functions. (Input)
FSCALE is used mainly in scaling the function residuals. In the absence of other information, set all entries to 1.0 .
Default: FSCALE = 1.0.
IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM (1) to zero for default values of IPARAM and RPARAM.
See Comment 4.
Default: IPARAM $=0$.
RPARAM - Parameter vector of length 5. (Input/Output)
See Comment 4.
FVEC - Vector of length $N$ containing the values of the functions at the approximate solution. (Output)

## FORTRAN 90 Interface

Generic: CALL NEQBJ (FCN, JAC, X [, ...])
Specific: The specific interface names are S_NEQBJ and D_NEQBJ.

## FORTRAN 77 Interface

Single: CALL NEQBJ (FCN, JAC, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC)

Double: The double precision name is DNEQBJ.

## Description

Routine $N E Q B J$ uses a secant algorithm to solve a system of nonlinear equations, i. e.,

$$
F(x)=0
$$

where $F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{n}$, and $x \in \mathbf{R}^{n}$.

From a current point, the algorithm uses a double dogleg method to solve the following subproblem approximately:

$$
\begin{gathered}
\min _{s \in \mathbf{R}^{n}}\left\|F\left(x_{c}\right)+J\left(x_{c}\right) s\right\|_{2} \\
\text { subject to }\|S\|_{2} \leq \delta_{c}
\end{gathered}
$$

to get a direction $s_{c}$, where $F\left(x_{c}\right)$ and $J\left(x_{c}\right)$ are the function values and the approximate Jacobian respectively evaluated at the current point $x_{c}$. Then, the function values at the point $x_{n}=x_{c}+s_{c}$ are evaluated and used to decide whether the new point $x_{n}$ should be accepted.
When the point $x_{n}$ is rejected, this routine reduces the trust region $\delta_{c}$ and goes back to solve the subproblem again. This procedure is repeated until a better point is found.

The algorithm terminates if the new point satisfies the stopping criterion. Otherwise, $\delta_{c}$ is adjusted, and the approximate Jacobian is updated by Broyden's formula,

$$
J_{n}=J_{c}+\frac{\left(y-J_{c} s_{c}\right) s_{c}^{T}}{s_{c}^{T} s_{c}}
$$

where $J_{n}=J\left(x_{n}\right), J_{c}=J\left(x_{c}\right)$, and $y=F\left(x_{n}\right)-F\left(x_{c}\right)$. The algorithm then continues using the new point as the current point, i.e. $x_{c} \leftarrow x_{n}$.

For more details, see Dennis and Schnabel (1983, Chapter 8).

## Comments

1. Workspace may be explicitly provided, if desired, by use of N2QBJ/DN2QBJ. The reference is:

CALL N2QBJ (FCN, JAC, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, WK, LWK)

The additional arguments are as follows:
WK - A work vector of length LWK. On output wK contains the following information: The third N locations contain the last step taken. The fourth N locations contain the last Newton step. The final $\mathrm{N}^{2}$ locations contain an estimate of the Jacobian at the solution.
$\boldsymbol{L W K}$ - Length of $w K$, which must be at least $2 * \mathrm{~N}^{2}+11 * N$. (Input)
2. Informational errors

| Type | Code |
| :---: | :--- |
| 3 | 1 |
|  | The last global step failed to decrease the 2-norm of $F(X)$ sufficiently; <br> either the current point is close to a root of $F(X)$ and no more <br> accuracy is possible, or the secant approximation to the Jacobian is <br> inaccurate, or the step tolerance is too large. |

3 The scaled distance between the last two steps is less than the step tolerance; the current point is probably an approximate root of $\mathrm{F}(\mathrm{X})$ (unless STEPTL is too large).
4 Maximum number of iterations exceeded.
5 Maximum number of function evaluations exceeded.
7 Five consecutive steps of length STEPMX have been taken; either the 2-norm of $F(X)$ asymptotes from above to a finite value in some direction or the maximum allowable stepsize STEPMX is too small.
3. The stopping criterion for $N E Q B J$ occurs when the scaled norm of the functions is less than the scaled function tolerance (RPARAM(1)).
4. If the default parameters are desired for $N E Q B J$, then set IPARAM(1) to zero and call routine NEQBJ. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling NEQBJ:

CALL N4QBJ (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to N4QBJ will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 6.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: not used in NEQBJ.

IPARAM(6) = Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values of XSCALE are set internally.
Default: 0.
RPARAM - Real vector of length 5.
$\operatorname{RPARAM}(1)=$ Scaled function tolerance.
The scaled norm of the functions is computed as

$$
\max _{i}\left(\left|f_{i}\right| * f s_{i}\right)
$$

where $f_{i}$ is the $i$-th component of the function vector F , and $f s_{i}$ is the $i$-th component of FSCALE.
Default:

$$
\sqrt{\varepsilon}
$$

where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. (STEPTL)
The scaled norm of the step between two points $x$ and $y$ is computed as

$$
\max _{i}\left\{\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}\right\}
$$

where $s_{i}$ is the $i$-th component of XSCALE.
Default: $\varepsilon^{2 / 3}$, where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ False convergence tolerance.
Default: not used in NEQBJ.
$\operatorname{RPARAM}(4)=$ Maximum allowable step size. (STEPMX)
Default: $1000 * \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$, where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=$ XSCALE, and $t=$ XGUESS.
$\operatorname{RPARAM}(5)=$ Size of initial trust region.
Default: based on the initial scaled Cauchy step.
If double precision is desired, then $D N 4 Q B J$ is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The following $3 \times 3$ system of nonlinear equations

$$
\begin{aligned}
& f_{1}(x)=x_{1}+e^{x_{1}-1}+\left(x_{2}+x_{3}\right)^{2}-27=0 \\
& f_{2}(x)=e^{x_{2}-2} / x_{1}+x_{3}^{2}-10=0 \\
& f_{3}(x)=x_{3}+\sin \left(x_{2}-2\right)+x_{2}^{2}-7=0
\end{aligned}
$$

is solved with the initial guess (4.0, 4.0, 4.0).

```
USE NEQBJ_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=3)
INTEGER K, NOUT
REAL X(N), XGUESS (N)
EXTERNAL FCN, JAC
Set values of initial guess
XGUESS = ( 4.0 4.0 4.0 )
DATA XGUESS/3*4.0/
CALL NEQBJ (FCN, JAC, X, XGUESS=XGUESS)
                                    Output
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) (X (K),K=1,N)
FORMAT (' The solution to the system is', /, ' X = (', 3F8.3, &
            ')')
END
SUBROUTINE FCN (N, X, F)
INTEGER N
REAL X(N), F(N)
REAL EXP, SIN
INTRINSIC EXP, SIN
F(1) = X(1) + EXP(X(1)-1.0) + (X(2) +X(3))*(X(2)+X(3)) - 27.0
F(2) = EXP(X(2)-2.0)/X(1) + X(3)*X(3) - 10.0
F(3) = X(3) + SIN(X(2)-2.0) + X(2)*X(2) - 7.0
RETURN
END
                                    User-supplied subroutine to
                                    compute Jacobian
SUBROUTINE JAC (N, X, FJAC, LDFJAC)
INTEGER N, LDFJAC
REAL X(N), FJAC(LDFJAC,N)
REAL COS, EXP
INTRINSIC COS, EXP
FJAC(1,1) = 1.0 + EXP (X(1)-1.0)
FJAC (1,2) = 2.0*(X(2)+X(3))
```

```
FJAC(1,3)=2.0*(X(2)+X(3))
FJAC}(2,1)=-\operatorname{EXP}(\textrm{X}(2)-2.0)*(1.0/X(1)**2
FJAC (2,2) = EXP (X (2)-2.0)* (1.0/X(1))
FJAC (2,3) = 2.0*X(3)
FJAC (3,1) = 0.0
FJAC (3,2) = COS (X (2)-2.0) + 2.0*X(2)
FJAC (3,3) = 1.0
RETURN
END
```


## Output

The solution to the system is
$X=\left(\begin{array}{lll}1.000 & 2.000 & 3.000\end{array}\right)$

## Chapter 8: Optimization

## Routines

### 8.1. Unconstrained Minimization

8.1.1 Univariate Function

Using function values only ...................................................UVMIF 1222
Using function and first derivative values ........................... UVMID 1225
Nonsmooth function...........................................................UVMGS 1229
8.1.2 Multivariate Function

Using finite-difference gradient............................................UMINF 1232
Using analytic gradient .......................................................UMING 1237
Using finite-difference Hessian............................................. UMIDH 1243
Using analytic Hessian ........................................................ UMIAH 1249
Using conjugate gradient with finite-difference gradient.....UMCGF 1255
Using conjugate gradient with analytic gradient ................ UMCGG 1259
Nonsmooth function........................................................... UMPOL 1263
$\begin{array}{ll}\text { 8.1.3 Nonlinear Least Squares } \\ & \text { Using finite-difference Jacobian.......................................... UNLSF } 1267\end{array}$
Using analytic Jacobian ........................................................UNLSJ 1273
8.2. Minimization with Simple Bounds

Using finite-difference gradient...........................................BCONF 1279
Using analytic gradient ......................................................BCONG 1286
Using finite-difference Hessian .......................................... BCODH 1293
Using analytic Hessian ..................................................... BCOAH 1299
Nonsmooth Function...........................................................BCPOL 1306
Nonlinear least squares using finite-difference Jacobian .... BCLSF 1310
Nonlinear least squares using analytic Jacobian..................BCLSJ 1317
Nonlinear least squares problem subject to bounds............BCNLS 1324
8.3. Linearly Constrained Minimization

Reads an MPS file containing a linear programming
problem or a quadratic programming problem ............READ_MPS 1333
Deallocates the space allocated for the IMSL
derived type s_MPS....................................................... MPS_FREE
1343
Dense linear programming ......................................... DENSE_LP 1346
Dense linear programming .................................................DLPRS 1351
Sparse linear programming SLPRS ..... 1355
Quadratic programming QPROG ..... 1361
General objective function with finite-difference gradient .LCONF ..... 1364
General objective function with analytic gradient .LCONG ..... 1370
8.4. Nonlinearly Constrained Minimization
Using a sequential equality constrained QP method NNLPF ..... 1377
Using a sequential equality constrained QP method ..... NNLPG ..... 1383
8.5. Service Routines
Central-difference gradient CDGRD ..... 1390
Forward-difference gradient FDGRD ..... 1392
Forward-difference Hessian ..... FDHES ..... 1394
Forward-difference Hessian using analytic gradient ..... GDHES ..... 1397
Forward-difference Jacobian ..... FDJAC ..... 1400
Check user-supplied gradient CHGRD ..... 1390
Check user-supplied Hessian CHHES ..... 1406
Check user-supplied Jacobian CHJAC ..... 1410
Generate starting points GGUES ..... 1414

## Usage Notes

## Unconstrained Minimization

The unconstrained minimization problem can be stated as follows:

$$
\min _{x \in \mathbf{R}^{n}} f(x)
$$

where $f: \mathbf{R}^{n} \rightarrow \mathbf{R}$ is at least continuous. The routines for unconstrained minimization are grouped into three categories: univariate functions (UV***), multivariate functions (UM***), and nonlinear least squares (UNLS*).

For the univariate function routines, it is assumed that the function is unimodal within the specified interval. Otherwise, only a local minimum can be expected. For further discussion on unimodality, see Brent (1973).

A quasi-Newton method is used for the multivariate function routines UMINF and UMING, whereas UMIDH and UMIAH use a modified Newton algorithm. The routines UMCGF and UMCGG make use of a conjugate gradient approach, and UMPOL uses a polytope method. For more details on these algorithms, see the documentation for the corresponding routines.
The nonlinear least squares routines use a modified Levenberg-Marquardt algorithm. If the nonlinear least squares problem is a nonlinear data-fitting problem, then software that is designed to deliver better statistical output may be useful; see IMSL (1991).

These routines are designed to find only a local minimum point. However, a function may have many local minima. It is often possible to obtain a better local solution by trying different initial points and intervals.

High precision arithmetic is recommended for the routines that use only function values. Also it is advised that the derivative-checking routines $\mathrm{CH} * * *$ be used to ensure the accuracy of the usersupplied derivative evaluation subroutines.

## Minimization with Simple Bounds

The minimization with simple bounds problem can be stated as follows:

$$
\begin{gathered}
\min _{x \in \mathbb{R}^{n}} f(x) \\
\text { subject to } \quad l_{i} \leq x_{i} \leq u_{i}, \quad \text { for } i-1,2, \ldots, n
\end{gathered}
$$

where $f: \mathbf{R}^{n} \rightarrow \mathbf{R}$, and all the variables are not necessarily bounded.
The routines BCO ** use the same algorithms as the routines UMI**, and the routines BCLS* are the corresponding routines of UNLS*. The only difference is that an active set strategy is used to ensure that each variable stays within its bounds. The routine BCPOL uses a function comparison method similar to the one used by UMPOL. Convergence for these polytope methods is not guaranteed; therefore, these routines should be used as a last alternative.

## Linearly Constrained Minimization

The linearly constrained minimization problem can be stated as follows:

$$
\begin{gathered}
\min _{x \in \mathrm{R}^{n}} f(x) \\
\text { subject to } A x=b
\end{gathered}
$$

where $f: \mathbf{R}^{n} \rightarrow \mathbf{R}, A$ is an $m \times n$ coefficient matrix, and $b$ is a vector of length $m$. If $f(x)$ is linear, then the problem is a linear programming problem; if $f(x)$ is quadratic, the problem is a quadratic programming problem.
The routine DLPRS uses an active set strategy to solve small- to medium-sized linear programming problems. No sparsity is assumed since the coefficients are stored in full matrix form.

The routine $Q P R O G$ is designed to solve convex quadratic programming problems using a dual quadratic programming algorithm. If the given Hessian is not positive definite, then QPROG modifies it to be positive definite. In this case, output should be interpreted with care.

The routines LCONF and LCONG use an iterative method to solve the linearly constrained problem with a general objective function. For a detailed description of the algorithm, see Powell (1988, 1989).

## Nonlinearly Constrained Minimization

The nonlinearly constrained minimization problem can be stated as follows:

$$
\min _{x \in \mathrm{R}^{n}} f(x)
$$

$$
\begin{aligned}
& \text { subject to } g_{i}(x)=0, \\
& g_{i}(x) \geq 0, \text { for } i-1,2, \ldots, m_{1} \\
& \text { for } i=m_{1}+1, \ldots, m_{1}
\end{aligned}
$$

where $f: \mathbf{R}^{n} \rightarrow \mathbf{R}$ and $g_{i}: \mathbf{R}^{n} \rightarrow \mathbf{R}$, for $i=1,2, \ldots, m$
The routines NNLPF and NNLPG use a sequential equality constrained quadratic programming method. A more complete discussion of this algorithm can be found in the documentation.

## Selection of Routines

The following general guidelines are provided to aid in the selection of the appropriate routine.

## Unconstrained Minimization

1. For the univariate case, use UVMID when the gradient is available, and use UVMIF when it is not. If discontinuities exist, then use UVMGS.
2. For the multivariate case, use UMCG* when storage is a problem, and use UMPOL when the function is nonsmooth. Otherwise, use UMI** depending on the availability of the gradient and the Hessian.
3. For least squares problems, use UnLS $J$ when the Jacobian is available, and use UNLSF when it is not.

## Minimization with Simple Bounds

1. Use BCONF when only function values are available. When first derivatives are available, use either BCONG or BCODH. If first and second derivatives are available, then use BCOAH.
2. For least squares, use BCLSF or BCLSJ depending on the availability of the Jacobian.
3. Use BCPOL for nonsmooth functions that could not be solved satisfactorily by the other routines.

The following charts provide a quick reference to routines in this chapter:



## UVMIF

Finds the minimum point of a smooth function of a single variable using only function evaluations.

## Required Arguments

$\boldsymbol{F}$ - User-supplied function to compute the value of the function to be minimized. The form is $F(X)$, where
X - The point at which the function is evaluated. (Input)
$X$ should not be changed by $F$.
F - The computed function value at the point X . (Output)

F must be declared EXTERNAL in the calling program.
XGUESS - An initial guess of the minimum point of F . (Input)
BOUND - A positive number that limits the amount by which x may be changed from its initial value. (Input)
$\boldsymbol{X}$ - The point at which a minimum value of F is found. (Output)

## Optional Arguments

STEP - An order of magnitude estimate of the required change in X . (Input) Default: $\mathrm{STEP}=1.0$.
$X A C C$ - The required absolute accuracy in the final value of x . (Input)
On a normal return there are points on either side of $x$ within a distance XACC at which $F$ is no less than $F(X)$.
Default: $\mathrm{XACC}=1 . \mathrm{e}-4$.

MAXFN - Maximum number of function evaluations allowed. (Input) Default: MAXFN $=1000$.

## FORTRAN 90 Interface

Generic: CALL UVMIF (F, XGUESS, BOUND, X [,...])

Specific: The specific interface names are S_UVMIF and D_UVMIF.

## FORTRAN 77 Interface

Single: CALL UVMIF (F, XGUESS, STEP, BOUND, XACC, MAXFN, X)
Double: The double precision name is DUVMIF.

## Description

The routine UVMIF uses a safeguarded quadratic interpolation method to find a minimum point of a univariate function. Both the code and the underlying algorithm are based on the routine ZXLSF written by M.J.D. Powell at the University of Cambridge.

The routine UVMIF finds the least value of a univariate function, $f$, that is specified by the function subroutine F . Other required data include an initial estimate of the solution, XGUESS, and a positive number BOUND. Let $x_{0}$ = XGUESS and $b=$ BOUND, then $x$ is restricted to the interval $\left[x_{0}-b, x_{0}+b\right]$. Usually, the algorithm begins the search by moving from $x_{0}$ to $x=x_{0}+s$, where
$s=$ STEP is also provided by the user and may be positive or negative. The first two function evaluations indicate the direction to the minimum point, and the search strides out along this direction until a bracket on a minimum point is found or until $x$ reaches one of the bounds $x_{0} \pm b$. During this stage, the step length increases by a factor of between two and nine per function evaluation; the factor depends on the position of the minimum point that is predicted by quadratic interpolation of the three most recent function values.

When an interval containing a solution has been found, we will have three points, $x_{1}, x_{2}$, and $x_{3}$, with $x_{1}<x_{2}<x_{3}$ and $f\left(x_{2}\right) \leq f\left(x_{1}\right)$ and $f\left(x_{2}\right) \leq f\left(x_{3}\right)$. There are three main ingredients in the technique for choosing the new $x$ from these three points. They are (i) the estimate of the minimum point that is given by quadratic interpolation of the three function values, (ii) a tolerance parameter $\varepsilon$, that depends on the closeness of $f$ to a quadratic, and (iii) whether $x_{2}$ is near the center of the range between $x_{1}$ and $x_{3}$ or is relatively close to an end of this range. In outline, the new value of $x$ is as near as possible to the predicted minimum point, subject to being at least $\varepsilon$ from $x_{2}$, and subject to being in the longer interval between $x_{1}$ and $x_{2}$ or $x_{2}$ and $x_{3}$ when $x_{2}$ is particularly close to $x_{1}$ or $x_{3}$. There is some elaboration, however, when the distance between these points is close to the required accuracy; when the distance is close to the machine precision; or when $\varepsilon$ is relatively large.
The algorithm is intended to provide fast convergence when $f$ has a positive and continuous second derivative at the minimum and to avoid gross inefficiencies in pathological cases, such as

$$
f(x)=x+1.001|x|
$$

The algorithm can make $\varepsilon$ large automatically in the pathological cases. In this case, it is usual for a new value of $x$ to be at the midpoint of the longer interval that is adjacent to the least calculated function value. The midpoint strategy is used frequently when changes to $f$ are dominated by computer rounding errors, which will almost certainly happen if the user requests an accuracy that is less than the square root of the machine precision. In such cases, the routine claims to have achieved the required accuracy if it knows that there is a local minimum point within distance $\delta$ of $x$, where $\delta=$ XACC, even though the rounding errors in $f$ may cause the existence of other local minimum points nearby. This difficulty is inevitable in minimization routines that use only function values, so high precision arithmetic is recommended.

## Comments

Informational errors

| Type | Code |  |
| :---: | :--- | :--- |
| 3 | 1 | Computer rounding errors prevent further refinement of X. |
| 3 | 2 | The final value of X is at a bound. The minimum is probably beyond the |
|  |  | bound. |

## Example

A minimum point of $e^{x}-5 x$ is found.

```
USE UVMIF INT
USE UMACH_INT
```

```
IMPLICIT NONE
! Declare variables
INTEGER MAXFN, NOUT
REAL BOUND, F, FX, STEP, X, XACC, XGUESS
EXTERNAL F
! Initialize variables
XGUESS = 0.0
XACC = 0.001
BOUND = 100.0
STEP = 0.1
MAXFN = 50
!
! Find minimum for F = EXP(X) - 5X
CALL UVMIF (F, XGUESS, BOUND, X, STEP=STEP, XACC=XACC, MAXFN=MAXFN)
FX = F(X)
! Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) X, FX
!
99999 FORMAT (' The minimum is at ', 7X, F7.3, //, ' The function ' &
    , 'value is ', F7.3)
!
    END
    REAL FUNCTION F (X)
    REAL X
    REAL EXP
    INTRINSIC EXP
    F = EXP (X) - 5.0E0*X
!
    RETURN
    END
```


## Output

The minimum is at 1.609
The function value is -3.047

## UVMID

Finds the minimum point of a smooth function of a single variable using both function evaluations and first derivative evaluations.

## Required Arguments

$\boldsymbol{F}$ - User-supplied function to define the function to be minimized. The form is $F(\mathrm{X})$, where $\boldsymbol{X}$ - The point at which the function is to be evaluated. (Input)
$\boldsymbol{F}$ - The computed value of the function at X . (Output)

F must be declared EXTERNAL in the calling program.
$\boldsymbol{G}$ - User-supplied function to compute the derivative of the function. The form is $\mathrm{G}(\mathrm{x})$, where
$\boldsymbol{X}$ - The point at which the derivative is to be computed. (Input)
$\boldsymbol{G}$ - The computed value of the derivative at X. (Output)
G must be declared EXTERNAL in the calling program.
$A$ - A is the lower endpoint of the interval in which the minimum point of F is to be located. (Input)
$\boldsymbol{B}-\mathrm{B}$ is the upper endpoint of the interval in which the minimum point of F is to be located. (Input)
$\boldsymbol{X}$ - The point at which a minimum value of F is found. (Output)

## Optional Arguments

XGUESS - An initial guess of the minimum point of F . (Input)
Default: XGUESS $=(a+b) / 2.0$.

ERRREL - The required relative accuracy in the final value of x . (Input)
This is the first stopping criterion. On a normal return, the solution X is in an interval that contains a local minimum and is less than or equal to $\operatorname{MAX}(1.0, \operatorname{ABS}(\mathrm{X})) *$ ERRREL. When the given ERRREL is less than machine epsilon, $S Q R T$ (machine epsilon) is used as ERRREL.
Default: ERRREL = 1.e-4.

GTOL - The derivative tolerance used to decide if the current point is a local minimum. (Input)
This is the second stopping criterion. X is returned as a solution when GX is less than or equal to GTOL. GTOL should be nonnegative, otherwise zero would be used.
Default: GTOL = 1.e-4.
MAXFN - Maximum number of function evaluations allowed. (Input) Default: $\operatorname{MAXFN}=1000$.
$\boldsymbol{F X}$ - The function value at point X . (Output)
$\boldsymbol{G} \boldsymbol{X}$ - The derivative value at point X . (Output)

## FORTRAN 90 Interface

Generic: CALL UVMID (F, G, A, B, X [, ...])
Specific: The specific interface names are S_UVMID and D_UVMID.

## FORTRAN 77 Interface

Single: CALL UVMID (F, G, XGUESS, ERRREL, GTOL, MAXFN, A, B, X, FX, GX)

Double: The double precision name is DUVMID.

## Description

The routine UVMID uses a descent method with either the secant method or cubic interpolation to find a minimum point of a univariate function. It starts with an initial guess and two endpoints. If any of the three points is a local minimum point and has least function value, the routine terminates with a solution. Otherwise, the point with least function value will be used as the starting point.

From the starting point, say $x_{c}$, the function value $f_{c}=f\left(x_{c}\right)$, the derivative value $g_{c}=g\left(x_{c}\right)$, and a new point $x_{n}$ defined by $x_{n}=x_{c}-g_{c}$ are computed. The function $f_{n}=f\left(x_{n}\right)$, and the derivative $g_{n}=g\left(x_{n}\right)$ are then evaluated. If either $f_{n} \geq f_{c}$ or $g_{n}$ has the opposite sign of $g_{c}$, then there exists a minimum point between $x_{c}$ and $x_{n}$; and an initial interval is obtained. Otherwise, since $x_{c}$ is kept as the point that has lowest function value, an interchange between $x_{n}$ and $x_{c}$ is performed. The secant method is then used to get a new point

$$
x_{s}=x_{c}-g_{c}\left(\frac{g_{n}-g_{c}}{x_{n}-x_{c}}\right)
$$

Let $x_{n} \leftarrow x_{s}$ and repeat this process until an interval containing a minimum is found or one of the convergence criteria is satisfied. The convergence criteria are as follows:

## Criterion 1:

$$
\left|x_{c}-x_{n}\right| \leq \varepsilon_{c}
$$

Criterion 2:

$$
\left|g_{c}\right| \leq \varepsilon_{g}
$$

where $\varepsilon_{c}=\max \left\{1.0,\left|x_{c}\right|\right\} \varepsilon, \varepsilon$ is a relative error tolerance and $\varepsilon_{g}$ is a gradient tolerance.
When convergence is not achieved, a cubic interpolation is performed to obtain a new point. Function and derivative are then evaluated at that point; and accordingly, a smaller interval that contains a minimum point is chosen. A safeguarded method is used to ensure that the interval reduces by at least a fraction of the previous interval. Another cubic interpolation is then performed, and this procedure is repeated until one of the stopping criteria is met.

## Comments

Informational errors
Type Code
31 The final value of X is at the lower bound. The minimum is probably beyond the bound.
32 The final value of X is at the upper bound. The minimum is probably beyond the bound.
43 The maximum number of function evaluations has been exceeded.

## Example

A minimum point of $e^{x}-5 x$ is found.

```
USE UVMID INT
USE UMACH_
    IMPLICIT NONE
! Declare variables
    INTEGER MAXFN, NOUT
    REAL A, B, ERRREL, F, FX, G, GTOL, GX, X, XGUESS, FTOL
    EXTERNAL F, G
    XGUESS = 0.0
    Set ERRREL to zero in order
    to use SQRT(machine epsilon)
    as relative error
    ERRREL = 0.0
    GTOL = 0.0
    A = -10.0
    B = 10.0
    MAXFN = 50
    CALL UVMID (F, G, A, B, X, XGUESS=XGUESS, ERRREL=ERRREL, &
                                    GTOL=FTOL, MAXFN=MAXFN, FX=FX, GX=GX)
                                    Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, FX, GX
99999 FORMAT (' The minimum is at ', 7X, F7.3, //, ' The function ' &
        , 'value is ', F7.3, //, ' The derivative is ', F7.3)
    END
! Real function: F = EXP(X) - 5.0*X
    REAL FUNCTION F (X)
    REAL X
    REAL EXP
    INTRINSIC EXP
    F = EXP(X) - 5.0E0*X
    RETURN
```

$!$
!
!
$!$

```
    END
!
    REAL FUNCTION G (X)
    REAL X
!
    REAL EXP
    INTRINSIC EXP
!
G = EXP (X) - 5.0E0
RETURN
END
```


## Output

```
The minimum is at 1.609
The function value is -3.047
The derivative is -0.001
```


## UVMGS

Finds the minimum point of a nonsmooth function of a single variable.

## Required Arguments

$\boldsymbol{F}$ - User-supplied function to compute the value of the function to be minimized. The form is $F(X)$, where

X - The point at which the function is evaluated. (Input)
$X$ should not be changed by $F$.
F - The computed function value at the point X . (Output)
F must be declared EXTERNAL in the calling program.
A- On input, A is the lower endpoint of the interval in which the minimum of $F$ is to be located. On output, A is the lower endpoint of the interval in which the minimum of F is located. (Input/Output)
$\boldsymbol{B}$ - On input, $B$ is the upper endpoint of the interval in which the minimum of $F$ is to be located. On output, $B$ is the upper endpoint of the interval in which the minimum of $F$ is located. (Input/Output)

XMIN - The approximate minimum point of the function F on the original interval (A, B). (Output)

## Optional Arguments

TOL - The allowable length of the final subinterval containing the minimum point. (Input)
Default: TOL = 1.e-4.

## FORTRAN 90 Interface

Generic: CALL UVMGS (F, A, B, XMIN [,...])

Specific: The specific interface names are S_UVMGS and D_UVMGS.

## FORTRAN 77 Interface

Single: CALL UVMGS (F, A, B, TOL, XMIN)
Double: The double precision name is DUVMGS.

## Description

The routine UVMGS uses the golden section search technique to compute to the desired accuracy the independent variable value that minimizes a unimodal function of one independent variable, where a known finite interval contains the minimum.

Let $\tau=$ TOL. The number of iterations required to compute the minimizing value to accuracy $\tau$ is the greatest integer less than or equal to

$$
\frac{\ln (\tau /(b-a))}{\ln (1-c)}+1
$$

where $a$ and $b$ define the interval and

$$
c=(3-\sqrt{5}) / 2
$$

The first two test points are $v_{1}$ and $v_{2}$ that are defined as

$$
v_{1}=a+c(b-a), \text { and } v_{2}=b-c(b-a)
$$

If $f\left(v_{1}\right)<f\left(v_{2}\right)$, then the minimizing value is in the interval $\left(a, v_{2}\right)$. In this case, $b \leftarrow v_{2}, v_{2} \leftarrow v_{1}$, and $v_{1} \leftarrow a+c(b-a)$. If $f\left(v_{1}\right) \geq f\left(v_{2}\right)$, the minimizing value is in ( $\left.v_{1}, b\right)$. In this case, $a \leftarrow v_{1}$, $v_{1} \leftarrow v_{2}$, and $v_{2} \leftarrow b-c(b-a)$.

The algorithm continues in an analogous manner where only one new test point is computed at each step. This process continues until the desired accuracy $\tau$ is achieved. XMIN is set to the point producing the minimum value for the current iteration.

Mathematically, the algorithm always produces the minimizing value to the desired accuracy; however, numerical problems may be encountered. If $f$ is too flat in part of the region of interest, the function may appear to be constant to the computer in that region. Error code 2 indicates that this problem has occurred. The user may rectify the problem by relaxing the requirement on $\tau$, modifying (scaling, etc.) the form of $f$ or executing the program in a higher precision.

## Comments

1. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | TOL is too small to be satisfied. |
| 4 | 2 | Due to rounding errors F does not appear to be unimodal. |

2. On exit from UVMGS without any error messages, the following conditions hold:
(B-A) $\leq$ TOL.
$\mathrm{A} \leq \mathrm{XMIN}$ and XMIN $\leq \mathrm{B}$
$F(X M I N) \leq F(A)$ and $F(X M I N) \leq F(B)$
3. On exit from UVMGS with error code 2, the following conditions hold:
$\mathrm{A} \leq$ XMIN and XMIN $\leq \mathrm{B}$
$F(X M I N) \geq F(A)$ and $F(X M I N) \geq F(B)$ (only one equality can hold).
Further analysis of the function $F$ is necessary in order to determine whether it is not unimodal in the mathematical sense or whether it appears to be not unimodal to the routine due to rounding errors in which case the $\mathrm{A}, \mathrm{B}$, and XMIN returned may be acceptable.

## Example

A minimum point of $3 x^{2}-2 x+4$ is found.

```
USE UVMGS_INT
USE UMACH_}\mp@subsup{}{}{-}\mathrm{ INT
IMPLICIT NONE
! Specification of variables
INTEGER NOUT
REAL A, B, FCN, FMIN, TOL, XMIN
EXTERNAL FCN
    A = 0.0E0
    B = 5.0E0
    TOL = 1.0E-3
! Minimize FCN
    CALL UVMGS (FCN, A, B, XMIN, TOL=TOL)
    FMIN = FCN(XMIN)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) XMIN, FMIN, A, B
99999 FORMAT (' The minimum is at ', F5.3, //, ' The ', &
    'function value is ', F5.3, //, ' The final ', &
    'interval is (', F6.4, ',', F6.4, ')', /)
    END
! REAL FUNCTION: F = 3*X**2 - 2*X + 4
    REAL FUNCTION FCN (X)
    REAL X
```

$!$
!

```
FCN = 3.0EO*X*X - 2.0EO*X + 4.0EO
RETURN
END
```

!

## Output

```
The minimum is at 0.333
The function value is 3.667
The final interval is (0.3331,0.3340)
```


## UMINF

Minimizes a function of N variables using a quasi-Newton method and a finite-difference gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
x - The point at which the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X}-$ Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing an initial guess of the computed solution. (Input) Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0.
Default: $\mathrm{XSCALE}=1.0$.

FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.

IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: IPARAM $=0$.
RPARAM - Parameter vector of length 7.(Input/Output)
See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMINF (FCN, X [, ...])
Specific: The specific interface names are S_UMINF and D_UMINF.

## FORTRAN 77 Interface

Single: CALL UMINF (FCN, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is Duminf.

## Description

The routine UMINF uses a quasi-Newton method to find the minimum of a function $f(x)$ of $n$ variables. Only function values are required. The problem is stated as follows:

$$
\min _{x \in \mathbf{R}^{n}} f(x)
$$

Given a starting point $x_{c}$, the search direction is computed according to the formula

$$
d=-B^{-1} g_{c}
$$

where $B$ is a positive definite approximation of the Hessian and $g_{c}$ is the gradient evaluated at $x_{c}$. A line search is then used to find a new point

$$
x_{n}=x_{c}+\lambda d, \quad \lambda>0
$$

such that

$$
f\left(x_{n}\right) \leq f\left(x_{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality condition $\|g(x)\|=\varepsilon$ is checked where $\varepsilon$ is a gradient tolerance.
When optimality is not achieved, $B$ is updated according to the BFGS formula

$$
B \leftarrow B-\frac{B s s^{T} B}{s^{T} B s}+\frac{y y^{T}}{y^{T} s}
$$

where $s=x_{n}-x_{c}$ and $y=g_{n}-g_{c}$. Another search direction is then computed to begin the next iteration. For more details, see Dennis and Schnabel (1983, Appendix A).

Since a finite-difference method is used to estimate the gradient, for some single precision calculations, an inaccurate estimate of the gradient may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact gradient can be easily provided, IMSL routine Uming should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2INF/DU2INF. The reference is:

CALL U2INF (FCN, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE, WK)
The additional argument is:
$\boldsymbol{W K}$ - Work vector of length $\mathrm{N}(\mathrm{N}+8)$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain the Cholesky factorization of a BFGS approximation to the Hessian at the solution.
2. Informational errors

Type Code
$4 \quad 2$ The iterates appear to be converging to a noncritical point.
43 Maximum number of iterations exceeded.
$4 \quad 4 \quad$ Maximum number of function evaluations exceeded.
45 Maximum number of gradient evaluations exceeded.
46 Five consecutive steps have been taken with the maximum step length.
27 Scaled step tolerance satisfied; the current point may be an approximate local solution, or the algorithm is making very slow progress and is not near a solution, or STEPTL is too big.
38 The last global step failed to locate a lower point than the current x value.
3. The first stopping criterion for UMINF occurs when the infinity norm of the scaled gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for Uminf occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for Uminf, then set IPARAM(1) to zero and call the routine UMINF. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UMINF:

CALL U4INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter.
If $\operatorname{IPARAM}(6)=0$, the Hessian is initialized to the identity matrix;
otherwise, it is initialized to a diagonal matrix containing

$$
\max \left(|f(t)|, f_{s}\right) * s_{i}^{2}
$$

on the diagonal where $t=\mathrm{XGUESS}, f_{s}=\mathrm{FSCALE}$, and $s=\mathrm{XSCALE}$.
Default: 0 .
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: Not used in UMINF.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at $x$ is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=\mathrm{XSCALE}$, and $f_{s}=\mathrm{FSCALE}$.

Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. (STEPTL)
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon 2 / 3$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: Not used in UMINF.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in UMINF.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: Not used in UMINF.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}, \varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}, \text { and } t=\mathrm{XGUESS}
$$

$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: Not used in UMINF.
If double precision is required, then DU4 INF is called, and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized.

```
USE UMINF_INT
USE U4INF }\mp@subsup{}{-}{INT
```

```
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
    INTEGER IPARAM(7), L, NOUT
    REAL F, RPARAM(7), X(N), XGUESS (N), &
        XSCALE (N)
    EXTERNAL ROSBRK
    DATA XGUESS/-1.2E0, 1.0E0/
    Relax gradient tolerance stopping
    criterion
    CALL U4INF (IPARAM, RPARAM)
    RPARAM(1) = 10.0E0*RPARAM(1)
                                    Minimize Rosenbrock function using
                                    initial guesses of -1.2 and 1.0
    CALL UMINF (ROSBRK, X, XGUESS=XGUESS, IPARAM=IPARAM, RPARAM=RPARAM, &
        FVALUE=F)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5)
!
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
    'value is ', F8.3, //, ' The number of iterations is ', &
    10X, I3, /, ' The number of function evaluations is ', &
    I3, /, ' The number of gradient evaluations is ', I3)
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
!
    RETURN
    END
```


## Output

```
The solution is 1.000 1.000
The function value is 0.000
The number of iterations is 15
The number of function evaluations is 40
The number of gradient evaluations is 19
```


## UMING

Minimizes a function of N variables using a quasi-Newton method and a user-supplied gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN ( $\mathrm{N}, \mathrm{X}, \mathrm{F}$ ), where

N - Length of X. (Input)

X - Vector of length N at which point the function is evaluated. (Input)
$X$ should not be changed by FCN.
F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
$\mathrm{X}-$ Vector of length N at which point the function is evaluated. (Input)
$X$ should not be changed by GRAD .
G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

XGUESS - Vector of length N containing the initial guess of the minimum. (Input)
Default: XGUESS $=0.0$.
XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0 .
Default: $\mathrm{XSCALE}=1.0$.

FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.
IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4. Default: $\operatorname{IPARAM}=0$.

RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMING (FCN, GRAD, X [,...])

Specific: The specific interface names are S_UMING and D_UMING.

## FORTRAN 77 Interface

Single: CALL UMING (FCN, GRAD, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, $X$, FVALUE)

Double: The double precision name is Duming.

## Description

The routine UMING uses a quasi-Newton method to find the minimum of a function $f(x)$ of $n$ variables. Function values and first derivatives are required. The problem is stated as follows:

$$
\min _{x \in \mathbb{R}^{n}} f(x)
$$

Given a starting point $x_{c}$, the search direction is computed according to the formula

$$
d=-B^{-1} g_{c}
$$

where $B$ is a positive definite approximation of the Hessian and $g_{c}$ is the gradient evaluated at $x_{c}$. A line search is then used to find a new point

$$
x_{n}=x_{c}+\lambda d, \quad \lambda>0
$$

such that

$$
f\left(x_{n}\right) \leq f\left(x_{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality condition $\|g(x)\|=\varepsilon$ is checked where $\varepsilon$ is a gradient tolerance.
When optimality is not achieved, $B$ is updated according to the BFGS formula

$$
B \leftarrow B-\frac{B s s^{T} B}{s^{T} B s}+\frac{y y^{T}}{y^{T} s}
$$

where $s=x_{n}-x_{c}$ and $y=g_{n}-g_{c}$. Another search direction is then computed to begin the next iteration. For more details, see Dennis and Schnabel (1983, Appendix A).

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2ING/DU2 ING. The reference is:
```
CALL U2ING (FCN, GRAD, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X,
FVALUE,WK)
```

The additional argument is
WK - Work vector of length N * ( $\mathrm{N}+8$ ). WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain the Cholesky factorization of a BFGS approximation to the Hessian at the solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 5 | Maximum number of gradient evaluations exceeded. |
| 4 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |
|  |  | 8 |
| The last global step failed to locate a lower point than the current X |  |  |
| value. |  |  |

3. The first stopping criterion for UMING occurs when the infinity norm of the scaled gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for UMING occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for UMIng, then set IPARAM(1) to zero and call routine UMING. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UMING:

CALL U4 INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4 INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
IPARAM(1) $=$ Initialization flag.
IPARAM(2) = Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter
If $\operatorname{IPARAM}(6)=0$, the Hessian is initialized to the identity matrix;
otherwise, it is initialized to a diagonal matrix containing

$$
\max \left(|f(t)|, f_{s}\right) * s_{i}^{2}
$$

on the diagonal where $t=\mathrm{XGUESS}, f_{s}=\mathrm{FSCALE}$, and $s=\mathrm{XSCALE}$.
Default: 0 .
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: Not used in UMING.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at
$x$ is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=\mathrm{XSCALE}$, and $f_{s}=\mathrm{FSCALE}$.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: Not used in UMING.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in Uming.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: Not used in Uming.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\begin{gathered}
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}} \\
\varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}, \text { and } t=\mathrm{XGUESS} .
\end{gathered}
$$

$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: Not used in Uming.
If double precision is required, then DU4 INF is called, and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized. Default values for parameters are used.

```
USE UMING_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=2)
INTEGER IPARAM(7), L, NOUT
REAL F, X(N), XGUESS(N)
EXTERNAL ROSBRK, ROSGRD
DATA XGUESS/-1.2E0, 1.0E0/
IPARAM(1) = 0
Minimize Rosenbrock function using
initial guesses of -1.2 and 1.0
```

!

```
CALL UMING (ROSBRK, ROSGRD, X, XGUESS=XGUESS, IPARAM=IPARAM, FVALUE=F)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5)
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
    'value is ', F8.3, //, ' The number of iterations is ', &
    10X, I3, /, ' The number of function evaluations is ', &
    I3, /, ' The number of gradient evaluations is ', I3)
END
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
    RETURN
    END
SUBROUTINE ROSGRD (N, X, G)
INTEGER N
REAL X(N), G(N)
G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
G(2) = 2.0E2*(X(2)-X(1)*X(1))
    RETURN
    END
```

!
$!$
$!$
$!$
!

## Output

| The solution is | $1.000 \quad 1.000$ |  |  |
| :--- | :--- | :--- | :--- |
| The function value is | 0.000 |  |  |
|  |  |  |  |
| The number of iterations is | 18 |  |  |
| The number of function evaluations is | 31 |  |  |
| The number of gradient evaluations is | 22 |  |  |

## UMIDH

Minimizes a function of N variables using a modified Newton method and a finite-difference Hessian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where N - Length of X . (Input)

X - Vector of length N at which point the function is evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function value at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.

GRAD - User-supplied subroutine to compute the gradient at the point X . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
X - The point at which the gradient is evaluated. (Input) $X$ should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)

GRAD must be declared EXTERNAL in the calling program.
$\boldsymbol{X}-$ Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
$\boldsymbol{X G U E S S}$ - Vector of length N containing initial guess. (Input)
Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0 .
Default: $\mathrm{XSCALE}=1.0$.
FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.

IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: $\operatorname{IPARAM}=0$.
RPARAM - Parameter vector of length 7. (Input/Output)
See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMIDH (FCN, GRAD, X [, ...])
Specific: The specific interface names are S_UMIDH and D_UMIDH.

## FORTRAN 77 Interface

Single: CALL UMIDH (FCN, GRAD, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DUMIDH.

## Description

The routine UMIDH uses a modified Newton method to find the minimum of a function $f(x)$ of $n$ variables. First derivatives must be provided by the user. The algorithm computes an optimal locally constrained step (Gay 1981) with a trust region restriction on the step. It handles the case that the Hessian is indefinite and provides a way to deal with negative curvature. For more details, see Dennis and Schnabel (1983, Appendix A) and Gay (1983).

Since a finite-difference method is used to estimate the Hessian for some single precision calculations, an inaccurate estimate of the Hessian may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact Hessian can be easily provided, IMSL routine UMIAH should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2IDH / DU2 IDH. The reference is:
```
CALL U2IDH (FCN, GRAD, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X,
FVALUE,WK)
```

The additional argument is:
$\boldsymbol{W K}$ - Work vector of length $\mathrm{N} *(\mathrm{~N}+9)$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain the Hessian at the approximate solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |


| 4 | 5 | Maximum number of gradient evaluations exceeded. <br> 4 |
| :--- | :--- | :--- |
| 2 | 7 | Five consecutive steps have been taken with the maximum step <br> length. <br> Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |
| 4 | 7 | Maximum number of Hessian evaluations exceeded. |
| 3 | 8 | The last global step failed to locate a lower point than the current X <br> value. |

3. The first stopping criterion for UMIDH occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for UMIDH occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for UMIDH, then set IPARAM(1) to zero and call routine UMIDH. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UMIDH:

CALL U4 INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter
Default: Not used in UMIDH.
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default:100

RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at $x$ is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=$ XSCALE, and $f_{s}=$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision. $\operatorname{RPARAM}(3)=$ Relative function tolerance.

Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in UMIDH.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}$, and $t=\mathrm{XGUESS}$. $\operatorname{RPARAM}(7)=$ Size of initial trust region radius.

Default: Based on initial scaled Cauchy step.

If double precision is required, then DU4INF is called, and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized. Default values for parameters are used.

```
    USE UMIDH_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
    INTEGER IPARAM(7), L, NOUT
    REAL F, X(N), XGUESS (N)
    EXTERNAL ROSBRK, ROSGRD
    DATA XGUESS/-1.2E0, 1.0E0/
    IPARAM(1) = 0
        Minimize Rosenbrock function using
        initial guesses of -1.2 and 1.0
    CALL UMIDH (ROSBRK, ROSGRD, X, XGUESS=XGUESS, IPARAM=IPARAM, FVALUE=F)
    Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5), IPARAM(7)
!
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
            'value is ', F8.3, //, ' The number of iterations is ', &
            10X, I3, /, ' The number of function evaluations is ', &
            I3, /, ' The number of gradient evaluations is ', I3, /, &
            ' The number of Hessian evaluations is ', I3)
!
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
!
    RETURN
    END
!
SUBROUTINE ROSGRD (N, X, G)
INTEGER N
```

```
    REAL X(N),G(N)
!
    G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2* (X(2)-X(1)*X(1))
!
    RETURN
    END
```


## Output

```
The solution is 1.000 1.000
The function value is 0.000
The number of iterations is 21
The number of function evaluations is 30
The number of gradient evaluations is 22
The number of Hessian evaluations is 21
```


## UMIAH

Minimizes a function of N variables using a modified Newton method and a user-supplied Hessian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
$X$ - Vector of length $N$ at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
X - Vector of length N at which point the gradient is evaluated. (Input)
$X$ should not be changed by GRAD.
G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.

HESS - User-supplied subroutine to compute the Hessian at the point x . The usage is CALL HESS ( $\mathrm{N}, \mathrm{X}, \mathrm{H}, \mathrm{LDH}$ ), where

N - Length of X . (Input)
$\mathrm{X}-$ Vector of length N at which point the Hessian is evaluated. (Input) $X$ should not be changed by HESS.

H - The Hessian evaluated at the point X . (Output)
LDH - Leading dimension of $H$ exactly as specified in the dimension statement of the calling program. (Input)

HESS must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length n containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing initial guess. (Input)
Default: $\mathrm{XGUESS}=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0.
Default: XSCALE $=1.0$.
FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\operatorname{FSCALE}=1.0$.
IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: IPARAM $=0$.
RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMIAH (FCN, GRAD, HESS, X, [, ...])
Specific: The specific interface names are S_UMIAH and D_UMIAH.

## FORTRAN 77 Interface

Single: CALL UMIAH (FCN, GRAD, HESS, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DUMIAH.

## Description

The routine UMIAH uses a modified Newton method to find the minimum of a function $f(x)$ of $n$ variables. First and second derivatives must be provided by the user. The algorithm computes an optimal locally constrained step (Gay 1981) with a trust region restriction on the step. This algorithm handles the case where the Hessian is indefinite and provides a way to deal with negative curvature. For more details, see Dennis and Schnabel (1983, Appendix A) and Gay (1983).

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2IAH / DU2 IAH. The reference is:

CALL U2IAH (FCN, GRAD, HESS, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE, WK)

The additional argument is:
$W K$ - Work vector of length $N *(N+9)$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $N^{2}$ locations contain the Hessian at the approximate solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 5 | Maximum number of gradient evaluations exceeded. |
| 4 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |

7 Scaled step tolerance satisfied; the current point may be an approximate local solution, or the algorithm is making very slow progress and is not near a solution, or STEPTL is too big.
7 Maximum number of Hessian evaluations exceeded.
38 The last global step failed to locate a lower point than the current X value.
3. The first stopping criterion for UMIAH occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for UMIAH occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for UMIAH, then set IPARAM(1) to zero and call the routine UMIAH. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UMIAH:

CALL U4INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter
Default: Not used in UMIAH.
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: 100.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at $x$ is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=$ XSCALE, and $f_{s}=$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in UMIAH.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=\operatorname{XSCALE}$, and $t=\mathrm{XGUESS}$.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is required, then DU4INF is called, and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized. Default values for parameters are used.

```
    USE UMIAH_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
    INTEGER IPARAM(7), L, NOUT
    REAL F, FSCALE, RPARAM(7), X(N), &
        XGUESS(N), XSCALE (N)
    EXTERNAL ROSBRK, ROSGRD, ROSHES
    DATA XGUESS/-1.2E0, 1.0E0/, XSCALE/1.0E0, 1.0E0/, FSCALE/1.0E0/
    IPARAM(1) = 0
        Minimize Rosenbrock function using
        initial guesses of -1.2 and 1.0
    CALL UMIAH (ROSBRK, ROSGRD, ROSHES, X, XGUESS=XGUESS, IPARAM=IPARAM, &
        FVALUE=F)
! Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5), IPARAM(7)
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
            'value is ', F8.3, //, ' The number of iterations is ', &
            10X, I3, /, ' The number of function evaluations is ', &
            I3, /, ' The number of gradient evaluations is ', I3, /, &
            ' The number of Hessian evaluations is ', I3)
    END
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F = 1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
    RETURN
    END
    SUBROUTINE ROSGRD (N, X, G)
    INTEGER N
    REAL X(N), G(N)
```

!
!
!
!
!
!

```
!
    G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
!
    SUBROUTINE ROSHES (N, X, H, LDH)
    INTEGER N, LDH
    REAL X(N), H(LDH,N)
!
    H(1,1) = -4.0E2*X(2) + 1.2E3*X(1)*X(1) + 2.0E0
    H(2,1) = -4.0E2*X(1)
    H(1, 2) = H(2,1)
    H}(2,2)=2.0E
!
RETURN
END
```


## Output

```
The solution is 1.000 1.000
The function value is 0.000
The number of iterations is 21
The number of function evaluations is 31
The number of gradient evaluations is 22
The number of Hessian evaluations is 21
```


## UMCGF

Minimizes a function of N variables using a conjugate gradient algorithm and a finite-difference gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
$X$ - The point at which the function is evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
DFPRED - A rough estimate of the expected reduction in the function. (Input) DFPRED is used to determine the size of the initial change to x .
$\boldsymbol{X}-$ Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

XGUESS - Vector of length N containing the initial guess of the minimum. (Input) Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
Default: $\mathrm{XSCALE}=1.0$.

GRADTL - Convergence criterion. (Input)
The calculation ends when the sum of squares of the components of G is less than GRADTL.
Default: GRADTL $=1 . \mathrm{e}-4$.
MAXFN - Maximum number of function evaluations. (Input)
If MAXFN is set to zero, then no restriction on the number of function evaluations is set. Default: MAXFN $=0$.
$\boldsymbol{G}$ - Vector of length N containing the components of the gradient at the final parameter estimates. (Output)

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMCGF (FCN, DFPRED, X [,...])
Specific: The specific interface names are S_UMCGF and D_UMCGF.

## FORTRAN 77 Interface

Single: CALL UMCGF (FCN, N, XGUESS, XSCALE, GRADTL, MAXFN, DFPRED, X, G, FVALUE)

Double: The double precision name is DUMCGF.

## Description

The routine UMCGF uses a conjugate gradient method to find the minimum of a function $f(x)$ of $n$ variables. Only function values are required.
The routine is based on the version of the conjugate gradient algorithm described in Powell (1977). The main advantage of the conjugate gradient technique is that it provides a fast rate of
convergence without the storage of any matrices. Therefore, it is particularly suitable for unconstrained minimization calculations where the number of variables is so large that matrices of dimension $n$ cannot be stored in the main memory of the computer. For smaller problems, however, a routine such as routine UMINF, is usually more efficient because each iteration makes use of additional information from previous iterations.
Since a finite-difference method is used to estimate the gradient for some single precision calculations, an inaccurate estimate of the gradient may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact gradient can be easily provided, routine UMCGG should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2CGF/DU2CGF. The reference is:
```
CALL U2CGF (FCN, N, XGUESS, XSCALE, GRADTL, MAXFN, DFPRED, X, G, FVALUE, S,
RSS,RSG, GINIT, XOPT, GOPT)
```

The additional arguments are as follows:
$\boldsymbol{S}$ - Vector of length N used for the search direction in each iteration.
$\boldsymbol{R S S}$ - Vector of length N containing conjugacy information.
$\boldsymbol{R S G}$ - Vector of length N containing conjugacy information.
GINIT - Vector of length N containing the gradient values at the start of an iteration.
XOPT - Vector of length $n$ containing the parameter values that yield the least calculated value for FVALUE.

GOPT - Vector of length N containing the gradient values that yield the least calculated value for FVALUE.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 1 | The line search of an integration was abandoned. This error may be <br> caused by an error in gradient. |
| 4 | 2 | The calculation cannot continue because the search is uphill. |
| 4 | 3 | The iteration was terminated because MAXFN was exceeded. |
| 3 | 4 | The calculation was terminated because two consecutive iterations <br> failed to reduce the function. |

3. Because of the close relation between the conjugate-gradient method and the method of steepest descent, it is very helpful to choose the scale of the variables in a way that balances the magnitudes of the components of a typical gradient vector. It can be particularly inefficient if a few components of the gradient are much larger than the rest.
4. If the value of the parameter GRADTL in the argument list of the routine is set to zero, then the subroutine will continue its calculation until it stops reducing the objective function. In this case, the usual behavior is that changes in the objective function become dominated by computer rounding errors before precision is lost in the gradient vector. Therefore, because the point of view has been taken that the user requires the least possible value of the function, a value of the objective function that is small due to computer rounding errors can prevent further progress. Hence, the precision in the final values of the variables may be only about half the number of significant digits in the computer arithmetic, but the least value of FVALUE is usually found to be quite accurate.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized and the solution is printed.

```
    USE UMCGF_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
!
    INTEGER I, MAXFN, NOUT
    REAL DFPRED, FVALUE, G(N), GRADTL, X(N), XGUESS (N)
    EXTERNAL ROSBRK
!
    DATA XGUESS/-1.2E0, 1.0E0/
!
    DFPRED = 0.2
    GRADTL = 1.0E-6
    MAXFN = 100
! Minimize the Rosenbrock function
    CALL UMCGF (ROSBRK, DFPRED, X, xguess=xguess, gradtl=gradtl, &
        g=g, fvalue=fvalue)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) (X (I), I=1,N), FVALUE, (G (I), I=1,N)
99999 FORMAT (' The solution is ', 2F8.3, //, ' The function ', &
        'evaluated at the solution is ', F8.3, //, ' The ', &
        'gradient is ', 2F8.3, /)
!
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
!
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
```

RETURN
END

## Output

```
The solution is 0.999 0.998
The function evaluated at the solution is 0.000
The gradient is -0.001 0.000
```


## UMCGG

Minimizes a function of $N$ variables using a conjugate gradient algorithm and a user-supplied gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
X - The point at which the function is evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
X - The point at which the gradient is evaluated. (Input) $X$ should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.

DFPRED - A rough estimate of the expected reduction in the function. (Input) DFPRED is used to determine the size of the initial change to X .
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
$\boldsymbol{X G U E S S}$ - Vector of length N containing the initial guess of the minimum. (Input) Default: XGUESS $=0.0$.

GRADTL - Convergence criterion. (Input)
The calculation ends when the sum of squares of the components of $G$ is less than GRADTL.
Default: GRADTL $=1 . \mathrm{e}-4$.

MAXFN - Maximum number of function evaluations. (Input)
Default: MAXFN $=100$.
$\boldsymbol{G}$ - Vector of length N containing the components of the gradient at the final parameter estimates. (Output)

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMCGG (FCN, GRAD, DFPRED, X [, ...])
Specific: The specific interface names are S_UMCGG and D_UMCGG.

## FORTRAN 77 Interface

Single: CALL UMCGG (FCN, GRAD, N, XGUESS, GRADTL, MAXFN, DFPRED, X, G, FVALUE)

Double: The double precision name is DUMCGG.

## Description

The routine UMCGG uses a conjugate gradient method to find the minimum of a function $f(x)$ of $n$ variables. Function values and first derivatives are required.

The routine is based on the version of the conjugate gradient algorithm described in Powell (1977). The main advantage of the conjugate gradient technique is that it provides a fast rate of convergence without the storage of any matrices. Therefore, it is particularly suitable for unconstrained minimization calculations where the number of variables is so large that matrices of dimension $n$ cannot be stored in the main memory of the computer. For smaller problems, however, a subroutine such as IMSL routine UMING, is usually more efficient because each iteration makes use of additional information from previous iterations.

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2CGG/DU2CGG. The reference is:

CALL U2CGG (FCN, GRAD, N, XGUESS, GRADTL, MAXFN, DFPRED, X, G, FVALUE, S, RSS, RSG, GINIT, XOPT, GOPT)

The additional arguments are as follows:
$\boldsymbol{S}$ - Vector of length N used for the search direction in each iteration.
$\boldsymbol{R S S}$ - Vector of length N containing conjugacy information.
$\boldsymbol{R S G}$ - Vector of length N containing conjugacy information.
GINIT - Vector of length n containing the gradient values at the start on an iteration.
$\boldsymbol{X O P T}$ - Vector of length N containing the parameter values which yield the least calculated value for FVALUE.

GOPT - Vector of length N containing the gradient values which yield the least calculated value for FVALUE.
2. Informational errors

| Type | Code |  |
| :---: | :--- | :--- |
| 4 | 1 | The line search of an integration was abandoned. This error may be <br> caused by an error in gradient. |
| 4 | 2 | The calculation cannot continue because the search is uphill. |
| 4 | 3 | The iteration was terminated because MAXFN was exceeded. |
| 3 | 4 | The calculation was terminated because two consecutive iterations <br> failed to reduce the function. |

3. The routine includes no thorough checks on the part of the user program that calculates the derivatives of the objective function. Therefore, because derivative calculation is a frequent source of error, the user should verify independently the correctness of the derivatives that are given to the routine.
4. Because of the close relation between the conjugate-gradient method and the method of steepest descent, it is very helpful to choose the scale of the variables in a way that balances the magnitudes of the components of a typical gradient vector. It can be particularly inefficient if a few components of the gradient are much larger than the rest.
5. If the value of the parameter GRADTL in the argument list of the routine is set to zero, then the subroutine will continue its calculation until it stops reducing the objective function. In this case, the usual behavior is that changes in the objective function become dominated by computer rounding errors before precision is lost in the gradient vector. Therefore, because the point of view has been taken that the user requires the
least possible value of the function, a value of the objective function that is small due to computer rounding errors can prevent further progress. Hence, the precision in the final values of the variables may be only about half the number of significant digits in the computer arithmetic, but the least value of FVALUE is usually found to be quite accurate.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized and the solution is printed.

```
    USE UMCGG_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
    INTEGER I, NOUT
    REAL DFPRED, FVALUE, G(N), GRADTL, X(N), &
        XGUESS (N)
    EXTERNAL ROSBRK, ROSGRD
!
    DATA XGUESS/-1.2E0, 1.0E0/
!
    DFPRED = 0.2
    GRADTL = 1.0E-7
! Minimize the Rosenbrock function
    CALL UMCGG (ROSBRK, ROSGRD, DFPRED, X, xguess=xguess, &
        gradtl=gradtl, g=g, fvalue=fvalue)
        Print the results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) (X (I), I=1,N), FVALUE, (G (I), I=1,N)
99999 FORMAT (' The solution is ', 2F8.3, //, ' The function ', &
        'evaluated at the solution is ', F8.3, //, ' The ', &
        'gradient is ', 2F8.3, /)
!
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
!
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
    RETURN
    END
!
SUBROUTINE ROSGRD (N, X, G)
INTEGER N
REAL X(N),G(N)
```

```
!
    G(1) = - 4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X (N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
    RETURN
    END
!
    SUBROUTINE ROSGRD (N, X, G)
    INTEGER N
    REAL X(N),G(N)
    G(1) = - 4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
```


## Output

```
The solution is 1.000 1.000
The function evaluated at the solution is 0.000
The gradient is 0.000 -0.000
```


## UMPOL

Minimizes a function of N variables using a direct search polytope algorithm.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
$X-$ Vector of length $N$ at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Real vector of length N containing the best estimate of the minimum found. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $N=\operatorname{SIZE}(x, 1)$.

XGUESS - Real vector of length N which contains an initial guess to the minimum. (Input) Default: XGUESS $=0.0$.
$\boldsymbol{S}$ - On input, real scalar containing the length of each side of the initial simplex.
(Input/Output)
If no reasonable information about $S$ is known, $S$ could be set to a number less than or equal to zero and UMPOL will generate the starting simplex from the initial guess with a random number generator. On output, the average distance from the vertices to the centroid that is taken to be the solution; see Comment 4.
Default: $\mathrm{S}=0.0$.
FTOL - First convergence criterion. (Input)
The algorithm stops when a relative error in the function values is less than $F T O L$, i.e. when $(F($ worst $)-F($ best $))<$ FTOL * $(1+\operatorname{ABS}(F($ best $)))$ where $F$ (worst) and $F$ (best) are the function values of the current worst and best points, respectively. Second convergence criterion. The algorithm stops when the standard deviation of the function values at the $\mathrm{N}+1$ current points is less than FTOL. If the subroutine terminates prematurely, try again with a smaller value for FTOL.
Default: FTOL = 1.e-7.
MAXFCN - On input, maximum allowed number of function evaluations. (Input/ Output) On output, actual number of function evaluations needed.
Default: $\operatorname{MAXFCN}=200$.
FVALUE - Function value at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL UMPOL (FCN, X [, ...])

Specific: The specific interface names are S_UMPOL and D_UMPOL.

## FORTRAN 77 Interface

Single: CALL UMPOL (FCN, N, XGUESS, S, FTOL, MAXFCN, X, FVALUE)
Double: $\quad$ The double precision name is DUMPOL.

## Description

The routine UMPOL uses the polytope algorithm to find a minimum point of a function $f(x)$ of $n$ variables. The polytope method is based on function comparison; no smoothness is assumed. It starts with $n+1$ points $x_{1}, x_{2}, \ldots, x_{n+1}$. At each iteration, a new point is generated to replace the worst point $x_{j}$, which has the largest function value among these $n+1$ points. The new point is constructed by the following formula:

$$
x_{k}=c+\alpha\left(c-x_{j}\right)
$$

where

$$
c=\frac{1}{n} \sum_{i \neq j} x_{i}
$$

and $\alpha(\alpha>0)$ is the reflection coefficient.
When $x_{k}$ is a best point, that is $f\left(x_{k}\right) \leq f\left(x_{i}\right)$ for $i=1, \ldots, n+1$, an expansion point is computed $x_{e}=c+\beta\left(x_{k}-c\right)$ where $\beta(\beta>1)$ is called the expansion coefficient. If the new point is a worst point, then the polytope would be contracted to get a better new point. If the contraction step is unsuccessful, the polytope is shrunk by moving the vertices halfway toward current best point. This procedure is repeated until one of the following stopping criteria is satisfied:
Criterion 1:

$$
f_{\text {best }}-f_{\text {worst }} \leq \varepsilon \mathrm{f}\left(1 .+\left|f_{\text {best }}\right|\right)
$$

Criterion 2:

$$
\sum_{i=1}^{n+1}\left(f_{i}-\frac{\sum_{j=1}^{n+1} f_{j}}{n+1}\right)^{2} \leq \varepsilon_{f}
$$

where $f_{i}=f\left(x_{i}\right), f_{j}=f\left(x_{j}\right)$, and $\varepsilon_{f}$ is a given tolerance. For a complete description, see Nelder and Mead (1965) or Gill et al. (1981).

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2POL/DU2POL. The reference is:

CALL U2POL (FCN, N, XGUESS, S, FTOL, MAXFCN, X, FVALUE, WK)
The additional argument is:
$\boldsymbol{W K}$ - Real work vector of length $\mathrm{N}^{* *} 2+5$ * $\mathrm{N}+1$.
2. Informational error

Type Code
41 Maximum number of function evaluations exceeded.
3. Since UMPOL uses only function value information at each step to determine a new approximate minimum, it could be quite ineficient on smooth problems compared to other methods such as those implemented in routine UMINF that takes into account
derivative information at each iteration. Hence, routine UMPOL should only be used as a last resort. Briefly, a set of $\mathrm{N}+1$ points in an N -dimensional space is called a simplex. The minimization process iterates by replacing the point with the largest function value by a new point with a smaller function value. The iteration continues until all the points cluster sufficiently close to a minimum.
4. The value returned in $S$ is useful for assessing the flatness of the function near the computed minimum. The larger its value for a given value of FTOL, the flatter the function tends to be in the neighborhood of the returned point.

## Example

The function

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

is minimized and the solution is printed.

```
USE UMPOL_INT
USE UMACH_INT
```

IMPLICIT NONE
!
Variable declarations
INTEGER N
PARAMETER ( $\mathrm{N}=2$ )
!
INTEGER K, NOUT
REAL FTOL, FVALUE, S, X(N), XGUESS (N)
EXTERNAL FCN
$!$
$!$
!
DATA XGUESS/-1.2, 1.0/
!
FTOL $=1.0 \mathrm{E}-10$
$\mathrm{S}=1.0$
!
CALL UMPOL (FCN, X, xguess=xguess, $s=s, f t o l=f t o l, \&$
fvalue=fvalue)
!
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) (X (K) , $\mathrm{K}=1, \mathrm{~N}$ ) , FVALUE
99999 FORMAT (' The best estimate for the minimum value of the', /, \&
' function is $X=(', 2(2 X, F 4.2), ~ ') ', ~ /, ~ ' w i t h ~ ', ~ \& ~$
'function value FVALUE $=$ ', E12.6)
$!$
END
! External function to be minimized
SUBROUTINE FCN (N, X, F)
INTEGER N
REAL $X(N), F$
$!$

```
F=100.0*(X(1)*X(1)-X(2))**2 + (1.0-X(1))**2
RETURN
END
```


## Output

The best estimate for the minimum value of the function is $X=(1.00 \quad 1.00)$
with function value $\operatorname{FVALUE}=0.502496 \mathrm{E}-10$

## UNLSF

Solves a nonlinear least-squares problem using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function that defines the least-squares problem. The usage is CALL FCN (M, N, X, F), where M - Length of F . (Input) N - Length of X. (Input)
$X-$ Vector of length $N$ at which point the function is evaluated. (Input) X should not be changed by FCN.
$\mathrm{F}-$ Vector of length M containing the function values at X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{M}$ - Number of functions. (Input)
$\boldsymbol{X}$ - Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Number of variables. N must be less than or equal to M. (Input) Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

XGUESS - Vector of length N containing the initial guess. (Input) Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. By
default, the values for XSCALE are set internally. See IPARAM(6) in Comment 4. Default: $\mathrm{XSCALE}=1.0$.

FSCALE - Vector of length M containing the diagonal scaling matrix for the functions.
(Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set all entries to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.
IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: I PARAM $=0$.

RPARAM - Parameter vector of length 7. (Input/Output)
See Comment 4.

FVEC - Vector of length m containing the residuals at the approximate solution. (Output)
FJAC — m by n matrix containing a finite difference approximate Jacobian at the approximate solution. (Output)

LDFJAC - Leading dimension of FJAC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFJAC = SIZE (FJAC,1).

## FORTRAN 90 Interface

Generic: CALL UNLSF (FCN, M, X [,...])
Specific: The specific interface names are S_UNLSF and D_UNLSF.

## FORTRAN 77 Interface

Single: CALL UNLSF (FCN, M, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, $X$, FVEC, FJAC, LDFJAC)

Double: $\quad$ The double precision name is DUNLSF.

## Description

The routine UNLSF is based on the MINPACK routine LMDIF by Moré et al. (1980). It uses a modified Levenberg-Marquardt method to solve nonlinear least squares problems. The problem is stated as follows:

$$
\min _{x \in \mathrm{R}^{n}} \frac{1}{2} F(x)^{T} F(x)=\frac{1}{2} \sum_{i=1}^{m} f_{i}(x)^{2}
$$

where $m \geq n, F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{m}$, and $f_{i}(x)$ is the $i$-th component function of $F(x)$. From a current point, the algorithm uses the trust region approach:

$$
\begin{gathered}
\min _{x_{n} \in \mathbf{R}^{n}}\left\|F\left(x_{c}\right)+J\left(x_{c}\right)\left(x_{n}-x_{c}\right)\right\|_{2} \\
\quad \text { subject to }\left\|x_{n}-x_{c}\right\|_{2} \leq \delta \mathrm{c}
\end{gathered}
$$

to get a new point $x_{n}$, which is computed as

$$
x_{n}=x_{c}-\left(J\left(x_{c}\right)^{T} J\left(x_{c}\right)+\mu_{c} I\right)^{-1} J\left(x_{c}\right)^{T} F\left(x_{c}\right)
$$

where $\mu_{c}=0$ if $\delta_{c} \geq\left\|\left(J\left(x_{c}\right)^{T} J\left(x_{c}\right)\right)^{-1} J\left(x_{c}\right)^{T} F\left(x_{c}\right)\right\|_{2}$ and $\mu_{c}>0$ otherwise. $F\left(x_{c}\right)$ and $J\left(x_{c}\right)$ are the function values and the Jacobian evaluated at the current point $x_{c}$. This procedure is repeated until the stopping criteria are satisfied. For more details, see Levenberg (1944), Marquardt (1963), or Dennis and Schnabel (1983, Chapter 10).
Since a finite-difference method is used to estimate the Jacobian for some single precision calculations, an inaccurate estimate of the Jacobian may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact Jacobian can be easily provided, routine UNLSJ should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2LSF / DU2LSF. The reference is:

CALL U2LSF (FCN, M, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC, WK, IWK)

The additional arguments are as follows:
WK - Real work vector of length $9 * N+3 * M-1$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Gauss-Newton step. The fourth N locations contain an estimate of the gradient at the solution.

IWK - Integer work vector of length $N$ containing the permutations used in the $Q R$ factorization of the Jacobian at the solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 3 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 3 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |

$7 \quad$ Scaled step tolerance satisfied; the current point may be an approximate local solution, or the algorithm is making very slow progress and is not near a solution, or STEPTL is too big.
3. The first stopping criterion for UNLSE occurs when the norm of the function is less than the absolute function tolerance (RPARAM(4)). The second stopping criterion occurs when the norm of the scaled gradient is less than the given gradient tolerance (RPARAM(1)). The third stopping criterion for UNLSF occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for UNLSF, then set IPARAM(1) to zero and call the routine UNLSF. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UNLSF:

CALL U4LSE (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4LSF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 6.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: Not used in UnLSF.
$\operatorname{IPARAM}(6)=$ Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values for XSCALE are set internally.
Default: 1.
RPARAM - Real vector of length 7.
RPARAM $(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\|F(x)\|_{2}^{2}}
$$

where

$$
g_{i}=\left(J(x)^{T} F(x)\right)_{i} *\left(f_{s}\right)_{i}^{2}
$$

$J(x)$ is the Jacobian, $s=$ XSCALE, and $f_{s}=$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE .
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: $\max \left(10^{-20}, \varepsilon^{2}\right), \max \left(10^{-40}, \varepsilon^{2}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=\operatorname{XSCALE}$, and $t=\mathrm{XGUESS}$.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is desired, then DU4LSF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The nonlinear least squares problem

$$
\min _{x \in \mathbf{R}^{2}} \frac{1}{2} \sum_{i=1}^{2} f_{i}(x)^{2}
$$

where

$$
f_{1}(x)=10\left(x_{2}-x_{1}^{2}\right) \text { and } f_{2}(x)=\left(1-x_{1}\right)
$$

is solved. RPARAM(4) is changed to a non-default value.

```
USE UMACH INT
USE U4LSF_INT
IMPLICIT NONE
! Declaration of variables
    INTEGER LDFJAC, M, N
    PARAMETER (LDFJAC=2, M=2, N=2)
    INTEGER IPARAM(6), NOUT
    REAL FVEC (M), RPARAM(7),X(N), XGUESS(N)
    EXTERNAL ROSBCK
! Compute the least squares for the
    DATA XGUESS/-1.2E0, 1.0E0/
    Relax the first stopping criterion by
    calling U4LSF and scaling the
    absolute function tolerance by 10.
    CALL U4LSF (IPARAM, RPARAM)
    RPARAM(4) = 10.0E0*RPARAM(4)
    CALL UNLSF (ROSBCK, M, X,xguess=xguess, iparam=iparam, rparam=rparam,&
    fvec=fvec)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, FVEC, IPARAM(3), IPARAM(4)
99999 FORMAT (' The solution is ', 2F9.4, //, ' The function ', &
            'evaluated at the solution is ', /, 18X, 2F9.4, //, &
            ' The number of iterations is ', 10X, I3, /, ' The ', &
            'number of function evaluations is ', I3, /)
    END
    SUBROUTINE ROSBCK (M, N, X, F)
    INTEGER M, N
    REAL X(N), F(M)
    F(1) = 10.0E0*(X(2)-X(1)*X(1))
```

!
!
!
!

```
F(2) = 1.0E0 - X(1)
RETURN
END
```


## Output

```
The solution is 1.0000 1.0000
The function evaluated at the solution is
0.0000 0.0000
The number of iterations is 24
The number of function evaluations is 33
```


## UNLSJ

Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function which defines the least-squares problem. The usage is
CALL FCN (M, N, X, F), where
M - Length of F . (Input)
N - Length of X . (Input)
X - Vector of length N at which point the function is evaluated. (Input)
$X$ should not be changed by FCN .
F - Vector of length M containing the function values at X . (Output)
FCN must be declared EXTERNAL in the calling program.
JAC - User-supplied subroutine to evaluate the Jacobian at a point X . The usage is
CALL JAC (M, N, X, FJAC, LDFJAC), where
M - Length of F . (Input)
N - Length of X . (Input)
$X$ - Vector of length $N$ at which point the Jacobian is evaluated. (Input)
X should not be changed by JAC.
FJAC - The computed M by n Jacobian at the point X. (Output)
LDFJAC - Leading dimension of FJAC. (Input)
JAC must be declared EXTERNAL in the calling program.
$\boldsymbol{M}$ - Number of functions. (Input)
$\boldsymbol{X}$ - Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Number of variables. N must be less than or equal to M . (Input)
Default: $\mathrm{N}=\operatorname{sizE}(\mathrm{x}, 1)$.

XGUESS - Vector of length N containing the initial guess. (Input) Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. By default, the values for XSCALE are set internally. See IPARAM(6) in Comment 4.
Default: $\mathrm{XSCALE}=1.0$.
FSCALE - Vector of length M containing the diagonal scaling matrix for the functions. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set all entries to 1.0 .
Default: $\operatorname{FSCALE}=1.0$.
IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4. Default: IPARAM $=0$.

RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVEC - Vector of length M containing the residuals at the approximate solution. (Output)
$\boldsymbol{F} \boldsymbol{J} \boldsymbol{A C}$ - m by n matrix containing a finite-difference approximate Jacobian at the approximate solution. (Output)

LDFJAC - Leading dimension of FJAC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFJAC = SIZE (FJAC, 1).

## FORTRAN 90 Interface

Generic: CALL UNLSJ (FCN, JAC, M, X [, ...])

Specific: The specific interface names are S_UNLSJ and D_UNLSJ.

## FORTRAN 77 Interface

Single: CALL UNLSJ (FCN, JAC, M, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC)

Double: $\quad$ The double precision name is DUNLSJ.

## Description

The routine UNLSJ is based on the MINPACK routine LMDER by Moré et al. (1980). It uses a modified Levenberg-Marquardt method to solve nonlinear least squares problems. The problem is stated as follows:

$$
\min _{x \in \mathbf{R}^{n}} \frac{1}{2} F(x)^{T} F(x)=\frac{1}{2} \sum_{i=1}^{m} f_{i}(x)^{2}
$$

where $m \geq n, F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{m}$, and $f_{i}(x)$ is the $i$-th component function of $F(x)$. From a current point, the algorithm uses the trust region approach:

$$
\begin{aligned}
& \min _{x_{n} \in \mathbf{R}^{n}}\left\|F\left(x_{c}\right)+J\left(x_{c}\right)\left(x_{n}-x_{c}\right)\right\|_{2} \\
& \quad \text { subject to }\left\|x_{n}-x_{c}\right\|_{2} \leq \delta_{c}
\end{aligned}
$$

to get a new point $x_{n}$, which is computed as

$$
x_{n}=x_{c}-\left(J\left(x_{c}\right)^{T} J\left(x_{c}\right)+\mu_{c} I\right)^{-1} J\left(x_{c}\right)^{T} F\left(x_{c}\right)
$$

where $\mu_{c}=0$ if $\delta_{c} \geq\left\|\left(J\left(x_{c}\right)^{T} J\left(x_{c}\right)\right)^{-1} J\left(x_{c}\right)^{T} F\left(x_{c}\right)\right\|_{2}$ and $\mu_{c}>0$ otherwise. $F\left(x_{c}\right)$ and $J\left(x_{c}\right)$ are the function values and the Jacobian evaluated at the current point $x_{c}$. This procedure is repeated until the stopping criteria are satisfied. For more details, see Levenberg (1944), Marquardt(1963), or Dennis and Schnabel (1983, Chapter 10).

## Comments

1. Workspace may be explicitly provided, if desired, by use of U2LSJ/DU2LSJ. The reference is:

CALL U2LSJ (FCN, JAC, M, N, XGUESS, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC, WK, IWK)
The additional arguments are as follows:
WK - Work vector of length $9 * \mathrm{~N}+3 * \mathrm{M}-1$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Gauss-Newton step. The fourth $N$ locations contain an estimate of the gradient at the solution.
$\boldsymbol{I W K}$ - Work vector of length $N$ containing the permutations used in the $Q R$ factorization of the Jacobian at the solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 3 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |


| 4 | 4 | Maximum number of function evaluations exceeded. |
| :--- | :--- | :--- |
| 4 | 5 | Maximum number of Jacobian evaluations exceeded. |
| 3 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |

3. The first stopping criterion for UNLS $J$ occurs when the norm of the function is less than the absolute function tolerance (RPARAM(4)). The second stopping criterion occurs when the norm of the scaled gradient is less than the given gradient tolerance (RPARAM(1)). The third stopping criterion for UNLSJ occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for UNLS J, then set IPARAM(1) to zero and call the routine UNLSJ. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling UNLSJ:

CALL U4LSF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4LSF will set IPARAM and RPARAM to their default values, so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 6.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: 100.
$\operatorname{IPARAM}(6)=$ Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values for XSCALE are set internally.
Default: 1.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\|F(x)\|_{2}^{2}}
$$

where

$$
g_{i}=\left(J(x)^{T} F(x)\right)_{i} *\left(f_{s}\right)_{i}^{2}
$$

$J(x)$ is the Jacobian, $s=$ XSCALE, and $f_{s}$ FSSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right)$, $\max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision. $\operatorname{RPARAM}(4)=$ Absolute function tolerance.

Default: $\max \left(10^{-20}, \varepsilon^{2}\right), \max \left(10^{-40}, \varepsilon^{2}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1} \sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}$, and $t=\mathrm{XGUESS}$.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.

If double precision is desired, then DU4LSF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The nonlinear least-squares problem

$$
\min _{x \in \mathbf{R}^{2}} \frac{1}{2} \sum_{i=1}^{2} f_{i}(x)^{2}
$$

where

$$
f_{1}(x)=10\left(x_{2}-x_{1}^{2}\right) \text { and } f_{2}(x)=\left(1-x_{1}\right)
$$

is solved; default values for parameters are used.

```
USE UNLSJ INT
USE UMACH_
IMPLICIT NONE
! Declaration of variables
    INTEGER LDFJAC, M, N
    PARAMETER (LDFJAC=2, M=2, N=2)
    INTEGER IPARAM(6), NOUT
    REAL FVEC (M), X(N), XGUESS (N)
    EXTERNAL ROSBCK, ROSJAC
        Compute the least squares for the
        Rosenbrock function.
    DATA XGUESS/-1.2E0, 1.0E0/
    IPARAM(1) = 0
    CALL UNLSJ (ROSBCK, ROSJAC, M, X, XGUESS=XGUESS, &
        IPARAM=IPARAM, FVEC=FVEC)
        Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, FVEC, IPARAM(3), IPARAM(4), IPARAM(5)
99999 FORMAT (' The solution is ', 2F9.4, //, ' The function ', &
            'evaluated at the solution is ', /, 18X, 2F9.4, //, &
            ' The number of iterations is ', 10X, I3, /, ' The ', &
            'number of function evaluations is ', I3, /, ' The ', &
            'number of Jacobian evaluations is ', I3, /)
            END
    SUBROUTINE ROSBCK (M, N, X, F)
    INTEGER M, N
    REAL X(N), F(M)
    F(1) = 10.0EO*(X(2)-X(1)*X(1))
    F(2) = 1.0E0 - X(1)
```

!
!
!
!

```
    RETURN
    END
!
    SUBROUTINE ROSJAC (M, N, X, FJAC, LDFJAC)
    INTEGER M, N, LDFJAC
    REAL X(N), FJAC(LDFJAC,N)
!
    FJAC(1,1) = -20.0E0*X(1)
    FJAC (2,1) = -1.0E0
    FJAC (1,2) = 10.0E0
    FJAC (2,2) = 0.0E0
    RETURN
    END
```


## Output

```
The solution is 1.0000 1.0000
The function evaluated at the solution is
0.0000 0.0000
The number of iterations is 23
The number of function evaluations is 32
The number of Jacobian evaluations is 24
```


## BCONF

Minimizes a function of N variables subject to bounds on the variables using a quasi-Newton method and a finite-difference gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
X - Vector of length N at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

$0 \quad$ User will supply all the bounds.

1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables. (Input, if $\operatorname{IBTYPE}=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ )
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on variables. (Input, if $\operatorname{IBTYPE}=0$; output, if IBTYPE $=1$ or 2 ; input/output, if IBTYPE $=3$ )
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length $N$ containing an initial guess of the computed solution. (Input) Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0.
Default: XSCALE $=1.0$.
FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.

IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: $\operatorname{IPARAM}=0$.

RPARAM - Parameter vector of length 7. (Input/Output)
See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL BCONF (FCN, IBTYPE, XLB, XUB, X [, ...])

Specific: The specific interface names are S_BCONF and D_BCONF.

## FORTRAN 77 Interface

Single: CALL BCONF (FCN, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DBCONF.

## Description

The routine BCONF uses a quasi-Newton method and an active set strategy to solve minimization problems subject to simple bounds on the variables. The problem is stated as follows:

$$
\min _{x \in \mathbf{R}^{n}} f(x)
$$

subject to $l \leq x \leq u$
From a given starting point $x^{c}$, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-B^{-1} g^{c}
$$

where $B$ is a positive definite approximation of the Hessian and $g^{c}$ is the gradient evaluated at $x^{c}$; both are computed with respect to the free variables. The search direction for the variables in IA is set to zero. A line search is used to find a new point $x^{n}$,

$$
x^{n}=x^{c}+\lambda d, \quad \lambda \in(0,1]
$$

such that

$$
f\left(x^{n}\right) \leq f\left(x^{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality conditions

$$
\begin{gathered}
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{i}<x_{i}<u_{i} \\
g\left(x_{i}\right)<0, x_{i}=u_{i} \\
g\left(x_{i}\right)>0, x_{i}=l_{i}
\end{gathered}
$$

are checked, where $\varepsilon$ is a gradient tolerance. When optimality is not achieved, $B$ is updated according to the BFGS formula:

$$
B \leftarrow B-\frac{B s s^{T} B}{s^{T} B s}+\frac{y y^{T}}{y^{T} s}
$$

where $s=x^{n}-x^{c}$ and $y=g^{n}-g^{c}$. Another search direction is then computed to begin the next iteration.

The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more details on the quasi-Newton method and line search, see Dennis and Schnabel (1983). For more detailed information on active set strategy, see Gill and Murray (1976).

Since a finite-difference method is used to estimate the gradient for some single precision calculations, an inaccurate estimate of the gradient may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact gradient can be easily provided, routine BCONG should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2ONF/DB2ONF. The reference is:
```
CALL B2ONF (FCN, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM,
RPARAM, X, FVALUE, WK, IWK)
```

The additional arguments are as follows:
WK - Real work vector of length $N$ * $(2 * N+8)$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain a BFGS approximation to the Hessian at the solution. Only the lower triangular portion of the matrix is stored in wK. The values returned in the upper triangle should be ignored.

IWK - Work vector of length N stored in column order.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :---: |
| 3 | 1 | Both the actual and predicted relative reductions in the function are less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 5 | Maximum number of gradient evaluations exceeded. |
| 4 | 6 | Five consecutive steps have been taken with the maximum step length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an approximate local solution, or the algorithm is making very slow progress and is not near a solution, or STEPTL is too big. |
| 3 | 8 | The last global step failed to locate a lower point than the current X value. |

3. The first stopping criterion for BCONF occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for BCONF occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for BCONF, then set IPARAM(1) to zero and call the routine BCONF. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling BCONF:
```
CALL U4INF (IPARAM, RPARAM)
```

Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter.
If $\operatorname{IPARAM}(6)=0$, the Hessian is initialized to the identity matrix;
otherwise,
it is initialized to a diagonal matrix containing

$$
\max \left(|f(t)|, f_{s}\right) * s_{i}^{2}
$$

on the diagonal where $t=\mathrm{XGUESS}, f_{s}=\operatorname{FSCALE}$, and $s=\mathrm{XSCALE}$.
Default: 0
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: Not used in BCONF.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=$ XSCALE, and $f_{s}=$ FSCALE.

Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(\operatorname{STEPTL})$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE .
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: Not used in BCONF.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in BCONF.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: Not used in BCONF.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1} \sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=$ XSCALE, and $t=$ XGUESS.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is required, then DU4INF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The problem

$$
\begin{gathered}
\min f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2} \\
\text { subject to } \quad-2 \leq x_{1} \leq 0.5 \\
-1 \leq x_{2} \leq 2
\end{gathered}
$$

is solved with an initial guess $(-1.2,1.0)$ and default values for parameters.

```
    USE BCONF INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
    INTEGER IPARAM(7), ITP, L, NOUT
    REAL F, FSCALE, RPARAM(7), X(N), XGUESS (N), &
        XLB(N), XSCALE(N), XUB(N)
    EXTERNAL ROSBRK
    DATA XGUESS/-1.2E0, 1.0E0/
    DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
                All the bounds are provided
                Default parameters are used
                Minimize Rosenbrock function using
                initial guesses of -1.2 and 1.0
    CALL BCONF (ROSBRK, ITP, XLB, XUB, X, XGUESS=XGUESS, &
            iparam=iparam, FVALUE=F)
                                    Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L), L=3,5)
!
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ s o l u t i o n ~ i s ~ ' , ~ 6 X , ~ 2 F 8 . 3 , ~ / / , ~ ' ~ T h e ~ f u n c t i o n ~ ' , ~ \& ~
                'value is ', F8.3, //, ' The number of iterations is ', &
                10X, I3, /, ' The number of function evaluations is ', &
                I3, /, ' The number of gradient evaluations is ', I3)
    END
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
!
    RETURN
END
```


## Output

| The solution is | 0.500 | 0.250 |
| :--- | :--- | :--- |
| The function value is | 0.250 |  |

```
The number of iterations is 24
The number of function evaluations is 34
The number of gradient evaluations is 26
```


## BCONG

Minimizes a function of N variables subject to bounds on the variables using a quasi-Newton method and a user-supplied gradient.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
$X$ - Vector of length $N$ at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)

X - Vector of length N at which point the gradient is evaluated. (Input) X should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

IBTYPE Action

0

1 All variables are nonnegative.

2

3
User will supply all the bounds.

All variables are nonpositive.

User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables. (Input, if IBTYPE $=0$; output, if IBTYPE $=1$ or 2 ; input/output, if $\operatorname{IBTYPE}=3$ )
$\boldsymbol{X U B}-$ Vector of length N containing the upper bounds on variables. (Input, if $\operatorname{IBTYPE}=0$; output, if IBTYPE = 1 or 2 ; input/output, if IBTYPE $=3$ )
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing the initial guess of the minimum. (Input) Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0.
Default: $\mathrm{XSCALE}=1.0$.

FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.

IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: $\operatorname{IPARAM}=0$.

RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL BCONG (FCN, GRAD, IBTYPE, XLB, XUB, X [,...])
Specific: The specific interface names are S_BCONG and D_BCONG.

## FORTRAN 77 Interface

Single: CALL BCONG (FCN, GRAD, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DBCONG.

## Description

The routine BCONG uses a quasi-Newton method and an active set strategy to solve minimization problems subject to simple bounds on the variables. The problem is stated as follows:

$$
\begin{gathered}
\min _{x \in \mathbf{R}^{n}} f(x) \\
\text { subject to } l \leq x \leq u
\end{gathered}
$$

From a given starting point $x^{c}$, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-B^{-1} g^{c}
$$

where $B$ is a positive definite approximation of the Hessian and $g^{c}$ is the gradient evaluated at $x^{c}$; both are computed with respect to the free variables. The search direction for the variables in IA is set to zero. A line search is used to find a new point $x^{n}$,

$$
x^{n}=x^{c}+\lambda d, \quad \lambda \in(0,1]
$$

such that

$$
f\left(x^{n}\right) \leq f\left(x^{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality conditions

$$
\begin{gathered}
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{i}<x_{i}<u_{i} \\
g\left(x_{i}\right)<0, x_{i}=u_{i} \\
g\left(x_{i}\right)>0, x_{i}=l_{i}
\end{gathered}
$$

are checked, where $\varepsilon$ is a gradient tolerance. When optimality is not achieved, B is updated according to the BFGS formula:

$$
B \leftarrow B-\frac{B s s^{T} B}{s^{T} B s}+\frac{y y^{T}}{y^{T} s}
$$

where $s=x^{n}-x^{c}$ and $y=g^{n}-g^{c}$. Another search direction is then computed to begin the next iteration.

The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more details on the quasi-Newton method and line search, see Dennis and Schnabel (1983). For more detailed information on active set strategy, see Gill and Murray (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2ONG/DB2ONG. The reference is:
```
CALL B2ONG (FCN, GRAD, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM,
RPARAM, X, FVALUE, WK, IWK)
```

The additional arguments are as follows:
$\boldsymbol{W} \boldsymbol{K}$ - Real work vector of length $\mathrm{N} *(2 * \mathrm{~N}+8)$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain a BFGS approximation to the Hessian at the solution. Only the lower triangular portion of the matrix is stored in $\omega K$. The values returned in the upper triangle should be ignored.

IWK - Work vector of length N stored in column order.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. <br> 4 |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 6 | Maximum number of gradient evaluations exceeded. <br> Five consecutive steps have been taken with the maximum step <br> length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. <br> The last global step failed to locate a lower point than the current X <br> value. |
| 3 | 8 |  |

3. The first stopping criterion for BCONG occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM (1) ). The second stopping criterion for BCONG occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM (2) ).
4. If the default parameters are desired for BCONG, then set IPARAM (1) to zero and call the routine BCONG. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling BCONG:

CALL U4INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:

IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter.
If IPARAM $(6)=0$, the Hessian is initialized to the identity matrix;
otherwise, it is initialized to a diagonal matrix containing

$$
\max \left(|f(t)|, f_{s}\right) * s_{i}^{2}
$$

on the diagonal where $t=$ XGUESS, $f s=$ FSCALE, and $s=$ XSCALE .
Default: 0.
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: Not used in BCONG.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=\mathrm{XSCALE}$, and $f_{s}=\mathrm{FSCALE}$.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(\operatorname{STEPTL})$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision. $\operatorname{RPARAM}(3)=$ Relative function tolerance.

Default: Not used in BCONG. $\operatorname{RPARAM}(4)=$ Absolute function tolerance.

Default: Not used in BCONG.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: Not used in BCONG. $\operatorname{RPARAM}(6)=$ Maximum allowable step size.

Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1} \sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=\operatorname{XSCALE}$, and $t=\mathrm{XGUESS}$.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is required, then DU4 INF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The problem

$$
\begin{aligned}
& \min f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2} \\
& \text { subject to } \quad-2 \leq x_{1} \leq 0.5 \\
&-1 \leq x_{2} \leq 2
\end{aligned}
$$

is solved with an initial guess $(-1.2,1.0)$, and default values for parameters.

```
USE BCONG INT
USE UMACH INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=2)
INTEGER IPARAM(7), ITP, L, NOUT
REAL F, X(N), XGUESS(N), XLB(N), XUB(N)
EXTERNAL ROSBRK, ROSGRD
```

```
!
    DATA XGUESS/-1.2E0, 1.0E0/
    DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
        All the bounds are provided
        Default parameters are used
        Minimize Rosenbrock function using
        initial guesses of -1.2 and 1.0
    CALL BCONG (ROSBRK, ROSGRD, ITP, XLB, XUB, X, XGUESS=XGUESS, &
        IPARAM=IPARAM, FVALUE=F)
        Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5)
!
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
            'value is ', F8.3, //, ' The number of iterations is ', &
            10X, I3, /, ' The number of function evaluations is ', &
            I3, /, ' The number of gradient evaluations is ', I3)
        END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
!
    RETURN
    END
    SUBROUTINE ROSGRD (N, X, G)
    INTEGER N
    REAL X(N),G(N)
!
    G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
```


## Output

| The solution is | 0.500 | 0.250 |
| :--- | :--- | :--- | :--- |
| The function value is | 0.250 |  |
|  |  |  |
| The number of iterations is | 22 |  |
| The number of function evaluations is | 32 |  |
| The number of gradient evaluations is | 23 |  |

## BCODH

Minimizes a function of N variables subject to bounds on the variables using a modified Newton method and a finite-difference Hessian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
X - Vector of length N at which point the function is evaluated. (Input) X should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point X . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
$\mathrm{X}-$ Vector of length N at which point the gradient is evaluated. (Input) X should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

0 User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on the variables. (Input)
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on the variables. (Input)
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

XGUESS - Vector of length N containing the initial guess of the minimum. (Input) Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0 .
Default: XSCALE $=1.0$.

FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.
IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: $\operatorname{IPARAM}=0$.
RPARAM - Parameter vector of length 7. (Input/Output)
See Comment 4.
FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL BCODH (FCN, GRAD, IBTYPE, XLB, XUB, X [,...])
Specific: The specific interface names are S_BCODH and D_BCODH.

## FORTRAN 77 Interface

Single: CALL BCODH (FCN, GRAD, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DBCODH.

## Description

The routine BCODH uses a modified Newton method and an active set strategy to solve minimization problems subject to simple bounds on the variables. The problem is stated as

$$
\min _{x \in \mathbf{R}^{n}} f(x)
$$

subject to $l \leq x \leq u$
From a given starting point $x^{c}$, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-H^{-1} g^{c}
$$

where $H$ is the Hessian and $g^{c}$ is the gradient evaluated at $x^{c}$; both are computed with respect to the free variables. The search direction for the variables in IA is set to zero. A line search is used to find a new point $x^{n}$,

$$
x^{n}=x^{c}+\lambda d, \quad \lambda \in(0,1]
$$

such that

$$
f\left(x^{n}\right) \leq f\left(x^{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality conditions

$$
\begin{gathered}
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{i}<x_{i}<u_{i} \\
g\left(x_{i}\right)<0, x_{i}=u_{i} \\
g\left(x_{i}\right)>0, x_{i}=l_{i}
\end{gathered}
$$

are checked where $\varepsilon$ is a gradient tolerance. When optimality is not achieved, another search direction is computed to begin the next iteration. This process is repeated until the optimality criterion is met.

The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more details on the modified Newton method and line search, see Dennis and Schnabel (1983). For more detailed information on active set strategy, see Gill and Murray (1976).

Since a finite-difference method is used to estimate the Hessian for some single precision calculations, an inaccurate estimate of the Hessian may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact Hessian can be easily provided, routine BCOAH should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2ODH/DB2ODH. The reference is:
```
CALL B2ODH (FCN, GRAD, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM,
RPARAM, X, FVALUE,WK, IWK)
```

The additional arguments are as follows:
WK - Real work vector of length $N$ * ( $N+8$ ). WK contains the following information on output: The second N locations contain the last step taken. The
third $N$ locations contain the last Newton step. The fourth $N$ locations contain an estimate of the gradient at the solution. The final $N^{2}$ locations contain the Hessian at the approximate solution.

IWK - Integer work vector of length $N$.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 5 | Maximum number of gradient evaluations exceeded. |
| 4 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |
| 4 | 7 | Maximum number of Hessian evaluations exceeded. |

3. The first stopping criterion for BCODH occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for BCODH occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for BCODH, then set IPARAM(1) to zero and call the routine BCODH. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM; then the following steps should be taken before calling BCODH:

CALL U4 INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4 INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.

Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.
Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter.
Default: Not used in BCODH.
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: 100.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=$ XSCALE, and $f_{s}=$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in BCODH.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
RPARAM(6) = Maximum allowable step size.

Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1} \sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$$
\varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}, \text { and } t=\mathrm{XGUESS} .
$$

$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is required, then DU4INF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The problem

$$
\begin{aligned}
& \min f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2} \\
& \text { subject to } \quad-2 \leq x_{1} \leq 0.5 \\
&-1 \leq x_{2} \leq 2
\end{aligned}
$$

is solved with an initial guess $(-1.2,1.0)$, and default values for parameters.

```
    USE BCODH_INT
    USE UMACH_
    IMPLICIT NONE
    INTEGER N
    PARAMETER (N=2)
!
    INTEGER IP, IPARAM(7), L, NOUT
    REAL F, X(N), XGUESS (N), XLB (N), XUB (N)
    EXTERNAL ROSBRK, ROSGRD
!
    DATA XGUESS/-1.2E0, 1.0E0/
    DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
!
    IPARAM(1) = 0
    IP = 0
        Minimize Rosenbrock function using
        initial guesses of -1.2 and 1.0
    CALL BCODH (ROSBRK, ROSGRD, IP, XLB, XUB, X, XGUESS=XGUESS, &
        IPARAM=IPARAM, FVALUE=F)
                                Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5)
!
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
    'value is ', F8.3, //, ' The number of iterations is ', &
```

```
        10X, I3, /, ' The number of function evaluations is ', &
        I3, /, ' The number of gradient evaluations is ', I3)
!
    END
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
    F=1.0E2*(X(2)-X(1)*X(1))**2 + (1.0EO-X(1))**2
!
    RETURN
    END
    SUBROUTINE ROSGRD (N, X, G)
    INTEGER N
    REAL X(N), G(N)
!
    G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
```


## Output

| The solution is | 0.500 | 0.250 |  |
| :--- | :--- | :--- | :--- |
| The function value is | 0.250 |  |  |
|  |  |  |  |
| The number of iterations is |  |  |  |
| The number of function evaluations is | 26 |  |  |
| The number of gradient evaluations is | 18 |  |  |

## BCOAH

Minimizes a function of N variables subject to bounds on the variables using a modified Newton method and a user-supplied Hessian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
$X$ - Vector of length $N$ at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.

GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
X - Vector of length N at which point the gradient is evaluated. (Input) $X$ should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.
HESS - User-supplied subroutine to compute the Hessian at the point x . The usage is CALL HESS ( $\mathrm{N}, \mathrm{X}, \mathrm{H}, \mathrm{LDH}$ ), where

N - Length of X . (Input)
X - Vector of length N at which point the Hessian is evaluated. (Input) $X$ should not be changed by HESS.

H - The Hessian evaluated at the point X . (Output)
LDH - Leading dimension of $H$ exactly as specified in the dimension statement of the calling program. (Input)

HESS must be declared EXTERNAL in the calling program.
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

$0 \quad$ User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X} \boldsymbol{L B}$ - Vector of length N containing the lower bounds on the variables. (Input)
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on the variables. (Input)
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{x}, 1)$.
XGUESS - Vector of length N containing the initial guess. (Input) Default: xGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. In the absence of other information, set all entries to 1.0 .
Default: $\operatorname{XSCALE}=1.0$.
FSCALE - Scalar containing the function scaling. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set FSCALE to 1.0.
Default: FSCALE $=1.0$.
IPARAM - Parameter vector of length 7. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4. Default: IPARAM $=0$.

RPARAM — Parameter vector of length 7. (Input/Output)
See Comment 4.
FVALUE - Scalar containing the value of the function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL BCOAH (FCN, GRAD, HESS, IBTYPE, XLB, XUB, X [, ...])
Specific: The specific interface names are S_BCOAH and D_BCOAH.

## FORTRAN 77 Interface

Single: CALL BCOAH (FCN, GRAD, HESS, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVALUE)

Double: The double precision name is DВСОАн.

## Description

The routine $\operatorname{BCOAH}$ uses a modified Newton method and an active set strategy to solve minimization problems subject to simple bounds on the variables. The problem is stated as follows:

$$
\min _{x \in \mathbf{R}^{\prime \prime}} f(x)
$$

$$
\text { subject to } l \leq x \leq u
$$

From a given starting point $x^{c}$, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-H^{-1} g^{c}
$$

where $H$ is the Hessian and $g^{c}$ is the gradient evaluated at $x^{c}$; both are computed with respect to the free variables. The search direction for the variables in IA is set to zero. A line search is used to find a new point $x^{n}$,

$$
x^{n}=x^{c}+\lambda d, \quad \lambda \in(0,1]
$$

such that

$$
f\left(x^{n}\right) \leq f\left(x^{c}\right)+\alpha g^{T} d, \quad \alpha \in(0,0.5)
$$

Finally, the optimality conditions

$$
\begin{gathered}
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{i}<x_{i}<u_{i} \\
g\left(x_{i}\right)<0, x_{i}=u_{i} \\
g\left(x_{i}\right)>0, x_{i}=l_{i}
\end{gathered}
$$

are checked where $\varepsilon$ is a gradient tolerance. When optimality is not achieved, another search direction is computed to begin the next iteration. This process is repeated until the optimality criterion is met.
The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more details on the modified Newton method and line search, see Dennis and Schnabel (1983). For more detailed information on active set strategy, see Gill and Murray (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2OAH/DB2OAH. The reference is:
```
CALL B2OAH (FCN, GRAD, HESS, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE,
IPARAM, RPARAM, X, FVALUE, WK, IWK)
```

The additional arguments are as follows:
WK - Work vector of length $\mathrm{N} *(\mathrm{~N}+8)$. WK contains the following information on output: The second $N$ locations contain the last step taken. The third $N$ locations contain the last Newton step. The fourth N locations contain an estimate of the gradient at the solution. The final $\mathrm{N}^{2}$ locations contain the Hessian at the approximate solution.

IWK - Work vector of length N.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 4 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 4 | 5 | Maximum number of gradient evaluations exceeded. |
| 4 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |
| 2 | 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |
| 4 | 7 | Maximum number of Hessian evaluations exceeded. <br> The last global step failed to locate a lower point than the current X <br> value. |

3. The first stopping criterion for $В С О А Н$ occurs when the norm of the gradient is less than the given gradient tolerance (RPARAM(1)). The second stopping criterion for BCOAH occurs when the scaled distance between the last two steps is less than the step tolerance (RPARAM(2)).
4. If the default parameters are desired for $\operatorname{BCOAH}$, then set $\operatorname{IPARAM}(1)$ to zero and call the routine ВСОАН. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling BCOAH:

CALL U4INF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4 INF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 7.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of gradient evaluations.

Default: 400.
$\operatorname{IPARAM}(6)=$ Hessian initialization parameter.
Default: Not used in BCOAH.
$\operatorname{IPARAM}(7)=$ Maximum number of Hessian evaluations.
Default: 100.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\max \left(|f(x)|, f_{s}\right)}
$$

where $g=\nabla f(x), s=$ XSCALE, and $f_{s}=$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(\operatorname{STEPTL})$
The $i$-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: Not used in BCOAH.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1} \sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$$
\varepsilon_{2}=\|s\|_{2}, s=\mathrm{XSCALE}, \text { and } t=\mathrm{XGUESS} .
$$

$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is required, then DU4INF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The problem

$$
\begin{gathered}
\min f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2} \\
\text { subject to } \quad-2 \leq x_{1} \leq 0.5 \\
-1 \leq x_{2} \leq 2
\end{gathered}
$$

is solved with an initial guess $(-1.2,1.0)$, and default values for parameters.

```
USE BCOAH INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=2)
    INTEGER IP, IPARAM(7), L, NOUT
    REAL F, X(N), XGUESS(N), XLB(N), XUB(N)
    EXTERNAL ROSBRK, ROSGRD, ROSHES
    DATA XGUESS/-1.2E0, 1.0E0/
    DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
    IPARAM(1) = 0
    IP = 0
        Minimize Rosenbrock function using
        initial guesses of -1.2 and 1.0
    CALL BCOAH (ROSBRK, ROSGRD, ROSHES, IP, XLB, XUB, X, &
        XGUESS=XGUESS,IPARAM=IPARAM, FVALUE=F)
                        Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, F, (IPARAM(L),L=3,5), IPARAM(7)
99999 FORMAT (' The solution is ', 6X, 2F8.3, //, ' The function ', &
        'value is ', F8.3, //, ' The number of iterations is ', &
        10X, I3, /, ' The number of function evaluations is ', &
        I3, /, ' The number of gradient evaluations is ', I3, /, &
        ' The number of Hessian evaluations is ', I3)
END
```

$!$
$!$
!
!

```
!
    SUBROUTINE ROSBRK (N, X, F)
    INTEGER N
    REAL X(N), F
!
    F = 1.0E2*(X(2)-X(1)*X(1))**2 + (1.0E0-X(1))**2
!
    RETURN
    END
!
    SUBROUTINE ROSGRD (N, X, G)
    INTEGER N
    REAL X(N), G(N)
!
    G(1) = -4.0E2*(X(2)-X(1)*X(1))*X(1) - 2.0E0*(1.0E0-X(1))
    G(2) = 2.0E2*(X(2)-X(1)*X(1))
!
    RETURN
    END
!
    SUBROUTINE ROSHES (N, X, H, LDH)
    INTEGER N, LDH
    REAL X(N), H(LDH,N)
!
    H(1,1)= -4.0E2*X(2) + 1.2E3*X(1)*X(1) + 2.0E0
    H(2,1) = -4.0E2*X(1)
    H(1,2) = H (2,1)
    H(2,2) = 2.0E2
!
    RETURN
    END
```


## Output

| The solution is | $0.500 \quad 0.250$ |  |  |
| :--- | :--- | :--- | :--- |
| The function value is | 0.250 |  |  |
|  |  |  |  |
| The number of iterations is | 18 |  |  |
| The number of function evaluations is | 29 |  |  |
| The number of gradient evaluations is | 19 |  |  |
| The number of Hessian evaluations is | 18 |  |  |

## BCPOL

Minimizes a function of N variables subject to bounds on the variables using a direct search complex algorithm.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN ( $\mathrm{N}, \mathrm{X}, \mathrm{F}$ ), where

N - Length of X . (Input)
$X-V e c t o r$ of length $N$ at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

0 User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on the first, variable. All other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on the variables. (Input, if IBTYPE $=0$; output, if IBTYPE $=1$ or 2 ; input/output, if IBTYPE $=3$ )
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on the variables. (Input, if IBTYPE $=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ )
$\boldsymbol{X}$ - Real vector of length N containing the best estimate of the minimum found. (Output)

## Optional Arguments

$N$ - The number of variables. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XGUESS}, 1)$.

XGUESS - Real vector of length N that contains an initial guess to the minimum. (Input) Default: XGUESS $=0.0$.

FTOL - First convergence criterion. (Input)
The algorithm stops when a relative error in the function values is less than FTOL, i.e. when $(F($ worst $)-F($ best $))<$ FTOL * $(1+\operatorname{ABS}(F$ (best) $))$ where $F$ (worst) and $F$ (best) are the function values of the current worst and best point, respectively. Second convergence criterion. The algorithm stops when the standard deviation of the function values at the $2 * N$ current points is less than FTOL. If the subroutine terminates prematurely, try again with a smaller value FTOL.
Default: $\mathrm{FTOL}=1.0 \mathrm{e}-4$ for single and $1.0 \mathrm{~d}-8$ for double precision.

MAXFCN - On input, maximum allowed number of function evaluations. (Input/ Output) On output, actual number of function evaluations needed.
Default: $\mathrm{MAXFCN}=300$.

FVALUE - Function value at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL BCPOL (FCN, IBTYPE, XLB, XUB, X [,...])
Specific: The specific interface names are S_BCPOL and D_BCPOL.

## FORTRAN 77 Interface

Single: CALL BCPOL (FCN, N, XGUESS, IBTYPE, XLB, XUB, FTOL, MAXFCN, X, FVALUE)

Double: $\quad$ The double precision name is DBCPOL.

## Description

The routine $B C P O L$ uses the complex method to find a minimum point of a function of $n$ variables. The method is based on function comparison; no smoothness is assumed. It starts with $2 n$ points $x_{1}, x_{2}, \ldots, x_{2 n}$. At each iteration, a new point is generated to replace the worst point $x_{j}$, which has the largest function value among these $2 n$ points. The new point is constructed by the following formula:

$$
x_{k}=c+\alpha(c-x \mathrm{j})
$$

where

$$
c=\frac{1}{2 n-1} \sum_{i \neq j} x_{i}
$$

and $\alpha(\alpha>0)$ is the reflection coefficient.
When $x_{k}$ is a best point, that is, when $f\left(x_{k}\right) \leq f\left(x_{i}\right)$ for $i=1, \ldots, 2 n$, an expansion point is computed $x_{e}=c+\beta\left(x_{k}-c\right)$, where $\beta(\beta>1)$ is called the expansion coefficient. If the new point is a worst point, then the complex would be contracted to get a better new point. If the contraction step is unsuccessful, the complex is shrunk by moving the vertices halfway toward the current best point. Whenever the new point generated is beyond the bound, it will be set to the bound. This procedure is repeated until one of the following stopping criteria is satisfied:
Criterion 1:

$$
f_{\text {best }}-f_{\text {worst }} \leq \varepsilon \mathrm{f}\left(1 .+\mid f_{\text {best }}\right)
$$

Criterion 2:

$$
\sum_{i=1}^{2 n}\left(f_{i}-\frac{\sum_{j=1}^{2 n} f_{j}}{2 n}\right)^{2} \leq \varepsilon_{f}
$$

where $f_{i}=f\left(x_{i}\right), f_{j}=f\left(x_{j}\right)$, and $\varepsilon_{f}$ is a given tolerance. For a complete description, see Nelder and Mead (1965) or Gill et al. (1981).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2POL/DB2 POL. The reference is:

CALL B2POL (FCN, N, XGUESS, IBTYPE, XLB, XUB, FTOL, MAXFCN, X, FVALUE, WK)
The additional argument is:
$\boldsymbol{W K}$ - Real work vector of length 2 * $\mathrm{N}^{* * 2}+5$ * N
2. Informational error

| Type | Code |  |
| :---: | :---: | :---: |
| 3 | 1 |  |
|  | The maximum number of function evaluations is exceeded. |  |

3. Since $B C P O L$ uses only function-value information at each step to determine a new approximate minimum, it could be quite inefficient on smooth problems compared to other methods such as those implemented in routine BCONF, which takes into account derivative information at each iteration. Hence, routine BCPOL should only be used as a last resort. Briefly, a set of 2 * N points in an N -dimensional space is called a complex. The minimization process iterates by replacing the point with the largest function value by a new point with a smaller function value. The iteration continues until all the points cluster sufficiently close to a minimum.

## Example

The problem

$$
\begin{aligned}
& \min f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2} \\
& \text { subject to } \quad-2 \leq x_{1} \leq 0.5 \\
&-1 \leq x_{2} \leq 2
\end{aligned}
$$

is solved with an initial guess $(-1.2,1.0)$, and the solution is printed.

```
USE BCPOL_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=2)
INTEGER IBTYPE, K, NOUT
REAL FTOL, FVALUE, X(N), XGUESS(N), XLB(N), XUB(N)
EXTERNAL FCN
```

!

```
!
!
XUB = ( 0.5, 2.0)
DATA XGUESS/-1.2, 1.0/, XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
FTOL = 1.0E-5
IBTYPE = 0
!
CALL BCPOL (FCN, IBTYPE, XLB, XUB, X, xguess=xguess, ftol=ftol, &
fvalue=fvalue)
!
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) (X (K),K=1,N), FVALUE
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ b e s t ~ e s t i m a t e ~ f o r ~ t h e ~ m i n i m u m ~ v a l u e ~ o f ~ t h e ' , ~ / , ~ \& ~
    ' function is X = (', 2(2X,F4.2), ')', /, ' with ', &
    'function value FVALUE = ', E12.6)
!
END
SUBROUTINE FCN (N, X, F)
INTEGER N
REAL X(N), F
F=100.0*(X(2)-X(1)*X(1))**2 + (1.0-X(1))**2
RETURN
END
```


## Output

The best estimate for the minimum value of the
function is $X=(0.500 .25)$
with function value FVALUE $=0.250002 \mathrm{E}+00$

## BCLSF

Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN ( $\mathrm{M}, \mathrm{N}, \mathrm{X}, \mathrm{F}$ ), where

M - Length of F . (Input)
N - Length of X . (Input)
x - The point at which the function is evaluated. (Input) $x$ should not be changed by FCN.
$\mathrm{F}-$ The computed function at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{M}-$ Number of functions. (Input)
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

$0 \quad$ User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables. (Input, if $\operatorname{IBTYPE}=0$; output, if $\operatorname{IBTYPE}=1$ or 2 ; input/output, if $I B T Y P E=3$ )
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on variables. (Input, if IBTYPE $=0$; output, if $\operatorname{IBTYPE}=1$ or 2 ; input/output, if $\operatorname{IBTYPE}=3$ )
$\boldsymbol{X}-$ Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Number of variables. (Input)
$N$ must be less than or equal to $M$.
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing the initial guess. (Input)
Default: XGUESS $=0.0$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. By default, the values for XSCALE are set internally. See IPARAM(6) in Comment 4.

FSCALE - Vector of length m containing the diagonal scaling matrix for the functions. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set all entries to 1.0 .
Default: FSCALE $=1.0$.
IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4. Default: $\operatorname{IPARAM}=0$.

RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVEC - Vector of length M containing the residuals at the approximate solution. (Output)
$\boldsymbol{F} \boldsymbol{J A C}-\mathrm{m}$ by N matrix containing a finite difference approximate Jacobian at the approximate solution. (Output)

LDFJAC - Leading dimension of FJAC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFJAC $=$ SIZE (FJAC , 1).

## FORTRAN 90 Interface

Generic: CALL BCLSF (FCN, M, IBTYPE, XLB, XUB, X [,...])
Specific: The specific interface names are $S_{\_} B C L S F$ and D_BCLSF.

## FORTRAN 77 Interface

Single: CALL BCLSF (FCN, M, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC)

Double: The double precision name is DBCLSF.

## Description

The routine BCLSF uses a modified Levenberg-Marquardt method and an active set strategy to solve nonlinear least squares problems subject to simple bounds on the variables. The problem is stated as follows:

$$
\begin{gathered}
\min _{x \in \mathbf{R}^{n}} \frac{1}{2} F(x)^{T} F(x)=\frac{1}{2} \sum_{i=1}^{m} f_{i}(x)^{2} \\
\text { subject to } l \leq x \leq u
\end{gathered}
$$

where $m \geq n, F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{m}$, and $f_{i}(x)$ is the $i$-th component function of $F(x)$. From a given starting point, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-\left(J^{T} J+\mu I\right)^{-1} J^{T} F
$$

where $\mu$ is the Levenberg-Marquardt parameter, $F=F(x)$, and $J$ is the Jacobian with respect to the free variables. The search direction for the variables in IA is set to zero. The trust region approach discussed by Dennis and Schnabel (1983) is used to find the new point. Finally, the optimality conditions are checked. The conditions are

$$
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{i}<x_{i}<u_{i}
$$

$$
\begin{aligned}
g\left(x_{i}\right) & <0, x_{i}
\end{aligned}=u_{i} .
$$

where $\varepsilon$ is a gradient tolerance. This process is repeated until the optimality criterion is achieved.
The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more detail on the Levenberg-Marquardt method, see Levenberg (1944), or Marquardt (1963). For more detailed information on active set strategy, see Gill and Murray (1976).

Since a finite-difference method is used to estimate the Jacobian for some single precision calculations, an inaccurate estimate of the Jacobian may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact Jacobian can be easily provided, routine BCLSJ should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2LSF / DB2LSF. The reference is:

CALL B2LSF (FCN, M, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC, WK, IWK)

The additional arguments are as follows:
WK - Work vector of length $11 * N+3 * M-1$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Gauss-Newton step. The fourth N locations contain an estimate of the gradient at the solution.
$I W K$ - Work vector of length $2 * N$ containing the permutations used in the $Q R$ factorization of the Jacobian at the solution.
2. Informational errors

Type Code
31 Both the actual and predicted relative reductions in the function are less than or equal to the relative function convergence tolerance.
32 The iterates appear to be converging to a noncritical point.
43 Maximum number of iterations exceeded.
44 Maximum number of function evaluations exceeded.
36 Five consecutive steps have been taken with the maximum step length.
27 Scaled step tolerance satisfied; the current point may be an approximate local solution, or the algorithm is making very slow progress and is not near a solution, or STEPTL is too big.
3. The first stopping criterion for BCLSF occurs when the norm of the function is less than the absolute function tolerance. The second stopping criterion occurs when the norm of the scaled gradient is less than the given gradient tolerance. The third stopping criterion
for BCLSF occurs when the scaled distance between the last two steps is less than the step tolerance.
4. If the default parameters are desired for BCLSF, then set IPARAM(1) to zero and call the routine BCLSF. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling BCLSF:

CALL U4LSF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.
Note that the call to U4LSF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 6.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=$ Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: 100.
$\operatorname{IPARAM}(6)=$ Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values for XSCALE are set internally.
Default: 1.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\|F(x)\|_{2}^{2}}
$$

where

$$
g_{i}=\left(J(x)^{T} F(x)\right)_{i} *\left(f_{s}\right)_{i}^{2}
$$

$J(x)$ is the Jacobian, $s=$ XSCALE, and $f_{s}$ = FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $(S T E P T L)$
The i-th component of the scaled step between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision. $\operatorname{RPARAM}(3)=$ Relative function tolerance.

Default: $\max \left(10^{-10,} \varepsilon^{2 / 3),} \max \left(10^{-20}, \varepsilon^{2 / 3}\right)\right.$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: $\max \left(10^{-20}, \varepsilon^{2}\right), \max \left(10^{-40}, \varepsilon^{2}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step size.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=$ XSCALE, and $t=$ XGUESS.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is desired, then DU4LSF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to "Error Handling" in the Introduction.

## Example

The nonlinear least squares problem

$$
\min _{x \in \mathbf{R}^{2}} \frac{1}{2} \sum_{i=1}^{2} f_{i}(x)^{2}
$$

subject to $-2 \leq x_{1} \leq 0.5$

$$
-1 \leq x_{2} \leq 2
$$

where

$$
f_{1}(x)=10\left(x_{2}-x_{1}^{2}\right) \text { and } f_{2}(x)=\left(1-x_{1}\right)
$$

is solved with an initial guess $(-1.2,1.0)$ and default values for parameters.

```
    USE BCLSF INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER M, N
    PARAMETER (M=2, N=2)
    INTEGER IPARAM(7), ITP, NOUT
    REAL FSCALE (M), FVEC (M), X (N), XGUESS (N), XLB (N), XS (N), XUB (N)
    EXTERNAL ROSBCK
! Compute the least squares for the
    DATA XGUESS/-1.2E0, 1.0E0/
    DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
                                    All the bounds are provided
    ITP = 0
    IPARAM(1) = 0
    CALL BCLSF (ROSBCK, M, ITP, XLB, XUB, X, xguess=xguess, &
            iparam=iparam, fvec=fvec)
                                Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) X, FVEC, IPARAM(3), IPARAM(4)
!
99999 FORMAT (' The solution is ', 2F9.4, //, ' The function ', &
            'evaluated at the solution is ', /, 18X, 2F9.4, //, &
            ' The number of iterations is ', 10X, I3, /, ' The ', &
            'number of function evaluations is ', I3, /)
    END
!
    SUBROUTINE ROSBCK (M, N, X, F)
    INTEGER M, N
    REAL X(N), F(M)
    F(1) = 1.0E1*(X(2)-X(1)*X(1))
    F(2) = 1.0E0 - X(1)
    RETURN
    END
```


## Output

```
The solution is 0.5000 0.2500
The function evaluated at the solution is
0.0000 0.5000
```

```
The number of iterations is 15
```

The number of function evaluations is 20

## BCLSJ

Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (M, N, X, F), where

M - Length of F . (Input)
N - Length of X . (Input)

X - The point at which the function is evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{J A C}$ - User-supplied subroutine to evaluate the Jacobian at a point x . The usage is
CALL JAC (M, N, X, FJAC, LDFJAC), where
$M-$ Length of $F$. (Input)
N - Length of X . (Input)

X - The point at which the function is evaluated. (Input)
$x$ should not be changed by FCN.

FJAC - The computed M by N Jacobian at the point X . (Output)
LDFJAC - Leading dimension of FJAC. (Input)
JAC must be declared EXTERNAL in the calling program.
$\boldsymbol{M}$ - Number of functions. (Input)
IBTYPE - Scalar indicating the types of bounds on variables. (Input)
IBTYPE Action
$0 \quad$ User will supply all the bounds.

1 All variables are nonnegative.

2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables. (Input, if $\operatorname{IBTYPE}=0$; output, if $\operatorname{IBTYPE}=1$ or 2 ; input/output, if $\operatorname{IBTYPE}=3$ )
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on variables. (Input, if IBTYPE $=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ )
$\boldsymbol{X}$ - Vector of length N containing the approximate solution. (Output)

## Optional Arguments

$N$ - Number of variables. (Input)
N must be less than or equal to M .
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
XGUESS - Vector of length N containing the initial guess. (Input)
Default: XGUESS $=0.0$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
XSCALE is used mainly in scaling the gradient and the distance between two points. By default, the values for XSCALE are set internally. See IPARAM(6) in Comment 4.

FSCALE - Vector of length M containing the diagonal scaling matrix for the functions. (Input)
FSCALE is used mainly in scaling the gradient. In the absence of other information, set all entries to 1.0 .
Default: $\mathrm{FSCALE}=1.0$.

IPARAM - Parameter vector of length 6. (Input/Output)
Set IPARAM(1) to zero for default values of IPARAM and RPARAM. See Comment 4.
Default: IPARAM=0.
RPARAM - Parameter vector of length 7. (Input/Output) See Comment 4.

FVEC - Vector of length M containing the residuals at the approximate solution. (Output)

FJAC - m by n matrix containing a finite difference approximate Jacobian at the approximate solution. (Output)

LDFJAC - Leading dimension of FJAC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFJAC SIZE $=($ FJAC, 1$)$.

## FORTRAN 90 Interface

Generic: CALL BCLSJ (FCN, JAC, M, IBTYPE, XLB, XUB, X [, ...])
Specific: The specific interface names are S_BCLSJ and D_BCLSJ.

## FORTRAN 77 Interface

Single: CALL BCLSU (FCN, JAC, M, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE, IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC)

Double: The double precision name is DBCLSJ.

## Description

The routine BCLSJ uses a modified Levenberg-Marquardt method and an active set strategy to solve nonlinear least squares problems subject to simple bounds on the variables. The problem is stated as follows:

$$
\min _{x \in \mathbf{R}^{n}} \frac{1}{2} F(x)^{T} F(x)=\frac{1}{2} \sum_{i=1}^{m} f_{i}(x)^{2}
$$

$$
\text { subject to } l \leq x \leq u
$$

where $m \geq n, F: \mathbf{R}^{n} \rightarrow \mathbf{R}^{m}$, and $f_{i}(x)$ is the $i$-th component function of $F(x)$. From a given starting point, an active set IA, which contains the indices of the variables at their bounds, is built. A variable is called a "free variable" if it is not in the active set. The routine then computes the search direction for the free variables according to the formula

$$
d=-\left(J^{T} J+\mu I\right)^{-1} J^{T} F
$$

where is the Levenberg-Marquardt parameter, $F=F(x)$, and $J$ is the Jacobian with respect to the free variables. The search direction for the variables in IA is set to zero. The trust region approach discussed by Dennis and Schnabel (1983) is used to find the new point. Finally, the optimality conditions are checked. The conditions are

$$
\begin{gathered}
\left\|g\left(x_{i}\right)\right\| \leq \varepsilon, l_{t}<x_{t}<u_{t} \\
g\left(x_{t}\right)<0, x_{t}=u_{t} \\
g\left(x_{t}\right)>0, x_{t}=l_{t}
\end{gathered}
$$

where $\varepsilon$ is a gradient tolerance. This process is repeated until the optimality criterion is achieved.
The active set is changed only when a free variable hits its bounds during an iteration or the optimality condition is met for the free variables but not for all variables in IA, the active set. In the latter case, a variable that violates the optimality condition will be dropped out of IA. For more
detail on the Levenberg-Marquardt method, see Levenberg (1944) or Marquardt (1963). For more detailed information on active set strategy, see Gill and Murray (1976).

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2LSJ/DB2LSJ. The reference is:

CALL B2LSJ (FCN, JAC, M, N, XGUESS, IBTYPE, XLB, XUB, XSCALE, FSCALE,
IPARAM, RPARAM, X, FVEC, FJAC, LDFJAC, WK, IWK)
The additional arguments are as follows:
$\boldsymbol{W K}$ - Work vector of length $11 * \mathrm{~N}+3 * \mathrm{M}-1$. WK contains the following information on output: The second N locations contain the last step taken. The third N locations contain the last Gauss-Newton step. The fourth N locations contain an estimate of the gradient at the solution.

IWK - Work vector of length $2 *^{N}$ containing the permutations used in the $Q \mathrm{R}$ factorization of the Jacobian at the solution.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | Both the actual and predicted relative reductions in the function are <br> less than or equal to the relative function convergence tolerance. |
| 3 | 2 | The iterates appear to be converging to a noncritical point. |
| 4 | 3 | Maximum number of iterations exceeded. |
| 4 | 4 | Maximum number of function evaluations exceeded. |
| 3 | 6 | Five consecutive steps have been taken with the maximum step <br> length. |
| 4 | 5 | Maximum number of Jacobian evaluations exceeded. <br> 2 |
| 7 | Scaled step tolerance satisfied; the current point may be an <br> approximate local solution, or the algorithm is making very slow <br> progress and is not near a solution, or STEPTL is too big. |  |

3. The first stopping criterion for BCLSJ occurs when the norm of the function is less than the absolute function tolerance. The second stopping criterion occurs when the norm of the scaled gradient is less than the given gradient tolerance. The third stopping criterion for BCLSJ occurs when the scaled distance between the last two steps is less than the step tolerance.
4. If the default parameters are desired for BCLSJ, then set IPARAM(1) to zero and call the routine BCLSJ. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling BCLSJ:

CALL U4LSF (IPARAM, RPARAM)
Set nondefault values for desired IPARAM, RPARAM elements.

Note that the call to U4LSF will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.

The following is a list of the parameters and the default values:
IPARAM - Integer vector of length 6.
IPARAM(1) = Initialization flag.
I PARAM(2) = Number of good digits in the function.
Default: Machine dependent.
$\operatorname{IPARAM}(3)=$ Maximum number of iterations.
Default: 100.
$\operatorname{IPARAM}(4)=$ Maximum number of function evaluations.
Default: 400.
$\operatorname{IPARAM}(5)=$ Maximum number of Jacobian evaluations.
Default: 100.
$\operatorname{IPARAM}(6)=$ Internal variable scaling flag.
If $\operatorname{IPARAM}(6)=1$, then the values for XSCALE are set internally.
Default: 1.
RPARAM - Real vector of length 7.
$\operatorname{RPARAM}(1)=$ Scaled gradient tolerance.
The $i$-th component of the scaled gradient at x is calculated as

$$
\frac{\left|g_{i}\right| * \max \left(\left|x_{i}\right|, 1 / s_{i}\right)}{\|F(x)\|_{2}^{2}}
$$

where

$$
g_{i}=\left(J(x)^{T} F(x)\right)_{i} *\left(f_{s}\right)_{i}^{2}
$$

$J(x)$ is the Jacobian, $s=$ XSCALE, and $f_{s}$ FSCALE.
Default:

$$
\sqrt{\varepsilon}, \sqrt[3]{\varepsilon}
$$

in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(2)=$ Scaled step tolerance. $($ STEPTL $)$
The $i$-th component of the scaled step
between two points $x$ and $y$ is computed as

$$
\frac{\left|x_{i}-y_{i}\right|}{\max \left(\left|x_{i}\right|, 1 / s_{i}\right)}
$$

where $s=$ XSCALE.
Default: $\varepsilon^{2 / 3}$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(3)=$ Relative function tolerance.
Default: $\max \left(10^{-10}, \varepsilon^{2 / 3}\right), \max \left(10^{-20}, \varepsilon^{2 / 3}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(4)=$ Absolute function tolerance.
Default: $\max \left(10^{-20}, \varepsilon^{2}\right), \max \left(10^{-40}, \varepsilon^{2}\right)$ in double where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(5)=$ False convergence tolerance.
Default: $100 \varepsilon$ where $\varepsilon$ is the machine precision.
$\operatorname{RPARAM}(6)=$ Maximum allowable step SIZE.
Default: $1000 \max \left(\varepsilon_{1}, \varepsilon_{2}\right)$ where

$$
\varepsilon_{1}=\sqrt{\sum_{i=1}^{n}\left(s_{i} t_{i}\right)^{2}}
$$

$\varepsilon_{2}=\|s\|_{2}, s=$ XSCALE, and $t=$ XGUESS.
$\operatorname{RPARAM}(7)=$ Size of initial trust region radius.
Default: based on the initial scaled Cauchy step.
If double precision is desired, then DU4LSF is called and RPARAM is declared double precision.
5. Users wishing to override the default print/stop attributes associated with error messages issued by this routine are referred to ERROR HANDLING in the Introduction.

## Example

The nonlinear least squares problem

$$
\begin{array}{r}
\min _{x \in \mathbf{R}^{2}} \frac{1}{2} \sum_{i=1}^{2} f_{i}(x)^{2} \\
\text { subject to }-2 \leq x_{1} \leq 0.5 \\
-1 \leq x_{2} \leq 2
\end{array}
$$

where

$$
f_{1}(x)=10\left(x_{2}-x_{1}^{2}\right) \text { and } f_{2}(x)=\left(1-x_{1}\right)
$$

is solved with an initial guess $(-1.2,1.0)$ and default values for parameters.

```
USE BCLSJ INT
USE UMACH_INT
IMPLICIT NONE
Declaration of variables
PARAMETER (LDFJAC=2, M=2, N=2)
INTEGER IPARAM(7), ITP, NOUT
REAL FVEC (M), RPARAM(7), X(N), XGUESS (N), XLB (N), XUB (N)
EXTERNAL ROSBCK, ROSJAC
    Compute the least squares for the
    Rosenbrock function.
DATA XGUESS/-1.2E0, 1.0E0/
DATA XLB/-2.0E0, -1.0E0/, XUB/0.5E0, 2.0E0/
    All the bounds are provided
ITP = 0
IPARAM(1) = 0
CALL BCLSJ (ROSBCK,ROSJAC,M,ITP,XLB,XUB,X,XGUESS=XGUESS, &
    IPARAM=IPARAM, FVEC=FVEC)
                                    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) X, FVEC, IPARAM(3), IPARAM(4)
99999 FORMAT (' The solution is ', 2F9.4, //, ' The function ', &
    'evaluated at the solution is ', /, 18X, 2F9.4, //, &
    ' The number of iterations is ', 10X, I3, /, ' The ', &
    'number of function evaluations is ', I3, /)
END
SUBROUTINE ROSBCK (M, N, X, F)
INTEGER M, N
REAL X(N), F(M)
F(1) = 1.0E1* (X(2)-X(1)*X(1))
F(2) = 1.0E0 - X(1)
RETURN
END
SUBROUTINE ROSJAC (M, N, X, FJAC, LDFJAC)
INTEGER M, N, LDFJAC
REAL X(N), FJAC (LDFJAC,N)
FJAC (1,1) = -20.0E0*X(1)
FJAC (2,1) = -1.0E0
FJAC (1,2) = 10.0E0
FJAC (2, 2) = 0.0E0
RETURN
END
```


## Output

```
The solution is 0.5000 0.2500
The function evaluated at the solution is
0.0000 0.5000
The number of iterations is 13
The number of function evaluations is 21
```


## BCNLS

Solves a nonlinear least-squares problem subject to bounds on the variables and general linear constraints.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (M, N, X, F), where
M - Number of functions. (Input)
N - Number of variables. (Input)
X - Array of length N containing the point at which the function will be evaluated. (Input)
F - Array of length M containing the computed function at the point X . (Output)
The routine FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{M}$ - Number of functions. (Input)
$\boldsymbol{C}$ — MCON $\times \mathrm{N}$ matrix containing the coefficients of the MCON general linear constraints. (Input)
$\boldsymbol{B L}$ - Vector of length MCON containing the lower limit of the general constraints. (Input).
$\boldsymbol{B} \boldsymbol{U}$ - Vector of length MCON containing the upper limit of the general constraints. (Input).
IRTYPE - Vector of length MCON indicating the types of general constraints in the matrix C.
(Input)
Let $R(I)=C(I, 1) \star X(1)+\ldots+C(I, N) \star X(N)$. Then the value of IRTYPE(I) signifies the following:

IRTYPE(I) I-th CONSTRAINT
0
BL(I).EQ.R(I).EQ.BU(I)
1 R(I).LE.BU(I)
2 R(I).GE.BL(I)
3 BL(I).LE.R(I).LE.BU(I)
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables; if there is no lower bound on a variable, then 1.0 E 30 should be set as the lower bound. (Input)
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on variables; if there is no upper bound on a variable, then -1.0 E 30 should be set as the upper bound. (Input)
$\boldsymbol{X}-$ Vector of length N containing the approximate solution. (Output)

## Optional Arguments

N - Number of variables. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{C}, 2)$.
MCON - The number of general linear constraints for the system, not including simple bounds. (Input)
Default: $\mathrm{MCON}=\operatorname{SIZE}(\mathrm{C}, 1)$.
LDC - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
LDC must be at least MCON.
Default: LDC $=\operatorname{SIZE}(C, 1)$.

XGUESS - Vector of length N containing the initial guess. (Input)
Default: XGUESS $=0.0$.
RNORM - The Euclidean length of components of the function $f(x)$ after the approximate solution has been found. (Output).

ISTAT - Scalar indicating further information about the approximate solution X . (Output) See the Comments section for a description of the tolerances and the vectors IPARAM and RPARAM.

## ISTATMeaning

1 The function $f(x)$ has a length less than TOLF $=\operatorname{RPARAM}(1)$. This is the expected value for ISTAT when an actual zero value of $f(x)$ is anticipated.

2 The function $f(x)$ has reached a local minimum. This is the expected value for ISTAT when a nonzero value of $f(x)$ is anticipated.

3 A small change (absolute) was noted for the vector $x$. A full model problem step was taken. The condition for ISTAT $=2$ may also be satisfied, so that a minimum has been found. However, this test is made before the test for ISTAT $=2$.

4 A small change (relative) was noted for the vector $x$. A full model problem step was taken. The condition for ISTAT $=2$ may also be satisfied, so that a minimum has been found. However, this test is made before the test for ISTAT $=2$.

The number of terms in the quadratic model is being restricted by the amount of storage allowed for that purpose. It is suggested, but not required, that additional storage be given for the quadratic model parameters. This is accessed through the vector IPARAM, documented below.

Return for evaluation of function and Jacobian if reverse communication is desired. See the Comments below.

## FORTRAN 90 Interface

Generic: CALL BCNLS (FCN, M, C, BL, BU, IRTYPE, XLB, XUB, X [, ...])

Specific: The specific interface names are S_BCNLS and D_BCNLS.

## FORTRAN 77 Interface

Single: CALL BCNLS (FCN, M, N, MCON, C, LDC, BL, BU, IRTYPE, XLB, XUB, XGUESS, X, RNORM, ISTAT)

Double: The double precision name is DBCNLS.

## Description

The routine BCNLS solves the nonlinear least squares problem

$$
\min \sum_{i=1}^{m} f_{i}(x)^{2}
$$

subject to

$$
\begin{aligned}
b_{l} & \leq C x \leq b_{u} \\
x_{l} & \leq x \leq x_{u}
\end{aligned}
$$

BCNLS is based on the routine DQED by R.J. Hanson and F.T. Krogh. The section of BCNLS that approximates, using finite differences, the Jacobian of $f(x)$ is a modification of JACBF by D.E. Salane.

## Comments

1. Workspace may be explicitly provided, if desired, by use of B2NLS/DB2NLS. The reference is:

CALL B2NLS (FCN, M, N, MCON, C, LDC, BL, BU, IRTYPE, XLB, XUB, XGUESS, X, RNORM, ISTAT, IPARAM, RPARAM, JAC, F, FJ, LDFJ, IWORK, LIWORK, WORK, LWORK)

The additional arguments are as follows:
IPARAM - Integer vector of length six used to change certain default attributes of BCNLS. (Input).

If the default parameters are desired for B2NLS, set IPARAM(1) to zero. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, the following steps should be taken before calling B2NLS:

## CALL B7NLS (IPARAM, RPARAM)

Set nondefault values for IPARAM and RPARAM.
If double precision is being used, DB7NLS should be called instead. Following is a list of parameters and the default values.
$\operatorname{IPARAM}(1)=$ Initialization flag.
$\operatorname{IPARAM}(2)=\operatorname{ITMAX}$, the maximum number of iterations allowed.
Default: 75
$\operatorname{IPARAM}(3)=$ a flag that suppresses the use of the quadratic model in the inner loop. If set to one, then the quadratic model is never used. Otherwise use the quadratic model where appropriate. This option decreases the amount of workspace as well as the computing overhead required. A user may wish to determine if the application really requires the use of the quadratic model.

Default: 0
$\operatorname{IPARAM}(4)=$ NTERMS, one more than the maximum number of terms used in the quadratic model.
Default: 5
$\operatorname{IPARAM}(5)=\operatorname{RCSTAT}$, a flag that determines whether forward or reverse communication is used. If set to zero, forward communication through functions FCN and JAC is used. If set to one, reverse communication is used, and the dummy routines B10LS / DB10LS and B11LS/DB11LS may be used in place of FCN and JAC, respectively. When BCNLS returns with ISTAT $=6$, arrays F and FJ are filled with $f(x)$ and the Jacobian of $f(x)$, respectively. BCNLS is then called again.
Default: 0
$\operatorname{IPARAM}(6)=$ a flag that determines whether the analytic Jacobian, as supplied in JAC, is used, or if a finite difference approximation is computed. If set to zero, JAC is not accessed and finite differences are used. If set to one, JAC is used to compute the Jacobian.

Default: 0
RPARAM - Real vector of length 7 used to change certain default attributes of
BCNLS. (Input)
For the description of RPARAM, we make the following definitions:
FC current value of the length of $f(x)$
FB best value of length of $f(x)$
FL value of length of $f(x)$ at the previous step
PV predicted value of length of $f(x)$, after the step is taken, using the
approximating model $\varepsilon$ machine epsilon $=\operatorname{amach}(4)$.
The conditions $|F B-P V| \leq T O L S N R * F B$ and $|F C-P V| \leq T O L P * F B$ and $|F C-F L| \leq T O L S N R * F B$ together with taking a full model step, must be satisfied before the condition ISTAT $=2$ is returned. (Decreasing any of the values for TOLF, TOLD, TOLX, TOLSNR, or TOLP will likely increase the number of iterations required for convergence.)
$\operatorname{RPARAM}(1)=T O L F$, tolerance used for stopping when $F C \leq T O L F$.
Default: $\min (1 . \mathrm{E}-5, \sqrt{\varepsilon})$
RPARAM $(2)=$ TOLX, tolerance for stopping when change to $x$ values has length less than or equal to TOLX*length of $x$ values.
Default : $\min (1 . \mathrm{E}-5, \sqrt{\varepsilon})$
$\operatorname{RPARAM}(3)=T O L D$, tolerance for stopping when change to $x$ values has length less than or equal to TOLD.
Default : $\min (1 . \mathrm{E}-5, \sqrt{\varepsilon})$
$\operatorname{RPARAM}(4)=\operatorname{TOLSNR}$, tolerance used in stopping condition ISTAT $=2$.
Default: 1.E-5
$\operatorname{RPARAM}(5)=T O L P$, tolerance used in stopping condition ISTAT $=2$.
Default: 1.E-5
$\operatorname{RPARAM}(6)=\operatorname{TOLUSE}$, tolerance used to avoid values of $x$ in the quadratic model's interpolation of previous points. Decreasing this value may result in more terms being included in the quadratic model.
Default : $\sqrt{\varepsilon}$
$\operatorname{RPARAM}(7)=$ COND, largest condition number to allow when solving for the quadratic model coefficients. Increasing this value may result in more terms being included in the quadratic model.

Default: 30
$\boldsymbol{J A C}$ - User-supplied subroutine to evaluate the Jacobian. The usage is
CALL JAC (M, N, X, FJAC, LDFJAC), where
$M$ - Number of functions. (Input)
N - Number of variables. (Input)
$X$ - Array of length $N$ containing the point at which the Jacobian will be evaluated.
(Input)
FJAC - The computed $M \times N$ Jacobian at the point X . (Output)
LDFJAC - Leading dimension of the array FJAC. (Input)
The routine JAC must be declared EXTERNAL in the calling program.
$\boldsymbol{F}$ - Real vector of length N used to pass $f(x)$ if reverse communication (IPARAM(4)) is enabled. (Input)
$\boldsymbol{F} \boldsymbol{J}$ - Real array of size $\mathrm{M} \times \mathrm{N}$ used to store the Jacobian matrix of $f(x)$ if reverse communication (IPARAM(4)) is enabled. (Input) Specifically,

$$
F J(i, j)=\frac{\partial f_{i}}{\partial x_{j}}
$$

$\boldsymbol{L D F J}$ - Leading dimension of $F J$ exactly as specified in the dimension statement of the calling program. (Input)

IWORK - Integer work vector of length LIWORK.
LIWORK - Length of work vector IWORK. LIWORK must be at least $5 \mathrm{MCON}+12 \mathrm{~N}+47+\operatorname{MAX}(\mathrm{M}, \mathrm{N})$

WORK - Real work vector of length LWORK
LWORK - Length of work vector WORK. LWORK must be at least
$41 \mathrm{~N}+6 \mathrm{M}+11 \mathrm{MCON}+(\mathrm{M}+\mathrm{MCON})(\mathrm{N}+1)+\mathrm{NA}(\mathrm{NA}+7)+8 \mathrm{MAX}(\mathrm{M}, \mathrm{N})+$ 99.

Where NA $=\mathrm{MCON}+2 \mathrm{~N}+6$.
2. Informational errors
Type Code

31 The function $f(x)$ has reached a value that may be a local minimum. However, the bounds on the trust region defining the size of the step are being hit at each step. Thus, the situation is suspect. (Situations of this type can occur when the solution is at infinity at some of the components of the unknowns, $x$ ).
32 The model problem solver has noted a value for the linear or quadratic model problem residual vector length that is greater than or equal to the current value of the function, i.e. the Euclidean length of $f(x)$. This situation probably means that the evaluation of $f(x)$ has more uncertainty or noise than is possible to account for in the tolerances used to not a local minimum. The value of $x$ is suspect, but a minimum has probably been found.
33 More than ITMAX iterations were taken to obtain the solution. The value obtained for $x$ is suspect, although it is the best set of $x$ values that occurred in the entire computation. The value of ITMAX can be increased though the IPARAM vector.

## Example 1

This example finds the four variables $x_{1}, x_{2}, x_{3}, x_{4}$ that are in the model function

$$
h(t)=x_{1} e^{x_{2} t}+x_{3} e^{x_{4} t}
$$

There are values of $h(t)$ at five values of $t$.

$$
\begin{aligned}
& h(0.05)=2.206 \\
& h(0.1)=1.994 \\
& h(0.4)=1.35 \\
& h(0.5)=1.216 \\
& h(1.0)=0.7358
\end{aligned}
$$

There are also the constraints that $x_{2}, x_{4} \leq 0, x_{1}, x_{3} \geq 0$, and $x_{2}$ and $x_{4}$ must be separated by at least 0.05 . Nothing more about the values of the parameters is known so the initial guess is 0 .

```
USE BCNLS INT
USE UMACH INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER MCON, N
PARAMETER (MCON=1, N=4)
PARAMETER (M=5, LDC=MCON)
INTEGER IRTYPE (MCON), NOUT
REAL BL (MCON), C(MCON,N), RNORM, X(N), XLB(N), &
        XUB (N)
    SPECIFICATIONS FOR SUBROUTINES
    SPECIFICATIONS FOR FUNCTIONS
EXTERNAL FCN
CALL UMACH (2, NOUT)
Define the separation between x(2)
and x(4)
C(1,1) = 0.0
C(1,2) = 1.0
C(1,3) = 0.0
C(1,4) = -1.0
BL(1) = 0.05
IRTYPE(1) = 2
XLB(1) = 0.0
XLB(2) = 1.0E30
XLB(3) = 0.0
XLB(4) = 1.0E30
XUB(1) = -1.0E30
XUB (2) = 0.0
XUB(3) = -1.0E30
XUB(4) = 0.0
CALL BCNLS (FCN, M, C, BL, BL, IRTYPE, XLB, XUB, X, RNORM=RNORM)
CALL WRRRN ('X', X, 1, N, 1)
```

```
    WRITE (NOUT,99999) RNORM
99999 FORMAT (/, 'rnorm = ', E10.5)
    END
!
    SUBROUTINE FCN (M, N, X, F)
! SPECIFICATIONS FOR ARGUMENTS
    INTEGER M, N
    REAL X(*), F(*)
    INTEGER I
    REAL H(5), T(5)
    SAVE H, T
    INTRINSIC EXP
    REAL EXP
!
    DATA T/0.05, 0.1, 0.4, 0.5, 1.0/
    DATA H/2.206, 1.994, 1.35, 1.216, 0.7358/
!
    DO 10 I=1, M
        F(I) = X(1)*EXP(X(2)*T(I)) + X(3)*EXP(X(4)*T(I)) - H(I)
    1 0 ~ C O N T I N U E ~
        RETURN
        END
```


## Output

```
                X
    rrrrr
rnorm = .42425E-03
```


## Additional Examples

## Example 2

This example solves the same problem as the last example, but reverse communication is used to evaluate $f(x)$ and the Jacobian of $f(x)$. The use of the quadratic model is turned off.

```
USE B2NLS INT
USE UMACH_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDC, LDFJ, M, MCON, N
PARAMETER (M=5, MCON=1, N=4, LDC=MCON, LDFJ=M)
                                    Specifications for local variables
INTEGER I, IPARAM(6), IRTYPE(MCON), ISTAT, IWORK(1000), &
    LIWORK, LWORK, NOUT
REAL BL(MCON), C(MCON,N), F(M), FJ(M,N), RNORM, RPARAM(7), &
    WORK(1000), X(N), XGUESS(N), XLB(N), XUB(N)
REAL H(5), T(5)
SAVE H, T
```

```
    INTRINSIC EXP
    REAL EXP
                                    Specifications for subroutines
                                    Specifications for functions
    EXTERNAL B10LS, B11LS
    DATA T/0.05, 0.1, 0.4, 0.5, 1.0/
    DATA H/2.206, 1.994, 1.35, 1.216, 0.7358/
    CALL UMACH (2, NOUT)
        Define the separation between x(2)
        and x(4)
    C(1,1) = 0.0
    C(1,2) = 1.0
    C(1,3) = 0.0
    C(1,4) = -1.0
    BL(1) = 0.05
    IRTYPE(1) = 2
!
    XLB(1) = 0.0
    XLB(2) = 1.0E30
    XLB(3) = 0.0
    XLB(4) = 1.0E30
! XUB (1) = -1.0E30
    XUB(1) = -1.0E30
    XUB(2) = 0.0
    XUB(3) = -1.0E30
    XUB(4) = 0.0
! Set initial guess to 0.0
    XGUESS = 0.0E0
!
    CALL B7NLS (IPARAM, RPARAM)
!
IPARAM(3) = 1
    IPARAM(5) = 1
    IPARAM(6) = 1
    LWORK = 1000
    LIWORK = 1000
!
    10 CONTINUE
    CALL B2NLS (B10LS, M, N, MCON, C, LDC, BL, BL, IRTYPE, XLB, &
                XUB, XGUESS, X, RNORM, ISTAT, IPARAM, RPARAM, &
                B11LS, F, FJ, LDFJ, IWORK, LIWORK, WORK, LWORK)
    Specify dummy routines for FCN
    and JAC since we are using reverse
    communication
!
!
                                    Evaluate functions if the routine
                                    returns with ISTAT = 6
    IF (ISTAT .EQ. 6) THEN
    DO 20 I=1, M
        FJ(I,1) = EXP(X(2)*T(I))
        FJ(I,2) = T(I)*X(1)*FJ(I,1)
```

```
        FJ(I, 3) = EXP(X(4)*T(I))
        FJ(I,4) = T(I)*X(3)*FJ(I,3)
                F(I) = X(1)*FJ(I,1) + X(3)*FJ(I, 3) - H(I)
        CONTINUE
        GO TO 10
    END IF
!
    CALL WRRRN ('X', X, 1, N, 1)
    WRITE (NOUT,99999) RNORM
99999 FORMAT (/, 'rnorm = ', E10.5)
END
```


## Output

|  | X | 2 | 4 | 4 |
| ---: | ---: | ---: | ---: | ---: |
| 1.999 | -1.000 | 0.500 | -9.954 |  |
| rnorm $=.42413 \mathrm{E}-03$ |  |  |  |  |

## READ_MPS

This subroutine reads an MPS file containing a linear programming problem or a quadratic programming problem.

## Required Arguments

FILENAME - Character string containing the name of the MPS file to be read. (Input)
MPS-A A structure of IMSL defined derived type s_MPS containing the data read from the MPS file. (Output)

The IMSL defined derived type s_MPS consists of the following components:

| Component | Description |
| :--- | :--- |
| character, allocatable :: filename | Name of the MPS file. |
| character $($ len $=8$ ) name | Name of the problem. |
| integer nrows | Number of rows in the constraint matrix. |
| integer ncolumns | Number of columns in the constraint <br> matrix. This is also the number of variables. |
| integer nonzeros | Number of non-zeros in the constraint <br> matrix. |
| integer nhessian | Number of non-zeros in the Hessian matrix. <br> If zero, then there is no Hessian matrix. |
| integer ninteger | Number of variables required to be integer. <br> This includes binary variables. |
| integer nbinary | Number of variables required to be binary |


| Component | Description |  |
| :---: | :---: | :---: |
|  | (0 or 1). |  |
| real (kind(1e0)), allocatable :: objective ( ) | A real array of length ncolumns containing the objective vector. |  |
| type (s_SparseMatrixElement), allocatable :: constraint(:) | A derived type array of length nonzeros and of type s_SparseMatrixElement containing the sparse matrix representation of the constraint matrix. See below for details. |  |
| type(s_SparseMatrixElement), allocatable :: hessian(:) | A derived type array of length nhessian and of type s_SparseMatrixElement containing the sparse matrix representation of the Hessian matrix. If nhessian is zero, then this field is not allocated. |  |
| real ( $\operatorname{kind}(1 \mathrm{e} 0))$, allocatable :: 1 ower_range ( : ) | A real array of length nrows containing the lower constraint bounds. If a constraint is unbounded below, the corresponding entry in lower_range is set to negative_infinity, defined below. |  |
| real (kind(1e0)), allocatable ::upper_range ( : ) | A real array of length nrows containing the upper constraint bounds. If a constraint is unbounded above, the corresponding entry in upper_range is set to positive infinity, defined below. |  |
| real $(\operatorname{kind}(1 \mathrm{e} 0))$, allocatable :: lower_bound ( : ) | A real array of length ncolumns containing the lower variable bounds. If a variable is unbounded below, the corresponding entry in lower_bound is set to negative_infinity, defined below. |  |
| real $(\operatorname{kind}(1 \mathrm{e} 0))$, allocatable :: upper_bound (:) | A real array of length ncolumns containing the upper variable bounds. If a variable is unbounded above, the corresponding entry in upper_bound is set to positive infinity, defined below. |  |
| integer, allocatable :: variable_type (:) | An integer array of length ncolumns containing the type of each variable. Variable types are: |  |
|  | 0 | Continous |
|  | 1 | Integer |
|  | 2 | Binary (0 or 1) |
|  | 3 | Semicontinuous |


| Component | Description |
| :---: | :---: |
| character (len=8) name_objective | Name of the set in ROWS used for the objective row. |
| character (len=8) name_rhs | Name of the RHS set used. |
| character (len=8) name_ranges | Name of the RANGES set used or the empty string if no RANGES section in the file. |
| character (len=8) name_bounds | Name of the BOUNDS set used or the empty string if no BOUNDS section in the file. |
| character (len=8), allocatable :: name_row ( ) | Array of length nrows containing the row names. The name of the $i$-th constraint row is name row(i). |
| character (len=8), allocatable :: name_column ( : | Array of length ncolumns containing the column names. The name of the $i$-th column and variable is name column(i). |
| real (kind (le0)) positive_infinity | Value used for a constraint or bound upper limit when the constraint or bound is unbounded above. This can be set using an optional argument. Default is $1.0 \mathrm{e}+30$. |
| real (kind (le0)) negative_infinity | Value used for a constraint or bound lower limit when the constraint or bound is unbounded below. This can be set using an optional argument. Default is $-1.0 \mathrm{e}+30$. |

This derived type stores the constraint and Hessian matrices in a simple sparse matrix format of derived type s_SparseMatrixElement defined in the interface module mp_types. s_SparseMatrixElement consists of three components; a row index, a column index, and a value. For each non-zero element in the constraint and Hessian matrices an element of derived type s_SparseMatrixElement is stored. The following code fragment expands the sparse constraint matrix of the derived type s_SparseMatrixElement contained in mps, a derived type of type s_MPS, into a dense matrix:

```
! allocate a matrix
integer nr = mps%nrows
integer nc = mps%ncolumns
real (kind(1e0)), allocatable :: matrix(:,:)
allocate(matrix(nr,nc))
matrix = 0.0e0
! expand the sparse matrix
do k = 1, mps%nonzeros
    i = mps%constraint(k) %row
    j =mps%constraint(k) %column
    matrix(i,j) = mps%constraint(k) %value
end do
```

The IMSL derived type d_MPS is the double precision counterpart to s_MPS. The IMSL derived type d_SparseMatrixElement is the double precision counterpart to s_SparseMatrixElement.

To release the space allocated for this derived type use the following statement:

```
call mps_free(mps)
```


## Optional Arguments

NUNIT - The unit number for reading an MPS file opened by the user. If NUNIT is not used, this subroutine opens the file indicated by FILENAME for reading and then closes it after reading. (Input) By default, 7 is used.
$\boldsymbol{O B J}$ - Character string of length 8 containing the name of the objective function set to be used. (Input)
An MPS file can contain multiple objective function sets.
By default, the first objective function set in the MPS file is used. This name is case sensitive.

RHS - Character string of length 8 containing the name of the RHS set to be used. (Input) An MPS file can contain multiple RHS sets. By default, the first RHS set in the MPS file is used. This name is case sensitive.

RANGES - Character string of length 8 containing the name of the RANGES set to be used. (Input)
An MPS file can contain multiple RANGES sets.
By default, the first RANGES set in the MPS file is used. This name is case sensitive.
BOUNDS - Character string of length 8 containing the name of the BOUNDS set to be used. (Input)
An MPS file can contain multiple BOUNDS sets.
By default, the first BOUNDS set in the MPS file is used. This name is case sensitive.
POS_INF - Value used for a constraint or bound upper limit when the constraint or bound is unbounded above. (Input)
Default: $1.0 \mathrm{e}+30$.
NEG_INF - Value used for a constraint or bound lower limit when the constraint or bound is unbounded below. (Input)
Default: $-1.0 \mathrm{e}+30$.

## FORTRAN 90 Interface

Generic: CALL READ_MPS (FILENAME, MPS [,...])
Specific: The specific interface names are S_READ_MPS and D_READ_MPS.

## Description

An MPS file defines a linear or quadratic programming problem.
A linear programming problem is assumed to have the form:

$$
\begin{gathered}
\min _{x} c^{T} x \\
b_{l} \leq A x \leq b_{u} \\
x_{l} \leq x \leq x_{u}
\end{gathered}
$$

A quadratic programming problem is assumed to have the form:

$$
\begin{gathered}
\min _{x} \frac{1}{2} x^{T} Q x+c^{T} x \\
b_{l} \leq A x \leq b_{u} \\
x_{l} \leq x \leq x_{u}
\end{gathered}
$$

The following table maps this notation into the components in the derived type returned by READ_MPS:

| $C$ | Objective |
| :--- | :--- |
| $A$ | Constraint |
| $Q$ | Hessian |
| $b_{l}$ | lower_range |
| $b_{u}$ | upper_range |
| $x_{l}$ | lower_bound |
| $x_{u}$ | upper_bound |

If the MPS file specifies an equality constraint or bound, the corresponding lower and upper values in the returned derived type will be exactly equal.

The problem formulation assumes that the constraints and bounds are two-sided. If a particular constraint or bound has no lower limit, then the corresponding component of the derived type is set to $-1.0 \mathrm{e}+30$. If the upper limit is missing, then the corresponding component of the derived type is set to $+1.0 \mathrm{e}+30$.

## MPS File Format

There is some variability in the MPS format. This section describes the MPS format accepted by this reader.

An MPS file consists of a number of sections. Each section begins with a name in column 1. With the exception of the NAME section, the rest of this line is ignored. Lines with a '*' or ' $\$$ ' in column 1 are considered comment lines and are ignored.

The body of each section consists of lines divided into fields, as follows:

| Field Number | Columns | Contents |
| :---: | :--- | :--- |
| 1 | $2-3$ | Indicator |
| 2 | $5-12$ | Name |
| 3 | $15-22$ | Name |
| 4 | $25-36$ | Value |
| 5 | $40-47$ | Name |
| 6 | $50-61$ | Value |

The format limits MPS names to 8 characters and values to 12 characters. The names in fields 2,3 and 5 are case sensitive. Leading and trailing blanks are ignored, but internal spaces are significant.
The sections in an MPS file are as follows.

- NAME
- ROWS
- COLUMNS
- RHS
- RANGES (optional)
- BOUNDS (optional)
- QUADRATIC (optional)
- ENDATA

Sections must occur in the above order.
MPS keywords, section names and indicator values, are case insensitive. Row, column and set names are case sensitive.

## NAME Section

The NAME section contains a single line. A problem name can occur anywhere on the line after NAME and before column 62 . The problem name is truncated to 8 characters.

## ROWS Section

The ROWS section defines the name and type for each row. Field 1 contains the row type and field 2 contains the row name. Row type values are not case sensitive. Row names are case sensitive. The following row types are allowed:

| Row Type | Meaning |
| :---: | :--- |
| E | Equality Constraint. |
| L | Less than or equal constraint |


| Row Type | Meaning |
| :---: | :--- |
| G | Greater than or equal constraint. |
| N | Objective or a free row. |

## COLUMNS Section

The COLUMNS section defines the nonzero entries in the objective and the constraint matrix. The row names here must have been defined in the ROWS section.

| Field | Contents |
| :---: | :--- |
| 2 | Column name. |
| 3 | Row name. |
| 4 | Value for the entry whose row and column <br> are given by fields 3 and 2. |
| 5 | Row name. |
| 6 | Value for the entry whose row and column <br> are given by fields 5 and 2. |

NOTE: Fields 5 and 6 are optional.
The COLUMNS section can also contain markers. These are indicated by the name 'MARKER' (with the quotes) in field 3 and the marker type in field 4 or 5.

Marker type 'INTORG' (with the quotes) begins an integer group. The marker type 'INTEND' (with the quotes) ends this group. The variables corresponding to the columns defined within this group are required to be integer.

## RHS Section

The RHS section defines the right-hand side of the constraints. An MPS file can contain more than one RHS set, distinguished by the RHS set name. The row names here must be defined in the ROWS section.

| Field | Contents |
| :---: | :--- |
| 2 | RHS set name. |
| 3 | Row name. |
| 4 | Value for the entry whose set and row are <br> given by fields 2 and 3. |
| 5 | Row name. |
| 6 | Value for the entry whose set and row are <br> given by fields 2 and 5. |

NOTE: Fields 5 and 6 are optional.

## RANGES Section

The optional RANGES section defines two-sided constraints. An MPS file can contain more than one range set, distinguished by the range set name. The row names here must have been defined in the ROWS section.

| Field | Contents |
| :---: | :--- |
| 2 | Range set name. |
| 3 | Row name. |
| 4 | Value for the entry whose set and row are <br> given by fields 2 and 3. |
| 5 | Row name. |
| 6 | Value for the entry whose set and row are <br> given by fields 2 and 5. |

NOTE: Fields 5 and 6 are optional.
Ranges change one-sided constraints, defined in the RHS section, into two-sided constraints. The two-sided constraint for row $i$ depends on the range value, $r_{i}$, defined in this section. The righthand side value, $b_{i}$, is defined in the RHS section. The two-sided constraints for row $i$ are given in the following table:

| Row Type | Lower Constraint | Upper Constraint |
| :---: | :--- | :--- |
| G | $b_{i}$ | $b_{i}+\left\|r_{i}\right\|$ |
| L | $b_{i}-\left\|r_{i}\right\|$ | $b_{i}$ |
| E | $b_{i}+\min \left(0, r_{i}\right)$ | $b_{i}+\max \left(0, r_{i}\right)$ |

## BOUNDS Section

The optional BOUNDS section defines bounds on the variables. By default, the bounds are $0 \leq x_{i} \leq \infty$. The bounds can also be used to indicate that a variable must be an integer.

More than one bound can be set for a single variable. For example, to set $2 \leq x_{i} \leq 6$ use a LO bound with value 2 to set $2 \leq x_{i}$ and a UP bound with value 6 to add the condition $x_{i} \leq 6$.
An MPS file can contain more than one bounds set, distinguished by the bound set name.

| Field | Contents |
| :---: | :--- |
| 1 | Bounds type. |
| 2 | Bounds set name. |
| 3 | Column name |
| 4 | Value for the entry whose set and column are <br> given by fields 2 and 3. |
| 5 | Column name. |
| 6 | Value for the entry whose set and column are <br> given by fields 2 and 5. |

NOTE: Fields 5 and 6 are optional.
The bound types are as follows. Here $b_{i}$ are the bound values defined in this section, the $x_{i}$ are the variables, and $I$ is the set of integers.

| Bounded Type | Definition | Formula |
| :---: | :--- | :--- |
| LO | Lower bound | $b_{j} \leq x_{i}$ |
| UP | Upper bound | $x_{i} \leq b_{i}$ |
| FX | Fixed variable | $x_{i}=b_{i}$ |
| FR | Free variable | $-\infty \leq x_{i} \leq \infty$ |
| MI | Lower bound is minus <br> infinity | $-\infty \leq x_{i}$ |
| PL | Upper bound is positive <br> infinity | $x_{i} \leq \infty$ |
| BV | Binary variable (variable <br> must be 0 or 1$).$ | $x_{i} \in\{0,1\}$ |
| LI | Upper bound and integer | $x_{i} \leq b_{i}$ and $x_{i} \in I$ |
| SC | Lower bound and <br> integer | $b_{i} \leq x_{i}$ and $x_{i} \in I$ |
|  | Semicontinuous | 0 or $b_{i} \leq x_{i}$ |

The bound type names are not case sensitive.
If the bound type is UP or UI and $b_{j}<0$ then the lower bound is set to $-\infty$.

## QUADRATIC Section

The optional QUADRATIC section defines the Hessian for quadratic programming problems. The names HESSIAN, QUADS, QUADOBJ, QSECTION, and QMATRIX are also recognized as beginning the QUADRATIC section.

| Field | Contents |
| :---: | :--- |
| 2 | Column name. |
| 3 | Column name. |
| 4 | Value for the entry specified by fields 2 and 3. |
| 5 | Column name. |
| 6 | Value for the entry specified by fields 2 and 5. |

NOTE: Fields 5 and 6 are optional.

## ENDATA Section

The ENDATA section ends the MPS file.

## Comments

Informational errors

| Type | Code |  |
| :--- | :--- | :--- |
| 3 | 5 | No objective coefficients found. |
| 3 | 6 | No RHS values found. |
| 3 | 8 | No range values found. |
| 3 | 9 | No bounds found. |
| 4 | 3 | Missing section title. |
| 4 | 4 | Error reading input file. |
| 4 | 7 | Invalid number. |
| 4 | 11 | Unexpected section header. |
| 4 | 12 | Unknown row type. |
| 4 | 13 | Out-of-order marker. |
| 4 | 14 | Unknown marker type. |
| 4 | 15 | Unknown column name. |
| 4 | 16 | Unknown bound type. |
| 4 | 17 | Unknown row name. |
| 4 | 18 | Unexpected section name. |

## Example 1

```
use read_mps_int
implicit none
TYPE (S_MPS) mps
CALL read_mps ('test.mps', mps)
End
```


## Additional Examples

## Example 2

See Example 2 of DENSE_LP.

## MPS_FREE

Deallocates the space allocated for the IMSL derived type s_MPS. This routine is usually used in conjunction with READ_MPS.

## Required Arguments

MPS - A structure of IMSL defined derived type s_MPS containing the data read from the MPS file. (Input/Output)
The allocated components of s_MPS will be deallocated on output.
The IMSL defined derived type s_MPS consists of the following components:

| Component | Description |
| :--- | :--- |
| character, allocatable :: filename | Name of the MPS file. |
| character $($ len $=8$ ) name | Name of the problem. |
| integer nrows | Number of rows in the constraint matrix. |
| integer ncolumns | Number of columns in the constraint <br> matrix. This is also the number of variables. |
| integer nonzeros | Number of non-zeros in the constraint <br> matrix. |
| integer nhessian | Number of non-zeros in the Hessian matrix. <br> If zero, then there is no Hessian matrix. |
| integer ninteger | Number of variables required to be integer. <br> This includes binary variables. |
| integer nbinary | Number of variables required to be binary <br> $(0$ or 1$)$. |
| real (kind(1e0)), allocatable :: objective (: ) | A real array of length ncolumns |


| Component | Description |
| :---: | :---: |
|  | containing the objective vector. |
| type (s_SparseMatrixElement), allocatable :: constraint(:) | A derived type array of length nonzeros and of type s_SparseMatrixElement containing the sparse matrix representation of the constraint matrix. See below for details. |
| type(s_SparseMatrixElement), allocatable :: hessian(:) | A derived type array of length nhessian and of type s_SparseMatrixElement containing the sparse matrix representation of the Hessian matrix. If nhessian is zero, then this field is not allocated. |
| real (kind(1e0)), allocatable ::lower_range (:) | A real array of length nrows containing the lower constraint bounds. If a constraint is unbounded below, the corresponding entry in lower_range is set to negative infinity, defined below. |
| real (kind(1e0)), allocatable ::upper_range ( : ) | A real array of length nrows containing the upper constraint bounds. If a constraint is unbounded above, the corresponding entry in upper_range is set to positive infinity, defined below. |
| real (kind(1e0)), allocatable :: lower_bound ( : ) | A real array of length ncolumns containing the lower variable bounds. If a variable is unbounded below, the corresponding entry in lower_bound is set to negative infinity, defined below. |
| real $(\operatorname{kind}(1 \mathrm{e} 0))$, allocatable :: upper_bound (:) | A real array of length ncolumns containing the upper variable bounds. If a variable is unbounded above, the corresponding entry in upper_bound is set to positive infinity, defined below. |
| integer, allocatable :: variable_type (:) | An integer array of length ncolumns containing the type of each variable. Variable types are: |
|  | 0 Continous |
|  | 1 Integer |
|  | 2 Binary (0 or 1) |
|  | 3 Semicontinuous |
| character (len=8) name_objective | Name of the set in ROWS used for the objective row. |


| Component | Description |
| :---: | :---: |
| character (len=8) name_rhs | Name of the RHS set used. |
| character (len=8) name_ranges | Name of the RANGES set used or the empty string if no RANGES section in the file. |
| character (len=8) name_bounds | Name of the BOUNDS set used or the empty string if no BOUNDS section in the file. |
| character (len=8), allocatable :: name_row ( ) | Array of length nrows containing the row names. The name of the $i$-th constraint row is name row(i). |
| character (len=8), allocatable :: name_column ( ) | Array of length ncolumns containing the column names. The name of the $i$-th column and variable is name column(i). |
| real (kind (le0)) positive_infinity | Value used for a constraint or bound upper limit when the constraint or bound is unbounded above. This can be set using an optional argument. Default is $1.0 \mathrm{e}+30$. |
| real (kind (le0)) negative_infinity | Value used for a constraint or bound lower limit when the constraint or bound is unbounded below. This can be set using an optional argument. Default is $-1.0 \mathrm{e}+30$. |

This derived type stores the constraint and Hessian matrices in a simple sparse matrix format of derived type s_SparseMatrixElement defined in the interface module mp_types.
s_SparseMatrixElement consists of three components; a row index, a column index, and a value. For each non-zero element in the constraint and Hessian matrices an element of derived type s_SparseMatrixElement is stored The following code fragment expands the sparse constraint matrix of the derived type s_SparseMatrixElement contained in mps, a derived type of type s_MPS, into a dense matrix:

```
! allocate a matrix
integer nr = mps%nrows
integer nc = mps%ncolumns
real (kind(le0)), allocatable :: matrix(:,:)
allocate(matrix(nr,nc))
matrix = 0.0e0
! expand the sparse matrix
do k = 1, mps%nonzeros
    i = mps%constraint(k)%row
    j = mps%constraint(k)%column
        matrix(i,j) = mps%constraint(k) %value
end do
```

The IMSL derived type d_MPS is the double precision counterpart to s_MPS. The IMSL derived type d_SparseMatrixElement is the double precision counterpart to s_SparseMatrixElement.

## FORTRAN 90 Interface

Generic: CALL MPS_FREE (MPS)
Specific: The specific interface names are S_MPS_FREE and D_MPS_FREE.

## Description

This subroutine simply issues deallocate statements for each of the arrays allocated in the IMSL derived type s_MPS defined above. It is supplied as a convenience utility to the user of READ_MPS.

## Example

In the following example, the space that had been allocated to accommodate the IMSL derived type S_MPS is deallocated with a call to MPS_FREE after a call to READ_MPS was made.

```
use read_mps_int
use mps_free_int
implici\overline{t}}\mathrm{ none
TYPE(S_MPS) mps
CALL read_mps ('test.mps', mps)
call mps_free (mps)
end
```


## DENSE_LP

Solves a linear programming problem.
NOTE: DENSE_LP is available in double precision only.

## Required Arguments

A-m by NVAR matrix containing the coefficients of the M constraints. (Input)
$\boldsymbol{B L}$ - Vector of length M containing the lower limit of the general constraints; if there is no lower limit on the $I$-th constraint, then $\mathrm{BL}(\mathrm{I})$ is not referenced. (Input)
$\boldsymbol{B} \boldsymbol{U}$ - Vector of length $M$ containing the upper limit of the general constraints; if there is no upper limit on the $I$-th constraint, then $B U(I)$ is not referenced; if there are no range constraints, BL and BU can share the same storage locations. (Input)
$\boldsymbol{C}$ - Vector of length NVAR containing the coefficients of the objective function. (Input)

IRTYPE - Vector of length M indicating the types of general constraints in the matrix A. (Input)
Let $R(I)=A(I, 1) * \operatorname{XSOL}(1)+\ldots+A(I, N V A R) * X S O L(N V A R)$. Then, the value of IRTYPE(I) signifies the following:

| Irtype[I] | I-th Constraint |
| :---: | :--- |
| 0 | $\mathrm{BL}(\mathrm{I})=\mathrm{R}(\mathrm{I})=\mathrm{BU}(\mathrm{I})$ |
| 1 | $\mathrm{R}(\mathrm{I}) \leq \mathrm{BU}(\mathrm{I})$ |
| 2 | $\mathrm{R}(\mathrm{I}) \geq \mathrm{BL}(\mathrm{I})$ |
| 3 | $\mathrm{BL}(\mathrm{I}) \leq \mathrm{R}(\mathrm{I}) \leq \mathrm{BU}(\mathrm{I})$ |
| 4 | Ignore this constraint |

$\mathbf{O B J}$ - Value of the objective function. (Output)
$\boldsymbol{X S O L}$ - Vector of length NVAR containing the primal solution.(Output)
DSOL - Vector of length M containing the dual solution. (Output)

## Optional Arguments

$\boldsymbol{M}$ - Number of constraints. (Input)
Default: $\mathrm{M}=\operatorname{SIZE}(\mathrm{A}, 1)$.
NVAR - Number of variables. (Input)
Default: NVAR $=$ SIZE (A,2).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
LDA must be at least M.
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{X L B}$ - Vector of length NVAR containing the lower bound on the variables; if there is no lower bound on a variable, then 1.0D30 should be set as the lower bound. (Input) Default: $\mathrm{XLB}=0.0 \mathrm{D} 0$.
$\boldsymbol{X U B}$ - Vector of length NVAR containing the upper bound on the variables; if there is no upper bound on a variable, then -1.0 D 30 should be set as the upper bound. (Input) Default: No upperbound enforced.

ITREF - The type if iterative refinement used. (Input)

| ITREF | Refinement |
| :---: | :--- |
| 0 | No refinement |
| 1 | Iterative refinement |
| 2 | Use extended refinement. Iterate until <br> no more progress. |

Default: ITREF $=0$.

ITERS - Number of iterations. (Output)
IERR - Status flag indicating which warning conditions were set upon completion.
(Output)

| IERR | Status |
| :---: | :--- |
| $\geq 0$ | Solution found. IERR $=0$ indicates there are no <br> warning conditions. If the solution was found <br> with warning conditions IERR is incremented by <br> the number given below. <br> 1 is added to the value returned if there are <br> multiple solutions giving essentially the same <br> minimum. <br> 2 is added to the value returned if there were <br> some constraints discarded because they were too <br> linearly dependent on other active constraints. |
| 2 | 4 is added to the value returned if the constraints <br> were not satisfied. L L minimization was applied <br> to all (including bounds on simple variables) but <br> the equalities, to approximate violated non- <br> equalities as well as possible. If a feasible <br> solution is possible then refinement may help <br> 8 is added to the value returned if the algorithm <br> appears to be cycling. Using refinement may <br> help. |
| 8 | 4 |

## FORTRAN 90 Interface

Generic: CALL DENSE_LP (A, BL, BU, C, IRTYPE, OBJ, XSOL, DSOL [,..])
Specific: The specific interface name is D_DENSE_LP. This subroutine is available in double precision only.

## Description

The routine DENSE_LP solves the linear programming problem

$$
\min _{x \in \mathbf{R}^{n}} c^{T} x
$$

$$
\begin{array}{ll}
\text { subject to } & b_{l} \leq A_{x} \leq b_{u} \\
& x_{l} \leq x \leq x_{u}
\end{array}
$$

where $c$ is the objective coefficient vector, $A$ is the coefficient matrix, and the vectors $b_{l}, b_{u}, x_{l}$ and $x_{u}$ are the lower and upper bounds on the constraints and the variables, respectively.

DENSE_LP uses an active set strategy.
Refer to the following paper for further information: Krogh, Fred, T. (2005), An Algorithm for Linear Programming, http://mathalacarte.com/fkrogh/pub/lp.pdf,Tujunga, CA.

## Comments

1. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 1 | 1 | Multiple solutions giving essentially the same solution exist. |
| 3 | 1 | Some constraints were discarded because they were too linearly <br> dependent on other active constraints. |
| 3 | 2 | All constraints are not satisfied. |
| 3 | 3 | The algorithm appears to be cycling. |
| 4 | 1 | The problem appears vacuous. |
| 4 | 2 | The problem is unbounded. |
| 4 | 3 | An acceptable pivot could not be found. |
| 4 | 4 | The constraint bounds are inconsistent. |
| 4 | 5 | The variable bounds are inconsistent. |

## Example 1

The linear programming problem in the standard form

$$
\begin{array}{rlr}
\min f(x)=-x_{1}-3 x_{2} & \\
\text { subject to } \left.\begin{array}{ll}
x_{1}+x_{2}+x_{3} & \\
x_{1}+x_{2}-x_{4} & =0.5 \\
x_{1} & =1.0 \\
x_{2} & +x_{5} \\
x_{i} \geq 0, \text { for } i=1, \ldots, 6 & +x_{6}
\end{array}\right) 1.0
\end{array}
$$

is solved.

```
USE UMACH INT
USE WRRRN_INT
USE DENSE LP INT
IMPLICIT N
INTEGER NOUT, M, NVAR
PARAMETER (M=4, NVAR=6)
DOUBLE PRECISION A(M, NVAR), B(M), C(NVAR), XSOL(NVAR), &
    DSOL(M), BL(M), BU(M), OBJ
INTEGER IRTYPE(M)
DATA A/1, 1, 1, 0, 1, 1, 0, 1, 1, 0, 0, 0, 0, -1, &
    0, 0, 0, 0, 1, 0, 0, 0, 0, 1/
```

```
    DATA B/1.5, 0.5, 1.0, 1.0/
    DATA C/-1.0, -3.0, 0.0, 0.0, 0.0, 0.0/
    DATA BL/1.5, 0.5, 1.0, 1.0/
    DATA BU/M*-1.D30/
    DATA IRTYPE/M*0/
    CALL UMACH(2, NOUT)
! Solve the LP problem
    CALL DENSE_LP (A, BL, BU, C, IRTYPE, OBJ, XSOL, DSOL)
    WRITE(NOUT, 99999) OBJ
    CALL WRRRN('Solution', XSOL, 1, NVAR, 1)
99999 FORMAT (' Objective', F9.4)
END
```


## Output

```
Objective -3.5000
```

|  | Solution |  |  |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 1 | 2 | 3 | 4 | 5 | 6 |  |

## Additional Examples

## Example 2

This example demonstrates how READ_MPS can be used together with DENSE_LP to solve a linear programming problem defined in an MPS file. The MPS file used in this example is an uncompressed version of the file 'afiro', available from http://www.netlib.org/lp/data/.

```
USE UMACH_INT
USE WRRRN INT
USE READ MPS INT
USE DENSE_LP_INT
IMPLICIT NONE
REAL(KIND(1D0)) OBJ
REAL(KIND(1D0)), ALLOCATABLE : : XSOL(:)
REAL(KIND(1D0)), ALLOCATABLE :: DSOL(:)
REAL (KIND(1D0)), ALLOCATABLE :: A(:,:)
INTEGER, ALLOCATABLE : : IRTYPE(:)
TYPE(D_MPS) PROBLEM
CHARACTER NAME*256
INTEGER I,J, K, NOUT
CALL UMACH (2, NOUT)
! READ LP PROBLEM FROM THE MPS FILE.
NAME = 'afiro'
CALL READ MPS (NAME, PROBLEM)
ALLOCATE (`
ALLOCATE (IRTYPE (PROBLEM%NROWS))
ALLOCATE (XSOL (PROBLEM%NCOLUMNS))
ALLOCATE (DSOL (PROBLEM%NROWS))
```

```
    A = 0
    IRTYPE = 3
! FILL DENSE A
    DO K = 1, PROBLEM%NONZEROS
        I = PROBLEM%CONSTRAINT (K) %ROW
        J = PROBLEM%CONSTRAINT (K) %COLUMN
        A(I,J) = PROBLEM%CONSTRAINT (K) %VALUE
    ENDDO
! CALL THE LP SOLVER
    CALL DENSE_LP (A, PROBLEM%LOWER_RANGE, PROBLEM%UPPER_RANGE, &
        PROBLEM%OBJECTIVE, IRTYPE, OBJ, XSOL, D
        XLB=PROBLEM%LOWER_BOUND, XUB=PROBLEM%UPPER_BOUND)
    WRITE(NOUT, 99999) OBJ
    CALL WRRRN('Solution', XSOL, 1, PROBLEM%NROWS, 1)
    DEALLOCATE (A)
    DEALLOCATE (IRTYPE)
    DEALLOCATE (XSOL)
    DEALLOCATE (DSOL)
9 9 9 9 9 ~ F O R M A T ( ' O b j e c t i v e : ~ ' , ~ E 1 6 . 7 )
    END
```


## Output

Objective: $\quad-0.4647531 \mathrm{E}+03$


## DLPRS

Solves a linear programming problem via the revised simplex algorithm.

## Required Arguments

A—M by NVAR matrix containing the coefficients of the M constraints. (Input)
$\boldsymbol{B L}$ - Vector of length $M$ containing the lower limit of the general constraints; if there is no lower limit on the $I$-th constraint, then $\mathrm{BL}(\mathrm{I})$ is not referenced. (Input)
$\boldsymbol{B} \boldsymbol{U}$ - Vector of length m containing the upper limit of the general constraints; if there is no upper limit on the I-th constraint, then $B U(I)$ is not referenced; if there are no range constraints, BL and BU can share the same storage locations. (Input)
$\boldsymbol{C}$ - Vector of length NVAR containing the coefficients of the objective function. (Input)

IRTYPE - Vector of length m indicating the types of general constraints in the matrix $A$. (Input)
Let $R(I)=A(I, 1) * X S O L(1)+\ldots+A(I, N V A R) * X S O L(N V A R)$. Then, the value of IRTYPE(I) signifies the following:

IRTYPE(I) I-th Constraint

0
BL(I).EQ.R(I).EQ.BU(I)
1
R(I).LE.BU(I)
2
R(I). GE. $\operatorname{BL}(I)$

3
BL(I).LE.R(I).LE.BU(I)
$\boldsymbol{O B J}$ - Value of the objective function. (Output)
$\boldsymbol{X S O L}$ - Vector of length NVAR containing the primal solution. (Output)
DSOL - Vector of length M containing the dual solution. (Output)

## Optional Arguments

M - Number of constraints. (Input)
Default: $\mathrm{M}=\operatorname{SIZE}(\mathrm{A}, 1)$.
NVAR - Number of variables. (Input)
Default: NVAR $=\operatorname{SIZE}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
LDA must be at least M.
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{X L B}$ - Vector of length NVAR containing the lower bound on the variables; if there is no lower bound on a variable, then 1.0E30 should be set as the lower bound. (Input) Default: $\mathrm{xLB}=0.0$.
$\boldsymbol{X U B}$ - Vector of length NVAR containing the upper bound on the variables; if there is no upper bound on a variable, then -1.0 E 30 should be set as the upper bound. (Input) Default: $\mathrm{XUB}=3.4 \mathrm{e} 38$ for single precision and $1.79 \mathrm{~d}+308$ for double precision.

## FORTRAN 90 Interface

Generic: CALL DLPRS (A, BL, BU, C, IRTYPE, OBJ, XSOL, DSOL [,...])
Specific: The specific interface names are S_DLPRS and D_DLPRS.

## FORTRAN 77 Interface

Single: CALL DLPRS (M, NVAR, A, LDA, BL, BU, C, IRTYPE, XLB, XUB, OBJ, XSOL, DSOL)

Double: The double precision name is DDLPRS.

## Description

The routine DLPRS uses a revised simplex method to solve linear programming problems, i.e., problems of the form

$$
\min _{x \in \mathbf{R}^{n}} c^{T} x
$$

$$
\text { subject to } b_{l} \leq A x \leq b_{u}
$$

$$
x_{l} \leq x \leq x_{u}
$$

where $c$ is the objective coefficient vector, $A$ is the coefficient matrix, and the vectors $b_{l}, b_{u}, x_{l}$ and $x_{u}$ are the lower and upper bounds on the constraints and the variables, respectively.
For a complete description of the revised simplex method, see Murtagh (1981) or Murty (1983).

## Comments

1. Workspace may be explicitly provided, if desired, by use of D2PRS/DD2PRS. The reference is:
```
CALL D2PRS (M, NVAR, A, LDA, BL, BU, C, IRTYPE, XLB, XUB, OBJ, XSOL, DSOL, AWK,
LDAWK, WK, IWK)
```

The additional arguments are as follows:
AWK - Real work array of dimension 1 by 1. (AWK is not used in the new implementation of the revised simplex algorithm. It is retained merely for calling sequence consistency.)

LDAWK - Leading dimension of AWK exactly as specified in the dimension statement of the calling program. LDAWK should be 1. (LDAWK is not used in the new implementation of the revised simplex algorithm. It is retained merely for calling sequence consistency.)
$\boldsymbol{W} \boldsymbol{K}$ - Real work vector of length $\mathrm{M} *(\mathrm{M}+28)$.
IWK - Integer work vector of length $29 * \mathrm{~m}+3 *$ NVAR.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 3 | 1 | The problem is unbounded. |
| 4 | 2 | Maximum number of iterations exceeded. |

44 Moved to a vertex that is poorly conditioned; using double precision may help.
$4 \quad 5 \quad$ The bounds are inconsistent.

## Example

A linear programming problem is solved.

```
USE DLPRS_INT
USE UMACH_
USE SSCAL__INT
IMPLICIT NONE
INTEGER LDA, M, NVAR
PARAMETER (M=2, NVAR=2, LDA=M)
                                    M = number of constraints
                                    NVAR = number of variables
INTEGER I, IRTYPE(M), NOUT
REAL A(LDA,NVAR), B (M), C (NVAR), DSOL (M), OBJ, XLB (NVAR), &
        XSOL (NVAR), XUB (NVAR)
                                    Set values for the following problem
                                    Max 1.0*XSOL(1) + 3.0*XSOL (2)
                                    XSOL(1) + XSOL(2) .LE. 1.5
                                    XSOL(1) + XSOL(2) .GE. 0.5
                            0 .LE. XSOL(1) .LE. 1
                            O .LE. XSOL(2) .LE. 1
DATA XLB/2*0.0/, XUB/2*1.0/
DATA A/4*1.0/, B/1.5, .5/, C/1.0, 3.0/
DATA IRTYPE/1, 2/
    To maximize, C must be multiplied by
    -1.
CALL SSCAL (NVAR, -1.0E0, C, 1)
    Solve the LP problem. Since there is
    no range constraint, only B is
    needed.
CALL DLPRS (A, B, B, C, IRTYPE, OBJ, XSOL, DSOL, &
            XUB=XUB)
        OBJ must be multiplied by -1 to get
        the true maximum.
OBJ = -OBJ
    DSOL must be multiplied by -1 for
    maximization.
CALL SSCAL (M, -1.0E0, DSOL, 1)
        Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) OBJ, (XSOL(I),I=1,NVAR), (DSOL (I),I=1,M)
99999 FORMAT (//, ' Objective = ', F9.4, //, ' Primal ',&
```

!

```
'Solution =', 2F9.4, //, ' Dual solution =', 2F9.4)
```

$!$

END

## Output

| Objective | $=3.5000$ |  |
| :--- | :--- | :--- | :--- |
| Primal Solution | $=0.5000$ | 1.0000 |
| Dual solution $=1.0000$ | 0.0000 |  |

## SLPRS

Solves a sparse linear programming problem via the revised simplex algorithm.

## Required Arguments

$\boldsymbol{A}$ - Vector of length NZ containing the coefficients of the M constraints. (Input)
IROW - Vector of length NZ containing the row numbers of the corresponding element in A. (Input)

JCOL - Vector of length NZ containing the column numbers of the corresponding elements in $A$. (Input)
$\boldsymbol{B L}$ — Vector of length M containing the lower limit of the general constraints; if there is no lower limit on the I-th constraint, then BL (I) is not referenced. (Input)
$\boldsymbol{B U}$ - Vector of length M containing the upper lower limit of the general constraints; if there is no upper limit on the I-th constraint, then $B U(I)$ is not referenced. (Input)
$\boldsymbol{C}$ - Vector of length NVAR containing the coefficients of the objective function. (Input)
IRTYPE - Vector of length M indicating the types of general constraints in the matrix A. (Input)
Let $R(I)=A(I, 1) \star X S O L(1)+\ldots+A(I, N V A R) \star X S O L(N V A R)$

IRTYPE(I)
0
1
2
3

## I-th CONSTRAINT

$B L(I)=R(I)=B U(I)$
$R(I) \leq B U(I)$
$R(I) \geq B L(I)$
$\mathrm{BL}(\mathrm{I}) \leq \mathrm{R}(\mathrm{I}) \leq \mathrm{BU}(\mathrm{I})$
$\boldsymbol{O B J}$ - Value of the objective function. (Output)
$\boldsymbol{X S O L}$ - Vector of length NVAR containing the primal solution. (Output)

DSOL — Vector of length M containing the dual solution. (Output)

## Optional Arguments

$\boldsymbol{M}$ - Number of constraints. (Input)
Default: M = SIZE (IRTYPE, 1 ).
NVAR - Number of variables. (Input)
Default: NVAR $=\operatorname{SIZE}(\mathrm{C}, 1)$.
$N Z$ - Number of nonzero coefficients in the matrix $A$. (Input)
Default: NZ $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{X L B}$ - Vector of length NVAR containing the lower bound on the variables; if there is no lower bound on a variable, then 1.0 E 30 should be set as the lower bound. (Input) Default: $\mathrm{xLB}=0.0$.
$\boldsymbol{X U B}$ - Vector of length NVAR containing the upper bound on the variables; if there is no upper bound on a variable, then -1.0 E 30 should be set as the upper bound. (Input) Default: $\mathrm{XLB}=3.4 \mathrm{e} 38$ for single precision and $1.79 \mathrm{~d}+308$ for double precision.

## FORTRAN 90 Interface

```
Generic: CALL SLPRS (A, IROW, JCOL, BL, BU, C, IRTYPE, OBJ, XSOL,
    DSOL [,...])
```

Specific: The specific interface names are S_SLPRS and D_SLPRS.

## FORTRAN 77 Interface

Single: CALL SLPRS (M, NVAR, NZ, A, IROW, JCOL, BL, BU, C, IRTYPE, XLB, XUB, OBJ, XSOL, DSOL)

Double: The double precision name is DSLPRS.

## Description

This subroutine solves problems of the form

$$
\min c \boldsymbol{T} x
$$

subject to

$$
\begin{aligned}
b_{l} & \leq A x \leq b_{u}, \\
x_{l} & \leq x \leq x_{u}
\end{aligned}
$$

where $c$ is the objective coefficient vector, $A$ is the coefficient matrix, and the vectors $b_{l}, b_{u}, x_{l}$, and $x_{u}$ are the lower and upper bounds on the constraints and the variables, respectively. SLPRS is designed to take advantage of sparsity in $A$. The routine is based on DPLO by Hanson and Hiebert.

## Comments

Workspace may be explicitly provided, if desired, by use of S2PRS/DS2PRS. The reference is:

CALL S2PRS (M, NVAR, NZ, A, IROW, JCOL, BL, BU, C, IRTYPE, XLB, XUB, OBJ, XSOL, DSOL, IPARAM, RPARAM, COLSCL, ROWSCL, WORK, LW, IWORK, LIW)

The additional arguments are as follows:
IPARAM - Integer parameter vector of length 12. If the default parameters are desired for SLPRS, then set IPARAM(1) to zero and call the routine SLPRS. Otherwise, if any nondefault parameters are desired for IPARAM or RPARAM, then the following steps should be taken before calling SLPRS:
CALL S5PRS (IPARAM, RPARAM)
Set nondefault values for IPARAM and RPARAM.

Note that the call to S5PRS will set IPARAM and RPARAM to their default values so only nondefault values need to be set above.
$\operatorname{IPARAM}(1)=0$ indicates that a minimization problem is solved. If set to $1, a$ maximization problem is solved.

Default: 0
$\operatorname{IPARAM}(2)=$ switch indicating the maximum number of iterations to be taken before returning to the user. If set to zero, the maximum number of iterations taken is set to $3 *($ NVARS +M ). If positive, that value is used as the iteration limit.

Default: $\operatorname{IPARAM}(2)=0$
$\operatorname{IPARAM}(3)=$ indicator for choosing how columns are selected to enter the basis. If set to zero, the routine uses the steepest edge pricing strategy which is the best local move. If set to one, the minimum reduced cost pricing strategy is used. The steepest edge pricing strategy generally uses fewer iterations than the minimum reduced cost pricing, but each iteration costs more in terms of the amount of calculation performed. However, this is very problem-dependent.

Default: $\operatorname{IPARAM}(3)=0$
$\operatorname{IPARAM}(4)=\operatorname{MXITBR}$, the number of iterations between recalculating the error in the primal solution is used to monitor the error in solving the linear system. This is an expensive calculation and every tenth iteration is generally enough.

Default: $\operatorname{IPARAM}(4)=10$
$\operatorname{IPARAM}(5)=$ NPP, the number of negative reduced costs (at most) to be found at each iteration of choosing a variable to enter the basis. If set to zero, NPP = NVARS will be used, implying that all of the reduced costs are computed at each such step. This "Partial pricing" may increase the total number of iterations required. However, it decreases the number of calculation required at each iteration. The effect on overall efficiency is very problem-dependent. If set to some positive number, that value is used as NPP.

Default: $\operatorname{IPARAM}(5)=0$
$\operatorname{IPARAM}(6)=\operatorname{IREDFQ}$, the number of steps between basis matrix redecompositions. Redecompositions also occur whenever the linear systems for the primal and dual systems have lost half their working precision.

Default: $\operatorname{IPARAM}(6)=50$
$\operatorname{IPARAM}(7)=$ LAMAT, the length of the portion of wORK that is allocated to sparse matrix storage and decomposition. LAMAT must be greater than NZ + NVARS +4 .

Default: LAMAT $=$ NZ + NVARS +5
$\operatorname{IPARAM}(8)=\operatorname{LBM}$, then length of the portion of IWORK that is allocated to sparse matrix storage and decomposition. LBM must be positive.

Default: $\mathrm{LBM}=8 * \mathrm{M}$
IPARAM(9) = switch indicating that partial results should be saved after the maximum number of iterations, IPARAM(2), or at the optimum. If IPARAM(9) is not zero, data essential to continuing the calculation is saved to a file, attached to unit number IPARAM(9). The data saved includes all the information about the sparse matrix A and information about the current basis. If $\operatorname{IPARAM}(9)$ is set to zero, partial results are not saved. It is the responsibility of the calling program to open the output file.
$\operatorname{IPARAM}(10)=$ switch indicating that partial results have been computed and stored on unit number IPARAM(10), if greater than zero. If IPARAM(10) is zero, a new problem is started.

Default: $\operatorname{IPARAM}(10)=0$
$\operatorname{IPARAM}(11)=$ switch indicating that the user supplies scale factors for the columns of the matrix $A$. If IPARAM $(11)=0, \operatorname{SLPRS}$ computes the scale factors as the reciprocals of the max norm of each column. If $\operatorname{IPARAM}(11)$ is set to one, element I of the vector COLSCL is used as the scale factor for column I of the matrix $A$. The scaling is implicit, so no input data is actually changed.

Default: $\operatorname{IPARAM}(11)=0$
$\operatorname{IPARAM}(12)=$ switch indicating that the user supplied scale factors for the rows of the matrix $A$. If IPARAM(12) is set to zero, no row scaling is one. If $\operatorname{IPARAM}(12)$ is set to 1 , element I of the vector ROWSCL is used as the scale factor for row I of the matrix $A$. The scaling is implicit, so no input data is actually changed.

Default: $\operatorname{IPARAM}(12)=0$
RPARAM - Real parameter vector of length 7.
$\operatorname{RPARAM}(1)=\operatorname{COSTSC}$, a scale factor for the vector of costs. Normally SLPRS computes this scale factor to be the reciprocal of the max norm if the vector costs after the column scaling has been applied. If RPARAM(1) is zero, SLPRS compute COSTSC.

Default: $\operatorname{RPARAM}(1)=0.0$
$\operatorname{RPARAM}(2)=\operatorname{ASMALL}$, the smallest magnitude of nonzero entries in the matrix $A$. If $\operatorname{RPARAM}(2)$ is nonzero, checking is done to ensure that all elements of $A$ are at least as large as RPARAM(2). Otherwise, no checking is done.

Default: $\operatorname{RPARAM}(2)=0.0$
$\operatorname{RPARAM}(3)=\operatorname{ABIG}$, the largest magnitude of nonzero entries in the matrix $A$. If $\operatorname{RPARAM}(3)$ is nonzero, checking is done to ensure that all elements of $A$ are no larger than RPARAM(3). Otherwise, no checking is done.
Default: $\operatorname{RPARAM}(3)=0.0$
$\operatorname{RPARAM}(4)=$ TOLLS , the relative tolerance used in checking if the residuals are feasible. RPARAM(4) is nonzero, that value is used as TOLLS, otherwise the default value is used.

Default: TOLLS $=1000.0^{*} \operatorname{amach}(4)$
$\operatorname{RPARAM}(5)=\mathrm{PHI}$, the scaling factor used to scale the reduced cost error estimates. In some environments, it may be necessary to reset PHI to the range $[0.01,0.1]$, particularly on machines with short word length and working precision when solving a large problem. If RPARAM(5) is nonzero, that value is used as PHI, otherwise the default value is used.

Default: PHI $=1.0$
$\operatorname{RPARAM}(6)=$ TOLABS, an absolute error test on feasibility. Normally a relative test is used with TOLLS (see RPARAM(4)). If this test fails, an absolute test will be applied using the value tOLABS.

Default: TOLABS $=0.0$
$\operatorname{RPARAM}(7)=$ pivot tolerance of the underlying sparse factorization routine. If RPARAM(7) is set to zero, the default pivot tolerance is used, otherwise, the RPARAM(7) is used.

Default: $\operatorname{RPARAM}(7)=0.1$

COLSCL - Array of length NVARS containing column scale factors for the matrix $A$. (Input).
COLSCL is not used if IPARAM(11) is set to zero.

ROWSCL - Array of length m containing row scale factors for the matrix $A$. (Input) ROWSCL is not used if IPARAM(12) is set to zero.

WORK - Work array of length LW.
$\boldsymbol{L W}$ - Length of real work array. LW must be at least

```
2 + 2NZ + 9NVAR + 27M + MAX(NZ + NVAR + 8, 4NVAR + 7).
```

IWORK - Integer work array of length LIW.

LIW - Length of integer work array. LIW must be at least $1+3 N V A R+41 M+\operatorname{MAX}(N Z+N V A R+8,4 N V A R+7)$.

## Example

Solve a linear programming problem, with

$$
A=\left[\begin{array}{ccccc}
0 & 0.5 & & & \\
& 1 & 0.5 & & \\
& & 1 & \ddots & \\
& & & \ddots & 0.5 \\
& & & & 1
\end{array}\right]
$$

defined in sparse coordinate format.

```
USE SLPRS_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER M, NVAR
PARAMETER (M=200, NVAR=200)
variables
INTEGER INDEX, IROW(3*M), J, JCOL(3*M), NOUT, NZ
REAL A(3*M), DSOL (M), OBJ, XSOL (NVAR)
INTEGER IRTYPE (M)
REAL B(M), C(NVAR), XL (NVAR), XU (NVAR)
                                    Specifications for subroutines
DATA B/199*1.7, 1.0/
DATA C/-1.0, -2.0, -3.0, -4.0, -5.0, -6.0, -7.0, -8.0, -9.0, &
-10.0, 190*-1.0/
DATA XL/200*0.1/
DATA XU/200*2.0/
DATA IRTYPE/200*1/
CALL UMACH (2, NOUT)
Define A
```

!
!
!

```
    INDEX = 1
    DO 10 J=2, M
! Superdiagonal element
        IROW(INDEX) = J - 1
        JCOL(INDEX) = J
        A(INDEX) = 0.5
        IROW(INDEX+1) = J
        JCOL (INDEX+1) = J
        A(INDEX+1) = 1.0
        INDEX = INDEX + 2
    10 CONTINUE
    NZ = INDEX - 1
!
!
    XL(4) = 0.2
    CALL SLPRS (A, IROW, JCOL, B, B, C, IRTYPE, OBJ, XSOL, DSOL, &
        NZ=NZ, XLB=XL, XUB=XU)
!
    WRITE (NOUT,99999) OBJ
!
99999 FORMAT (/, 'The value of the objective function is ', E12.6)
!
    END
```


## Output

The value of the objective function is $-.280971 \mathrm{E}+03$

## QPROG

Solves a quadratic programming problem subject to linear equality/inequality constraints.

## Required Arguments

NEQ - The number of linear equality constraints. (Input)
A - NCON by NVAR matrix. (Input)
The matrix contains the equality contraints in the first NEQ rows followed by the inequality constraints.

B - Vector of length NCON containing right-hand sides of the linear constraints. (Input)
$\boldsymbol{G}$ - Vector of length NVAR containing the coefficients of the linear term of the objective function. (Input)
$\boldsymbol{H}$ - NVAR by NVAR matrix containing the Hessian matrix of the objective function. (Input) H should be symmetric positive definite; if H is not positive definite, the algorithm attempts to solve the QP problem with H replaced by a H + DIAGNL * I such that H + DIAGNL * I is positive definite. See Comment 3.
$\boldsymbol{S O L}-$ Vector of length NVAR containing solution. (Output)

## Optional Arguments

NVAR - The number of variables. (Input)
Default: NVAR $=\operatorname{SIZE}(A, 2)$.
NCON - The number of linear constraints. (Input)
Default: NCON = SIZE (A,1).
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
$\mathbf{L D H}$ - Leading dimension of H exactly as specified in the dimension statement of the calling program. (Input)
Default: LDH = SIZE (H, 1).

DIAGNL - Scalar equal to the multiple of the identity matrix added to $H$ to give a positive definite matrix. (Output)

NACT - Final number of active constraints. (Output)
$\boldsymbol{I A C T}$ - Vector of length NVAR containing the indices of the final active constraints in the first NACT positions. (Output)

ALAMDA - Vector of length NVAR containing the Lagrange multiplier estimates of the final active constraints in the first NACT positions. (Output)

## FORTRAN 90 Interface

Generic: CALL QPROG (NEQ, A, B, G, H, SOL [,...])
Specific: The specific interface names are $S_{-Q P R O G}$ and $D \_Q P R O G$.

## FORTRAN 77 Interface

Single: CALL QPROG (NVAR, NCON, NEQ, A, LDA, B, G, H, LDH, DIAGNL, SOL, NACT, IACT, ALAMDA)

Double: $\quad$ The double precision name is DQPROG.

## Description

The routine QPROG is based on M.J.D. Powell's implementation of the Goldfarb and Idnani (1983) dual quadratic programming ( QP ) algorithm for convex QP problems subject to general linear equality/inequality constraints, i.e., problems of the form

$$
\begin{aligned}
& \min _{x \in \mathbf{R}^{g}} g^{T} x+\frac{1}{2} x^{T} H x \\
& \text { subject to } \quad A_{1} x=b_{1} \\
& \quad A_{2} x \geq b_{2}
\end{aligned}
$$

given the vectors $b_{1}, b_{2}$, and $g$ and the matrices $H, A_{1}$, and $A_{2}$. $H$ is required to be positive definite. In this case, a unique $x$ solves the problem or the constraints are inconsistent. If $H$ is not positive definite, a positive definite perturbation of $H$ is used in place of $H$. For more details, see Powell (1983, 1985).

## Comments

1. Workspace may be explicitly provided, if desired, by use of Q2ROG/DQ2ROG. The reference is:

CALL Q2ROG (NVAR, NCON, NEQ, A, LDA, B, G, H, LDH, DIAGNL, SOL, NACT, IACT, ALAMDA, WK)

The additional argument is:
WK - Work vector of length $(3$ * NVAR**2 +11 * NVAR $) / 2+$ nCon.
2. Informational errors

| Type | Code |
| :---: | :---: | :--- |
| 3 |  | | 1 | Due to the effect of computer rounding error, a change in the <br> variables fail to improve the objective function value; usually the |
| :--- | :--- |
| solution is close to optimum. |  |

 H + DIAGNL * I should also be used in the definition of the Lagrange multipliers.

## Example

The quadratic programming problem

$$
\begin{array}{ll}
\min f(x)= & x_{1}^{2}+x_{2}^{2}+x_{3}^{2}+x_{4}^{2}+x_{5}^{2}-2 x_{2} x_{3}-2 x_{4} x_{5}-2 x_{1} \\
\text { subject to } \quad & x_{1}+x_{2}+x_{3}+x_{4}+x_{5}=5 \\
& x_{3}-2 x_{4}-2 x_{5}=-3
\end{array}
$$

is solved.

```
USE QPROG_INT
USE UMACH_
IMPLICIT NONE
INTEGER LDA, LDH, NCON, NEQ, NVAR
PARAMETER (NCON=2, NEQ=2, NVAR=5, LDA=NCON, LDH=NVAR)
```

```
!
    INTEGER K, NACT, NOUT
    REAL A(LDA,NVAR), ALAMDA (NVAR), B (NCON), G(NVAR), &
        H(LDH,LDH), SOL (NVAR)
                        Set values of A, B, G and H.
                                A =(1.0 1.0 1.0 1.0 1.0)
                                ( 0.0 0.0 1.0 -2.0 -2.0)
                B = ( 5.0 -3.0)
                G =(\begin{array}{lllll}{-2.0}&{0.0}&{0.0}&{0.0}&{0.0}\end{array})
                H}=(\begin{array}{lllll}{2.0}&{0.0}&{0.0}&{0.0}&{0.0}\end{array}
            ( 0.0 2.0 -2.0 0.0 0.0)
            ( 0.0 -2.0 2.0 0.0 0.0)
            ( 0.0 0.0 0.0 2.0 -2.0)
            (0.0 0.0 0.0 -2.0 2.0)
    DATA A/1.0, 0.0, 1.0, 0.0, 1.0, 1.0, 1.0, -2.0, 1.0, -2.0/
    DATA B/5.0, -3.0/
DATA G/-2.0, 4*0.0/
DATA H/2.0, 5*0.0, 2.0, -2.0, 3*0.0, -2.0, 2.0, 5*0.0, 2.0, &
    -2.0, 3*0.0, -2.0, 2.0/
CALL QPROG (NEQ, A, B, G, H, SOL)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) (SOL(K),K=1,NVAR)
99999 FORMAT (' The solution vector is', /, ' SOL = (', 5F6.1, &
    ' )')
END
```


## Output

The solution vector is
$\mathrm{SOL}=\left(\begin{array}{lllll}1.0 & 1.0 & 1.0 & 1.0 & 1.0\end{array}\right)$

## LCONF

Minimizes a general objective function subject to linear equality/inequality constraints.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Value of NVAR. (Input)
$\mathrm{X}-$ Vector of length N at which point the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
NEQ - The number of linear equality constraints. (Input)
A - NCON by NVAR matrix. (Input)
The matrix contains the equality constraint gradients in the first NEQ rows, followed by the inequality constraint gradients.
$\boldsymbol{B}$ - Vector of length NCON containing right-hand sides of the linear constraints. (Input) Specifically, the constraints on the variables $X(I), I=1, \ldots$, NVAR are $A(K, 1) * X(1)+\ldots+A(K, N V A R) * X(N V A R) . E Q \cdot B(K), K=1, \ldots$, NEQ.A $(K, 1) * X(1)+\ldots+A(K, N V A R) * X(N V A R) . L E . B(K), K=N E Q+1, \ldots$, NCON. Note that the data that define the equality constraints come before the data of the inequalities.
$\boldsymbol{X L B}$ - Vector of length NVAR containing the lower bounds on the variables; choose a very large negative value if a component should be unbounded below or set $\mathrm{XLB}(\mathrm{I})=\mathrm{XUB}(\mathrm{I})$ to freeze the I -th variable. (Input) Specifically, these simple bounds are XLB(I). LE.X(I), I = 1, $\ldots$, NVAR.
$\boldsymbol{X U B}$ - Vector of length NVAR containing the upper bounds on the variables; choose a very large positive value if a component should be unbounded above. (Input) Specifically, these simple bounds are X(I).LE. XUB(I), I =1, $\ldots$, NVAR.
$\boldsymbol{S O L}-$ Vector of length NVAR containing solution. (Output)

## Optional Arguments

NVAR - The number of variables. (Input)
Default: NVAR $=\operatorname{SIZE}(A, 2)$.
$N C O N$ - The number of linear constraints (excluding simple bounds). (Input)
Default: NCON $=\operatorname{SIZE}(A, 1)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

XGUESS - Vector of length NVAR containing the initial guess of the minimum. (Input) Default: XGUESS $=0.0$.

ACC - The nonnegative tolerance on the first order conditions at the calculated solution. (Input)
Default: $\mathrm{ACC}=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.

MAXFCN - On input, maximum number of function evaluations allowed. (Input/ Output) On output, actual number of function evaluations needed.
Default: $\operatorname{MAXFCN}=400$.
$\mathbf{O B J}$ - Value of the objective function. (Output)
$\boldsymbol{N A C T}$ - Final number of active constraints. (Output)
IACT - Vector containing the indices of the final active constraints in the first NACT positions. (Output)
Its length must be at least $\mathrm{NCON}+2$ * NVAR.
ALAMDA - Vector of length NVAR containing the Lagrange multiplier estimates of the final active constraints in the first NACT positions. (Output)

## FORTRAN 90 Interface

Generic: CALL LCONF (FCN, NEQ, A, B, XLB, XUB, SOL [,...])
Specific: The specific interface names are S_LCONF and D_LCONF.

## FORTRAN 77 Interface

Single: CALL LCONF (FCN, NVAR, NCON, NEQ, A, LDA, B, XLB, XUB, XGUESS, ACC, MAXFCN, SOL, OBJ, NACT, IACT, ALAMDA)

Double: $\quad$ The double precision name is DLCONF.

## Description

The routine LCONF is based on M.J.D. Powell's TOLMIN, which solves linearly constrained optimization problems, i.e., problems of the form

$$
\begin{array}{ll}
\min _{x \in \mathbf{R}^{n}} f(x) \\
\text { subject to } & A_{1} x=b_{1} \\
& A_{2} x \leq b_{2} \\
& x \mathrm{l} \leq x \leq x \mathrm{u}
\end{array}
$$

given the vectors $b_{1}, b_{2}, x_{l}$ and $x_{u}$ and the matrices $A_{1}$, and $A_{2}$.
The algorithm starts by checking the equality constraints for inconsistency and redundancy. If the equality constraints are consistent, the method will revise $x^{0}$, the initial guess provided by the user, to satisfy

$$
A_{1} x=b_{1}
$$

Next, $x^{0}$ is adjusted to satisfy the simple bounds and inequality constraints. This is done by solving a sequence of quadratic programming subproblems to minimize the sum of the constraint or bound violations.

Now, for each iteration with a feasible $x^{k}$, let $J_{k}$ be the set of indices of inequality constraints that have small residuals. Here, the simple bounds are treated as inequality constraints. Let $I_{k}$ be the set of indices of active constraints. The following quadratic programming problem

$$
\begin{gathered}
\min f\left(x^{k}\right)+d^{T} \nabla f\left(x^{k}\right)+\frac{1}{2} d^{T} B^{k} d \\
\text { subject to } \quad a_{j} d=0 j \in I \mathrm{k} \\
a_{j} d \leq 0 j \in J \mathrm{k}
\end{gathered}
$$

is solved to get $\left(d^{k}, \lambda^{k}\right)$ where $a_{j}$ is a row vector representing either a constraint in $A_{1}$ or $A_{2}$ or a bound constraint on $x$. In the latter case, the $a_{j}=e_{i}$ for the bound constraint $x_{i} \leq\left(x_{u}\right)_{i}$ and $a_{j}=-e_{i}$ for the constraint $-x_{i} \leq\left(-x_{l}\right)_{i}$. Here, $e_{i}$ is a vector with a 1 as the $i$-th component, and zeroes elsewhere. $\lambda^{k}$ are the Lagrange multipliers, and $B^{k}$ is a positive definite approximation to the second derivative $\nabla^{2} f\left(x^{k}\right)$.

After the search direction $d^{k}$ is obtained, a line search is performed to locate a better point. The new point $x^{k+1}=x^{k}+\alpha^{k} d^{k}$ has to satisfy the conditions

$$
f\left(x^{k}+\alpha^{k} d^{k}\right) \leq f\left(x^{k}\right)+0.1 \alpha^{k}\left(d^{k}\right)^{T} \nabla f\left(x^{k}\right)
$$

and

$$
\left(d^{k}\right)^{T} \nabla f\left(x^{k}+\alpha^{k} d^{k}\right) \geq 0.7\left(d^{k}\right)^{T} \nabla f\left(x^{k}\right)
$$

The main idea in forming the set $J_{k}$ is that, if any of the inequality constraints restricts the steplength $\alpha^{k}$, then its index is not in $J_{k}$. Therefore, small steps are likely to be avoided.
Finally, the second derivative approximation, $B^{k}$, is updated by the BFGS formula, if the condition

$$
\left(d^{k}\right)^{T} \nabla f\left(x^{k}+\alpha^{k} d^{k}\right)-\nabla f\left(x^{k}\right)>0
$$

holds. Let $x^{k} \leftarrow x^{k+l}$, and start another iteration.
The iteration repeats until the stopping criterion

$$
\left\|\nabla f\left(x^{k}\right)-A^{k} \lambda^{k}\right\|_{2} \leq \tau
$$

is satisfied; here, $\tau$ is a user-supplied tolerance. For more details, see Powell $(1988,1989)$.
Since a finite-difference method is used to estimate the gradient for some single precision calculations, an inaccurate estimate of the gradient may cause the algorithm to terminate at a noncritical point. In such cases, high precision arithmetic is recommended. Also, whenever the exact gradient can be easily provided, routine LCONG should be used instead.

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2ONF/DL2ONF. The reference is:

CALL L2ONF (FCN, NVAR, NCON, NEQ, A, LDA, B, XLB, XUB, XGUESS, ACC, MAXFCN, SOL, OBJ, NACT, IACT, ALAMDA, IPRINT, INFO, WK)

The additional arguments are as follows:
IPRINT — Print option (see Comment 3). (Input)
INFO - Informational flag (see Comment 3). (Output)
$\boldsymbol{W K}-$ Real work vector of length NVAR**2 +11 * NVAR + NCON.
2. Informational Errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 4 | The equality constraints are inconsistent. <br> 4 |
| 4 | 5 | The equality constraints and the bounds on the variables are found to <br> be inconsistent. |
| 4 | No vector X satisfies all of the constraints. In particular, the current <br> active constraints prevent any change in X that reduces the sum of <br> constraint violations. |  |
| 4 | 7 | Maximum number of function evaluations exceeded. <br> 4 |

3. The following are descriptions of the arguments IPRINT and INFO:

IPRINT - This argument must be set by the user to specify the frequency of printing during the execution of the routine LCONF. There is no printed output if $\operatorname{IPRINT}=0$. Otherwise, after ensuring feasibility, information is given every IABS(IPRINT) iterations and whenever a parameter called TOL is reduced. The printing provides the values of $X(),. F($.$) and G()=.\operatorname{GRAD}(F)$ if IPRINT is positive. If IPRINT is negative, this information is augmented by the current values of $\operatorname{IACT}(\mathrm{K}) \mathrm{K}=1, \ldots$, $\operatorname{NACT}, \operatorname{PAR}(K) \quad K=1, \ldots, \operatorname{NACT}$ and $\operatorname{RESKT}(I) I=1, \ldots, N$. The reason for returning to the calling program is also displayed when IPRINT is nonzero.

INFO - On exit from L2ONF, INFO will have one of the following integer values to indicate the reason for leaving the routine:
$I N F O=1$ SOL is feasible, and the condition that depends on ACC is satisfied.
$I N F O=2$ SOL is feasible, and rounding errors are preventing further progress.
$I N F O=3 S O L$ is feasible, but the objective function fails to decrease although a decrease is predicted by the current gradient vector.

INFO $=4$ In this case, the calculation cannot begin because LDA is less than NCON or because the lower bound on a variable is greater than the upper bound.
$I N F O=5$ This value indicates that the equality constraints are inconsistent. These constraints include any components of $x($.$) that are frozen by setting$ $\mathrm{XL}(\mathrm{I})=\mathrm{XU}(\mathrm{I})$.
$I N F O=6$ In this case there is an error return because the equality constraints and the bounds on the variables are found to be inconsistent.

INFO $=7$ This value indicates that there is no vector of variables that satisfies all of the constraints. Specifically, when this return or an INFO $=6$ return occurs, the current active constraints (whose indices are $\operatorname{IACT}(\mathrm{K}), \mathrm{K}=1, \ldots, \mathrm{NACT}$ ) prevent any change in $\mathrm{X}($.$) that reduces the sum of constraint violations. Bounds are$ only included in this sum if $\operatorname{INFO}=6$.
$I N F O=8$ Maximum number of function evaluations exceeded.
$I N F O=9$ The variables are determined by the equality constraints.

## Example

The problem from Schittkowski (1987)

$$
\begin{gathered}
\min f(x)=-x_{1} x_{2} x_{3} \\
\text { subject to } \quad-x_{1}-2 x_{2}-2 x_{3} \leq 0 \\
x_{1}+2 x_{2}+2 x_{3} \leq 72 \\
0 \leq x_{1} \leq 20 \\
0
\end{gathered}
$$

is solved with an initial guess $x_{1}=10, x_{2}=10$ and $x_{3}=10$.

```
USE LCONF INT
USE UMACH_
IMPLICIT NONE
INTEGER NCON, NEQ, NVAR
PARAMETER (NCON=2, NEQ=0, NVAR=3)
    INTEGER MAXFCN, NOUT
    REAL A(NCON,NVAR), ACC, B (NCON), OBJ, &
        SOL (NVAR), XGUESS (NVAR), XLB (NVAR), XUB (NVAR)
    EXTERNAL FCN
        Set values for the following problem.
```

$!$

```
!
!
!
!
!
    DATA A/-1.0, 1.0, -2.0, 2.0, -2.0, 2.0/, B/0.0, 72.0/
    DATA XLB/3*0.0/, XUB/20.0, 11.0, 42.0/, XGUESS/3*10.0/
    DATA ACC/0.0/, MAXFCN/400/
!
    CALL UMACH (2, NOUT)
    CALL LCONF (FCN, NEQ, A, B, XLB, XUB, SOL, XGUESS=XGUESS, &
        MAXFCN=MAXFCN, ACC=ACC, OBJ=OBJ)
!
    WRITE (NOUT,99998) 'Solution:'
    WRITE (NOUT,99999) SOL
    WRITE (NOUT,99998) 'Function value at solution:'
    WRITE (NOUT,99999) OBJ
    WRITE (NOUT,99998) 'Number of function evaluations:', MAXFCN
    STOP
99998 FORMAT (//, ' ', A, I4)
99999 FORMAT (1X, 5F16.6)
    END
!
    SUBROUTINE FCN (N, X, F)
    INTEGER N
    REAL X(*), F
!
    F}=-X(1)*X(2)*X(3
    RETURN
    END
```


## Output

```
Solution:
\[
\begin{array}{lll}
20.000000 & 11.000000 & 15.000000
\end{array}
\]
Function value at solution:
\(-3300.000000\)
Number of function evaluations: 5
```


## LCONG

Minimizes a general objective function subject to linear equality/inequality constraints.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Value of NVAR. (Input)

X - Vector of length N at which point the function is evaluated. (Input) $x$ should not be changed by FCN.

F - The computed function value at the point x . (Output)
FCN must be declared EXTERNAL in the calling program.
GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where
N - Value of NVAR. (Input)

X - Vector of length N at which point the function is evaluated. (Input) $X$ should not be changed by GRAD.
$\mathrm{G}-$ Vector of length N containing the values of the gradient of the objective function evaluated at the point X . (Output)

GRAD must be declared EXTERNAL in the calling program.

NEQ - The number of linear equality constraints. (Input)
A - NCON by NVAR matrix. (Input)
The matrix contains the equality constraint gradients in the first NEQ rows, followed by the inequality constraint gradients.
$\boldsymbol{B}$ - Vector of length NCON containing right-hand sides of the linear constraints. (Input) Specifically, the constraints on the variables X(I), $I=1, \ldots$, NVAR are $A(K, 1) * X(1)+\ldots+A(K, N V A R) * X(N V A R) \cdot E Q \cdot B(K), K=1, \ldots$, NEQ.A(K, 1) * X(1) + $\ldots+$ A (K, NVAR) * X(NVAR).LE.B(K), $K=N E Q+1, \ldots, N C O N$. Note that the data that define the equality constraints come before the data of the inequalities.
$\boldsymbol{X L B}$ - Vector of length NVAR containing the lower bounds on the variables; choose a very large negative value if a component should be unbounded below or set $\mathrm{XLB}(\mathrm{I})=\mathrm{XUB}(\mathrm{I})$ to freeze the I -th variable. (Input) Specifically, these simple bounds are XLB(I).LE.X(I), I = $1, \ldots$, NVAR.
$\boldsymbol{X U B}$ - Vector of length NVAR containing the upper bounds on the variables; choose a very large positive value if a component should be unbounded above. (Input) Specifically, these simple bounds are X(I). LE. XUB(I), I = $1, \ldots$, NVAR.

SOL - Vector of length NVAR containing solution. (Output)

## Optional Arguments

NVAR - The number of variables. (Input)
Default: NVAR = SIZE (A,2).
NCON - The number of linear constraints (excluding simple bounds). (Input)
Default: NCON $=\operatorname{SIZE}(A, 1)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
XGUESS - Vector of length NVAR containing the initial guess of the minimum. (Input)
Default: XGUESS $=0.0$.
$A C C$ - The nonnegative tolerance on the first order conditions at the calculated solution.
(Input)
Default: $\mathrm{ACC}=1 . \mathrm{e}-4$ for single precision and 1.d-8 for double precision.
MAXFCN - On input, maximum number of function evaluations allowed.(Input/ Output) On output, actual number of function evaluations needed.
Default: $\operatorname{MAXFCN}=400$.
$\mathbf{O B J}$ - Value of the objective function. (Output)
NACT - Final number of active constraints. (Output)
IACT - Vector containing the indices of the final active constraints in the first NACT positions. (Output)
Its length must be at least $\mathrm{NCON}+2$ * NVAR.
ALAMDA - Vector of length NVAR containing the Lagrange multiplier estimates of the final active constraints in the first NACT positions. (Output)

## FORTRAN 90 Interface

Generic: CALL LCONG (FCN, GRAD, NEQ, A, B, XLB, XUB, SOL [,...])
Specific: The specific interface names are $S_{\text {_LCONG }}$ and D_LCONG.

## FORTRAN 77 Interface

Single: CALL LCONG (FCN, GRAD, NVAR, NCON, NEQ, A, LDA, B, XLB,
XUB, XGUESS, ACC, MAXFCN, SOL, OBJ, NACT, IACT, ALAMDA)
Double: $\quad$ The double precision name is DLCONG.

## Description

The routine LCONG is based on M.J.D. Powell's TOLMIN, which solves linearly constrained optimization problems, i.e., problems of the form

$$
\min _{x \in \mathrm{R}^{n}} f(x)
$$

subject to $\quad A_{1} x=b_{1}$

$$
\begin{gathered}
A_{2} x \leq b_{2} \\
x_{l} \leq x \leq x_{u}
\end{gathered}
$$

given the vectors $b_{1}, b_{2}, x_{l}$ and $x_{u}$ and the matrices $A_{1}$, and $A_{2}$.
The algorithm starts by checking the equality constraints for inconsistency and redundancy. If the equality constraints are consistent, the method will revise $x^{0}$, the initial guess provided by the user, to satisfy

$$
A_{1} x=b_{1}
$$

Next, $x^{0}$ is adjusted to satisfy the simple bounds and inequality constraints. This is done by solving a sequence of quadratic programming subproblems to minimize the sum of the constraint or bound violations.

Now, for each iteration with a feasible $x_{k}$, let $J_{k}$ be the set of indices of inequality constraints that have small residuals. Here, the simple bounds are treated as inequality constraints. Let $I_{k}$ be the set of indices of active constraints. The following quadratic programming problem

$$
\begin{aligned}
& \min f\left(x^{k}\right)+d^{T} \nabla f\left(x^{k}\right)+\frac{1}{2} d^{T} B^{k} d \\
& \text { subject to } \quad a_{j} d=0 \quad j \in I_{k} \\
& a_{j} d \leq 0 \quad j \in J_{k}
\end{aligned}
$$

is solved to get $\left(d^{k}, \lambda^{k}\right)$ where $a_{j}$ is a row vector representing either a constraint in $A_{1}$ or $A_{2}$ or a bound constraint on $x$. In the latter case, the $a_{j}=e_{i}$ for the bound constraint $x_{i} \leq\left(x_{u}\right)_{i}$ and $a_{j}=-e_{i}$ for the constraint $-x_{i} \leq\left(-x_{l}\right)_{i}$. Here, $e_{i}$ is a vector with a 1 as the $i$-th component, and zeroes elsewhere. $\lambda^{k}$ are the Lagrange multipliers, and $B^{k}$ is a positive definite approximation to the second derivative $\nabla^{2} f\left(x^{k}\right)$.
After the search direction $d^{k}$ is obtained, a line search is performed to locate a better point. The new point $x^{k+1}=x^{k}+\alpha^{k} d^{k}$ has to satisfy the conditions

$$
f\left(x^{k}+\alpha^{k} d^{k}\right) \leq f\left(x^{k}\right)+0.1 \alpha^{k}\left(d^{k}\right)^{T} \nabla f\left(x^{k}\right)
$$

and

$$
\left(d^{k}\right)^{T} \nabla f\left(x^{k}+\alpha^{k} d^{k}\right) \geq 0.7\left(d^{k}\right)^{T} \nabla f\left(x^{k}\right)
$$

The main idea in forming the set $J_{k}$ is that, if any of the inequality constraints restricts the steplength $\alpha^{k}$, then its index is not in $J_{k}$. Therefore, small steps are likely to be avoided.

Finally, the second derivative approximation, $B^{k}$, is updated by the BFGS formula, if the condition

$$
\left(d^{k}\right)^{T} \nabla f\left(x^{k}+\alpha^{k} d^{k}\right)-\nabla f\left(x^{k}\right)>0
$$

holds. Let $x^{k} \leftarrow x^{k+1}$, and start another iteration.
The iteration repeats until the stopping criterion

$$
\left\|\nabla f\left(x^{k}\right)-A^{k} \lambda^{k}\right\|_{2} \leq \tau
$$

is satisfied; here, $\tau$ is a user-supplied tolerance. For more details, see Powell $(1988,1989)$.

## Comments

1. Workspace may be explicitly provided, if desired, by use of L2ONG/DL2ONG. The reference is:

CALL L2ONG (FCN, GRAD, NVAR, NCON, NEQ, A, LDA, B, XLB, XUB, XGUESS, ACC, MAXFCN, SOL, OBJ, NACT, IACT, ALAMDA, IPRINT, INFO, WK)

The additional arguments are as follows:
IPRINT — Print option (see Comment 3). (Input)
INFO - Informational flag (see Comment 3). (Output)
$\boldsymbol{W K}-$ Real work vector of length NVAR**2 +11 * NVAR + NCON.
2. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 4 | The equality constraints are inconsistent. |
| 4 | 5 | The equality constraints and the bounds on the variables are found to <br> be inconsistent. |
| 4 | 6 | No vector X satisfies all of the constraints. In particular, the current <br> active constraints prevent any change in X that reduces the sum of <br> constraint violations. |
| 4 | 7 | Maximum number of function evaluations exceeded. |
| 4 | 9 | The variables are determined by the equality constraints. |

3. The following are descriptions of the arguments IPRINT and INFO:

IPRINT - This argument must be set by the user to specify the frequency of printing during the execution of the routine LCONG. There is no printed output if IPRINT $=0$. Otherwise, after ensuring feasibility, information is given every IABS(IPRINT) iterations and whenever a parameter called TOL is reduced. The printing provides the values of $X(),. F($.$) and G()=.\operatorname{GRAD}(F)$ if IPRINT is positive. If IPRINT is negative, this information is augmented by the current values of $\operatorname{IACT}(\mathrm{K}) \mathrm{K}=1, \ldots, \mathrm{NACT}, \operatorname{PAR}(\mathrm{K}) \mathrm{K}=1, \ldots$, NACT and $\operatorname{RESKT}(\mathrm{I}) \mathrm{I}=1, \ldots, \mathrm{~N}$. The reason for returning to the calling program is also displayed when IPRINT is nonzero.

INFO - On exit from L2ONG, INFO will have one of the following integer values to indicate the reason for leaving the routine:
$I N F O=1 \quad S O L$ is feasible and the condition that depends on ACC is satisfied.
$I N F O=2 S O L$ is feasible and rounding errors are preventing further progress.
$I N F O=3 S O L$ is feasible but the objective function fails to decrease although a decrease is predicted by the current gradient vector.
$\operatorname{INFO}=4$ In this case, the calculation cannot begin because LDA is less than NCON or because the lower bound on a variable is greater than the upper bound.
$I N F O=5$ This value indicates that the equality constraints are inconsistent. These constraints include any components of $x($.$) that are frozen$ by setting $\mathrm{XL}(\mathrm{I})=\mathrm{XU}(\mathrm{I})$.

INFO $=6$ In this case, there is an error return because the equality constraints and the bounds on the variables are found to be inconsistent.

INFO $=7$ This value indicates that there is no vector of variables that satisfies all of the constraints. Specifically, when this return or an INFO $=6$ return occurs, the current active constraints (whose indices are $\operatorname{IACT}(\mathrm{K}), \mathrm{K}=1, \ldots, \mathrm{NACT})$ prevent any change in $\mathrm{X}($. that reduces the sum of constraint violations, where only bounds are included in this sum if $I N F O=6$.
$I N F O=8$ Maximum number of function evaluations exceeded.
$I N F O=9$ The variables are determined by the equality constraints.

## Example

The problem from Schittkowski (1987)

$$
\min f(x)=-x_{1} x_{2} x_{3}
$$

subject to $\quad-x_{1}-2 x_{2}-2 x_{3} \leq 0$

$$
\begin{aligned}
& x_{1}+2 x_{2}+2 x_{3} \leq 72 \\
& 0 \leq x_{1} \leq 20 \\
& 0 \leq x_{2} \leq 11 \\
& 0 \leq x_{3} \leq 42
\end{aligned}
$$

is solved with an initial guess $x_{1}=10, x_{2}=10$ and $x_{3}=10$.

```
    USE LCONG_INT
    USE UMACH_INT
    IMPLICIT NONE
!
    INTEGER NCON, NEQ, NVAR
    PARAMETER (NCON=2, NEQ=0, NVAR=3)
!
    INTEGER MAXFCN, NOUT
    REAL A(NCON,NVAR), ACC, B (NCON), OBJ, &
        SOL (NVAR), XGUESS (NVAR), XLB (NVAR), XUB (NVAR)
    EXTERNAL FCN, GRAD
                                    Set values for the following problem.
                                    Min -X(1)*X(2)*X(3)
                -X(1) - 2*X(2) - 2*X(3) .LE. 0
                        X(1) + 2*X(2) + 2*X(3) .LE. 72
                0 .LE. X(1) .LE. 20
                0 .LE. X(2) .LE. 11
                O .LE. X(3) .LE. 42
    DATA A/-1.0, 1.0, -2.0, 2.0, -2.0, 2.0/, B/0.0, 72.0/
    DATA XLB/3*0.0/, XUB/20.0, 11.0, 42.0/, XGUESS/3*10.0/
    DATA ACC/0.0/, MAXFCN/400/
    CALL UMACH (2, NOUT)
    CALL LCONG (FCN, GRAD, NEQ, A, B, XLB, XUB, SOL, XGUESS=XGUESS, &
            ACC=ACC, MAXFCN=MAXFCN, OBJ=OBJ)
!
    WRITE (NOUT,99998) 'Solution:'
    WRITE (NOUT,99999) SOL
    WRITE (NOUT,99998) 'Function value at solution:'
    WRITE (NOUT,99999) OBJ
    WRITE (NOUT,99998) 'Number of function evaluations:', MAXFCN
    STOP
99998 FORMAT (//, ' ', A, I4)
99999 FORMAT (1X, 5F16.6)
    END
!
    SUBROUTINE FCN (N, X, F)
    INTEGER N
    REAL X(*), F
!
    F = -X(1)*X(2)*X(3)
    RETURN
    END
!
    SUBROUTINE GRAD (N, X, G)
    INTEGER N
    REAL X(*), G(*)
!
```

```
G(1) = -X(2)*X(3)
G(2) = -X(1)*X(3)
G(3) = -X(1)*X(2)
RETURN
END
```


## Output

```
Solution:
20.000000 11.000000 15.000000
Function value at solution:
-3300.000000
Number of function evaluations: 5
```


## NNLPF

Solves a general nonlinear programming problem using a sequential equality constrained quadratic programming method.

## Required Arguments

FCN - User-supplied subroutine to evaluate the objective function and constraints at a given point. The internal usage is CALL FCN (X, IACT, RESULT, IERR), where
$\boldsymbol{X}$ - The point at which the objective function or constraint is evaluated. (Input)
IACT - Integer indicating whether evaluation of the objective function is requested or evaluation of a constraint is requested. If $\operatorname{IACT}$ is zero, then an objective function evaluation is requested. If IACT is nonzero then the value if IACT indicates the index of the constraint to evaluate. (Input)

RESULT - If IACT is zero, then RESULT is the computed function value at the point x. If IACT is nonzero, then RESULT is the computed constraint value at the point X . (Output)

IERR - Logical variable. On input IERR is set to .FALSE. If an error or other undesirable condition occurs during evaluation, then IERR should be set to .TRUE. Setting IERR to . TRUE. will result in the step size being reduced and the step being tried again. (If IERR is set to . TRUE. for xGuess, then an error is issued.)

The routine FCN must be use-associated in a user module that uses NNLPF_INT, or else declared EXTERNAL in the calling program. If FCN is a separately compiled routine, not in a module, then it must be declared External.

M - Total number of constraints. (Input)

ME - Number of equality constraints. (Input)
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

$0 \quad$ User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable; all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on variables. (Input, if IBTYPE $=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ )
If there is no lower bound for a variable, then the corresponding XLB value should be set to -Huge(X(1)).
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on variables. (Input, if IBTYPE $=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ ).
If there is no upper bound for a variable, then the corresponding XUB value should be set to Huge(X(1)).
$\boldsymbol{X}-$ Vector of length N containing the computed solution. (Output)

## Optional Arguments

N - Number of variables. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X})$.
$\boldsymbol{X G U E S S}$ - Vector of length N containing an initial guess of the solution. (Input) Default: XGUESS $=x$, (with the smallest value of $\|x\|_{2}$ ) that satisfies the bounds.

XSCALE - Vector of length N setting the internal scaling of the variables. The initial value given and the objective function and gradient evaluations however are always in the original unscaled variables. The first internal variable is obtained by dividing values X(I) by XSCALE (I). (Input)
In the absence of other information, set all entries to 1.0.
Default: XSCALE (: ) = 1.0.
IPRINT - Parameter indicating the desired output level. (Input)

## IPRINT Action

0 No output printed.

One line of intermediate results is printed in each iteration.
2 Lines of intermediate results summarizing the most important data for each step are printed.

3 Lines of detailed intermediate results showing all primal and dual variables, the relevant values from the working set, progress in the backtracking and etc are printed

4 Lines of detailed intermediate results showing all primal and dual variables, the relevant values from the working set, progress in the backtracking, the gradients in the working set, the quasi-Newton updated and etc are printed.

Default: $\operatorname{IPRINT}=0$.

MAXITN - Maximum number of iterations allowed. (Input)
Default: MAXITN $=200$.
EPSDIF - Relative precision in gradients. (Input)
Default: EPSDIF $=\operatorname{EPSILON}(x(1))$
TAUO - A universal bound describing how much the unscaled penalty-term may deviate from zero. (Input)
NNLPF assumes that within the region described by

$$
\sum_{i=1}^{M_{e}}\left|g_{i}(x)\right|-\sum_{i=M_{e}+1}^{M} \min \left(0, g_{i}(x)\right) \leq \text { TAU } 0
$$

all functions may be evaluated safely. The initial guess, however, may violate these requirements. In that case an initial feasibility improvement phase is run by NNLPF until such a point is found. A small TAU0 diminishes the efficiency of NNLPF, because the iterates then will follow the boundary of the feasible set closely. Conversely, a large TAUO may degrade the reliability of the code.
Default TAU0 $=1 . \mathrm{E} 0$
DELO - In the initial phase of minimization a constraint is considered binding if

$$
\frac{g_{i}(x)}{\max \left(1,\left\|\nabla g_{i}(x)\right\|\right)} \leq \text { DELO } \quad i=M_{e}+1, \ldots, M
$$

Good values are between .01 and 1.0. If DEL0 is chosen too small then identification of the correct set of binding constraints may be delayed. Contrary, if DEL0 is too large, then the method will often escape to the full regularized SQP method, using individual slack variables for any active constraint, which is quite costly. For well-scaled problems $D E L 0=1.0$ is reasonable. (Input)
Default: DELO $=.5{ }^{*}$ TAUO

EPSFCN - Relative precision of the function evaluation routine. (Input)
Default: EPSFCN = epsilon(x(1))
IDTYPE - Type of numerical differentiation to be used. (Input)
Default: $\operatorname{IDTYPE}=1$

## IDTYPE Action

1 Use a forward difference quotient with discretization stepsize $0.1\left(\right.$ EPSFCN $^{1 / 2}$ ) componentwise relative.

2 Use the symmetric difference quotient with discretization stepsize $0.1\left(\right.$ EPSFCN $^{1 / 3}$ ) componentwise relative

3 Use the sixth order approximation computing a Richardson extrapolation of three symmetric difference quotient values. This uses a discretization stepsize $0.01\left(\right.$ EPSFCN $^{1 / 7}$ )

TAUBND - Amount by which bounds may be violated during numerical differentiation. Bounds are violated by TAUBND (at most) only if a variable is on a bound and finite differences are taken for gradient evaluations. (Input)
Default: $T$ AUBND $=1 . E 0$
SMALLW - Scalar containing the error allowed in the multipliers. For example, a negative multiplier of an inequality constraint is accepted (as zero) if its absolute value is less than SMALLW. (Input)
Default: SMALLW $=\exp (2 * \log (\operatorname{epsilon}(x(1) / 3)))$
DELMIN - Scalar which defines allowable constraint violations of the final accepted result. Constraints are satisfied if $\left|\mathrm{g}_{\mathrm{i}}(\mathrm{x})\right| \leq$ DELMIN , and $\mathrm{g}_{\mathrm{j}}(\mathrm{x}) \geq(-$ DELMIN $)$ respectively. (Input)
Default: DELMIN $=\min ($ DEL0/10, $\max (E P S D I F, \min ($ DEL0/10, $\max (1 . E-6 *$ DEL0, SMALLW) )) )

SCFMAX - Scalar containing the bound for the internal automatic scaling of the objective function. (Intput)
Default: SCFMAX $=1.0 \mathrm{E} 4$
FVALUE - Scalar containing the value of the objective function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL NNLPF (FCN, M, ME, IBTYPE, XLB, XUB, X [,...])
Specific: The specific interface names are $S_{-}$NNLPF and $D_{-}$NNLPF .

## Description

The routine NNLPF provides an interface to a licensed version of subroutine DONLP2, a FORTRAN code developed by Peter Spellucci (1998). It uses a sequential equality constrained quadratic programming method with an active set technique, and an alternative usage of a fully regularized mixed constrained subproblem in case of nonregular constraints (i.e. linear dependent gradients in the "working sets"). It uses a slightly modified version of the Pantoja-Mayne update for the Hessian of the Lagrangian, variable dual scaling and an improved Armjijo-type stepsize algorithm. Bounds on the variables are treated in a gradient-projection like fashion. Details may be found in the following two papers:
P. Spellucci: An SQP method for general nonlinear programs using only equality constrained subproblems. Math. Prog. 82, (1998), 413-448.
P. Spellucci: A new technique for inconsistent problems in the SQP method. Math. Meth. of Oper. Res. 47, (1998), 355-500. (published by Physica Verlag, Heidelberg, Germany).
The problem is stated as follows:

$$
\begin{array}{lll}
\text { subject to } & g_{j}(x)=0, \text { for } & j=1, \ldots, m_{e} \\
& g_{j}(x) \geq 0, \text { for } & j=m_{e}+1, \ldots, m \\
& x_{l} \leq x \leq x_{u} &
\end{array}
$$

$$
\min _{x \in \mathbf{R}^{n}} f(x)
$$

Although default values are provided for optional input arguments, it may be necessary to adjust these values for some problems. Through the use of optional arguments, NNLPF allows for several parameters of the algorithm to be adjusted to account for specific characteristics of problems. The DONLP2 Users Guide provides detailed descriptions of these parameters as well as strategies for maximizing the perfomance of the algorithm. The DONLP2 Users Guide is available in the "help" subdirectory of the main IMSL product installation directory. In addition, the following are a number of guidelines to consider when using NNLPF.

- A good initial starting point is very problem specific and should be provided by the calling program whenever possible. See optional argument XGUESS.
- Gradient approximation methods can have an effect on the success of NNLPF. Selecting a higher order appoximation method may be necessary for some problems. See optional argument IDTYPE.
- If a two sided constraint $l_{i} \leq g_{i}(x) \leq u_{i}$ is transformed into two constraints $g_{2 i}(x) \geq 0$ and $g_{2 i+1}(x) \geq 0$, then choose $\operatorname{DEL} 0<\frac{1}{2}\left(u_{i}-l_{i}\right) / \max \left\{1,\left\|\nabla g_{i}(x)\right\|\right\}$, or at least try to provide an estimate for that value. This will increase the efficiency of the algorithm. See optional argument DELO.
- The parameter IERR provided in the interface to the user supplied function FCN can be very useful in cases when evaluation is requested at a point that is not possible or reasonable. For example, if evaluation at the requested point would result in a floating point exception, then setting IERR to .TRUE. and returning without performing the evaluation will avoid the exception. NNLPF will then reduce the stepsize and try the step again. Note, if IERR is set to .TRUE. for the initial guess, then an error is issued.


## Example

The problem

$$
\begin{gathered}
\min F(x)=\left(x_{1}-2\right)^{2}+\left(x_{2}-1\right)^{2} \\
\text { subject to } \quad g_{1}(x)=x_{1}-2 x_{2}+1=0 \\
g_{2}(x)=-x_{1}^{2} / 4-x_{2}^{2}+1 \geq 0
\end{gathered}
$$

is solved.

```
USE NNLPF_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER IBTYPE, M, ME
PARAMETER (IBTYPE=0, M=2, ME=1)
!
REAL(KIND(1E0)) FVALUE, X(2), XGUESS(2), XLB(2), XUB(2)
EXTERNAL FCN
!
XLB = -HUGE (X(1))
XUB = HUGE(X(1))
!
CALL NNLPF (FCN, M, ME, IBTYPE, XLB, XUB, X)
!
CALL WRRRN ('The solution is', X)
END
SUBROUTINE FCN (X, IACT, RESULT, IERR)
INTEGER IACT
REAL(KIND(1E0)) X(*), RESULT
LOGICAL IERR
!
SELECT CASE (IACT)
CASE (0)
        RESULT = (X(1)-2.0E0)**2 + (X(2)-1.0E0)**2
CASE (1)
        RESULT = X(1) - 2.0E0*X(2) + 1.0E0
CASE (2)
    RESULT = -(X(1)**2)/4.0E0 - X(2)**2 + 1.0E0
END SELECT
RETURN
END
```


## Output

```
The solution is
        0.8229
    20.9114
```


## NNLPG

Solves a general nonlinear programming problem using a sequential equality constrained quadratic programming method with user supplied gradients.

## Required Arguments

FCN - User-supplied subroutine to evaluate the objective function and constraints at a given point. The internal usage is CALL FCN (X, IACT, RESULT, IERR), where
$\boldsymbol{X}$ - The point at which the objective function or constraint is evaluated. (Input)
IACT - Integer indicating whether evaluation of the objective function is requested or evaluation of a constraint is requested. If IACT is zero, then an objective function evaluation is requested. If IACT is nonzero then the value if IACT indicates the index of the constraint to evaluate. (Input)

RESULT - If IACT is zero, then RESULT is the computed objective function value at the point $X$. If IACT is nonzero, then RESULT is the computed constraint value at the point X . (Output)

IERR - Logical variable. On input IERR is set to .FALSE. If an error or other undesirable condition occurs during evaluation, then IERR should be set to .TRUE. Setting IERR to . TRUE. will result in the step size being reduced and the step being tried again. (If IERR is set to . TRUE. for XGUESS, then an error is issued.)

The routine FCN must be use-associated in a user module that uses NNLPG_INT, or else declared EXTERNAL in the calling program. If FCN is a separately compiled routine, not in a module, then it must be declared EXTERNAL.

GRAD - User-supplied subroutine to evaluate the gradients at a given point. The usage is CALL GRAD (X, IACT, RESULT), where
$\boldsymbol{X}$ - The point at which the gradient of the objective function or gradient of a constraint is evaluated. (Input)

IACT - Integer indicating whether evaluation of the function gradient is requested or evaluation of a constraint gradient is requested. If IACT is zero, then an objective function gradient evaluation is requested. If $I A C T$ is nonzero then the value if IACT indicates the index of the constraint gradient to evaluate. (Input)RESULT - If IACT is zero, then RESULT is the computed gradient of the objective function at the point X . If IACT is nonzero, then RESULT is the computed gradient of the requested constraint value at the point X . (Output)

The routine GRAD must be use-associated in a user module that uses NNLPG_INT, or else declared EXTERNAL in the calling program. If GRAD is a separately compiled routine, not in a module, then is must be declared EXTERNAL
$\boldsymbol{M}$ — Total number of constraints. (Input)
$M E$ - Number of equality constraints. (Input)
IBTYPE - Scalar indicating the types of bounds on variables. (Input)

## IBTYPE Action

0 User will supply all the bounds.
1 All variables are nonnegative.
2 All variables are nonpositive.
3 User supplies only the bounds on 1st variable, all other variables will have the same bounds.
$\boldsymbol{X L B}$ - Vector of length N containing the lower bounds on the variables. (Input, if IBTYPE $=0$; output, if $\operatorname{IBTYPE}=1$ or 2 ; input/output, if $I B T Y P E=3$ ) If there is no lower bound on a variable, then the corresponding XLB value should be set to -huge(x(1)).
$\boldsymbol{X U B}$ - Vector of length N containing the upper bounds on the variables. (Input, if $I B T Y P E=0$; output, if $I B T Y P E=1$ or 2 ; input/output, if $I B T Y P E=3$ ) If there is no upper bound on a variable, then the corresponding XUB value should be set to huge(x(1)).
$\boldsymbol{X}$ - Vector of length N containing the computed solution. (Output)

## Optional Arguments

$N$ - Number of variables. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X})$.
IPRINT - Parameter indicating the desired output level. (Input)

## IPRINT Action

$0 \quad$ No output printed.
1 One line of intermediate results is printed in each iteration.
2 Lines of intermediate results summarizing the most important data for each step are printed.

Lines of detailed intermediate results showing all primal and dual variables, the relevant values from the working set, progress in the backtracking and etc are printed

Lines of detailed intermediate results showing all primal and dual variables, the relevant values from the working set, progress in the backtracking, the gradients in the working set, the quasi-Newton updated and etc are printed.

Default: $\operatorname{IPRINT}=0$.

MAXITN - Maximum number of iterations allowed. (Input)
Default: $\operatorname{MAXITN}=200$.

XGUESS - Vector of length N containing an initial guess of the solution. (Input)
Default: XGUESS $=x$, (with the smallest value of $\|x\|_{2}$ ) that satisfies the bounds.

TAUO - A universal bound describing how much the unscaled penalty-term may deviate from zero. (Input)
NNLPG assumes that within the region described by

$$
\sum_{i=1}^{M_{e}}\left|g_{i}(x)\right|-\sum_{i=M_{e}+1}^{M} \min \left(0, g_{i}(x)\right) \leq \text { TAU } 0
$$

all functions may be evaluated safely. The initial guess however, may violate these requirements. In that case an initial feasibility improvement phase is run by NNLPG until such a point is found. A small TAU0 diminishes the efficiency of NNLPG, because the iterates then will follow the boundary of the feasible set closely. Conversely, a large TAUO may degrade the reliability of the code.
Default: TAUO = 1.E0

DELO - In the initial phase of minimization a constraint is considered binding if

$$
\frac{g_{i}(x)}{\max \left(1,\left\|\nabla g_{i}(x)\right\|\right)} \leq \mathrm{DELO} \quad i=M_{e}+1, \ldots, M
$$

Good values are between .01 and 1.0. If DEL0 is chosen too small then identification of the correct set of binding constraints may be delayed. Contrary, if DELO is too large, then the method will often escape to the full regularized SQP method, using individual slack variables for any active constraint, which is quite costly. For well-scaled problems $D E L 0=1.0$ is reasonable. (Input)
Default: DELO $=.5{ }^{*}$ TAUO
SMALLW - Scalar containing the error allowed in the multipliers. For example, a negative multiplier of an inequality constraint is accepted (as zero) if its absolute value is less than SMALLW. (Input)
Default: SMALLW $=\exp (2 * \log (\operatorname{epsilon}(x(1) / 3)))$

DELMIN - Scalar which defines allowable constraint violations of the final accepted result. Constraints are satisfied if $\left|\mathrm{g}_{\mathrm{i}}(\mathrm{x})\right| \leq$ DELMIN , and $\mathrm{g}_{\mathrm{j}}(\mathrm{x}) \geq(-$ DELMIN $)$ respectively. (Input)
Default: DELMIN $=\min ($ DEL0 $/ 10, \max (1 . E-6 *$ DEL0, SMALLW $))$
SCFMAX - Scalar containing the bound for the internal automatic scaling of the objective function. (Intput)
Default: $\operatorname{SCFMAX}=1.0 \mathrm{E} 4$
FVALUE - Scalar containing the value of the objective function at the computed solution. (Output)

## FORTRAN 90 Interface

Generic: CALL NNLPG (FCN, GRAD, M, ME, IBTYPE, XLB, XUB, X [,$\ldots]$ )
Specific: The specific interface names are S_NNLPG and D_NNLPG.

## Description

The routine NNLPG provides an interface to a licensed version of subroutine DONLP2, a FORTRAN code developed by Peter Spellucci (1998). It uses a sequential equality constrained quadratic programming method with an active set technique, and an alternative usage of a fully regularized mixed constrained subproblem in case of nonregular constraints (i.e. linear dependent gradients in the "working sets"). It uses a slightly modified version of the Pantoja-Mayne update for the Hessian of the Lagrangian, variable dual scaling and an improved Armjijo-type stepsize algorithm. Bounds on the variables are treated in a gradient-projection like fashion. Details may be found in the following two papers:
P. Spellucci: An SQP method for general nonlinear programs using only equality constrained subproblems. Math. Prog. 82, (1998), 413-448.
P. Spellucci: A new technique for inconsistent problems in the SQP method. Math. Meth. of Oper. Res. 47, (1998), 355-500. (published by Physica Verlag, Heidelberg, Germany).
The problem is stated as follows:

$$
\min _{x \in \mathbb{R}^{n}} f(x)
$$

subject to

$$
\begin{array}{ll}
g_{j}(x)=0, \text { for } & j=1, \ldots, m_{e} \\
g_{j}(x) \geq 0, \text { for } & j=m_{e}+1, \ldots, m \\
x_{l} \leq x \leq x_{u} &
\end{array}
$$

Although default values are provided for optional input arguments, it may be necessary to adjust these values for some problems. Through the use of optional arguments, NNLPG allows for several parameters of the algorithm to be adjusted to account for specific characteristics of problems. The DONLP2 Users Guide provides detailed descriptions of these parameters as well as strategies for maximizing the perfomance of the algorithm. The DONLP2 Users Guide is available in the
"help" subdirectory of the main IMSL product installation directory. In addition, the following are a number of guidelines to consider when using NNLPG.

- A good initial starting point is very problem specific and should be provided by the calling program whenever possible. See optional argument XGUESS.
- If a two sided constraint $l_{i} \leq g_{i}(x) \leq u_{i}$ is transformed into two constraints $g_{2 i}(x) \geq 0$ and $g_{2 i+1}(x) \geq 0$, then choose DEL $0<\frac{1}{2}\left(u_{i}-l_{i}\right) / \max \left\{1,\left\|\nabla g_{i}(x)\right\|\right\}$, or at least try to provide an estimate for that value. This will increase the efficiency of the algorithm. See optional argument DELO.
- The parameter IERR provided in the interface to the user supplied function FCN can be very useful in cases when evaluation is requested at a point that is not possible or reasonable. For example, if evaluation at the requested point would result in a floating point exception, then setting IERR to .TRUE. and returning without performing the evaluation will avoid the exception. NNLPG will then reduce the stepsize and try the step again. Note, if IERR is set to . TRUE. for the initial guess, then an error is issued.


## Comments

1. Informational errors

| Type | Code |  |
| :---: | :---: | :--- |
| 4 | 1 | Constraint evaluation returns an error with current point. |
| 4 | 2 | Objective evaluation returns an error with current point. |
| 4 | 3 | Working set is singular in dual extended QP. |
| 4 | 4 | QP problem is seemingly infeasible. |
| 4 | 5 | A stationary point located. |
| 4 | 6 | A stationary point located or termination criteria too strong. |
| 4 | 7 | Maximum number of iterations exceeded. |
| 4 | 8 | Stationary point not feasible. |
| 4 | 9 | Very slow primal progress. |
| 4 | 10 | The problem is singular. |
| 4 | 11 | Matrix of gradients of binding constraints is singular or very ill- |
| 4 | 12 | conditioned. |
| 4 | Small changes in the penalty function. |  |

## Example 1

The problem

$$
\begin{gathered}
\min F(x)=\left(x_{1}-2\right)^{2}+\left(x_{2}-1\right)^{2} \\
\text { subject to } \quad g_{1}(x)=x_{1}-2 x_{2}+1=0 \\
g_{2}(x)=-x_{1}^{2} / 4-x_{2}^{2}+1 \geq 0
\end{gathered}
$$

is solved.

```
USE NNLPG_INT
USE WRRRN_INT
```

```
IMPLICIT NONE
INTEGER IBTYPE, M, ME
PARAMETER (IBTYPE=0, M=2, ME=1)
!
REAL(KIND(1E0)) FVALUE, X(2), XGUESS(2), XLB(2), XUB(2)
EXTERNAL FCN, GRAD
!
XLB = -HUGE (X(1))
XUB = HUGE (X(1))
CALL NNLPG (FCN, GRAD, M, ME, IBTYPE, XLB, XUB, X)
!
CALL WRRRN ('The solution is', X)
END
SUBROUTINE FCN (X, IACT, RESULT, IERR)
INTEGER IACT
REAL(KIND(1E0)) X(*), RESULT
LOGICAL IERR
SELECT CASE (IACT)
CASE (0)
    RESULT = (X(1)-2.0E0)**2 + (X(2)-1.0E0)**2
CASE (1)
    RESULT = X(1) - 2.0E0*X(2) + 1.0E0
CASE (2)
    RESULT = -(X(1)**2)/4.0E0 - X(2)**2 + 1.0E0
END SELECT
RETURN
END
SUBROUTINE GRAD (X, IACT, RESULT)
INTEGER IACT
REAL(KIND(1E0)) X(*),RESULT(*)
!
SELECT CASE (IACT)
CASE (0)
    RESULT (1) = 2.0E0*(X(1)-2.0E0)
    RESULT (2) = 2.0EO*(X(2)-1.0E0)
CASE (1)
    RESULT (1) = 1.0E0
    RESULT (2) = -2.0E0
CASE (2)
    RESULT (1) = -0.5E0*X(1)
    RESULT (2) = -2.0E0*X(2)
END SELECT
RETURN
END
```


## Output

```
The solution is
1 0.8229
20.9114
```


## Additional Examples

## Example 2

The same problem from Example 1 is solved, but here we use central differences to compute the gradient of the first constraint. This example demonstrates how NNLPG can be used in cases when analytic gradients are known for only a portion of the constraints and/or objective function. The subroutine CDGRD is used to compute an approximation to the gradient of the first constraint.

```
USE NNLPG INT
USE CDGRD_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER IBTYPE, M, ME
PARAMETER (IBTYPE=0, M=2, ME=1)
REAL(KIND(1E0)) FVALUE, X(2), XGUESS (2), XLB(2), XUB(2)
EXTERNAL FCN, GRAD
XLB = -HUGE (X (1))
XUB = HUGE(X(1))
CALL NNLPG (FCN, GRAD, M, ME, IBTYPE, XLB, XUB, X)
CALL WRRRN ('The solution is', X)
END
SUBROUTINE FCN (X, IACT, RESULT, IERR)
INTEGER IACT
REAL (KIND (1E0)) X(2), RESULT
LOGICAL IERR
EXTERNAL CONSTR1
SELECT CASE (IACT)
CASE (0)
    RESULT = (X(1)-2.OEO)**2 + (X(2)-1.0EO)**2
CASE (1)
    CALL CONSTR1 (2, X, RESULT)
CASE (2)
    RESULT = - (X (1)**2)/4.0E0 - X(2)**2 + 1.0E0
END SELECT
RETURN
END
SUBROUTINE GRAD (X, IACT, RESULT)
USE CDGRD_INT
INTEGER IACT
REAL(KIND(1E0)) X (2),RESULT (2)
EXTERNAL CONSTR1
SELECT CASE (IACT)
CASE (0)
    RESULT (1) = 2.0E0*(X(1)-2.0E0)
    RESULT (2) = 2.0E0*(X(2)-1.0E0)
```

$!$

```
CASE (1)
    CALL CDGRD(CONSTR1, X, RESULT)
CASE (2)
    RESULT (1) = -0.5E0*X(1)
    RESULT (2) = -2.0E0*X(2)
END SELECT
RETURN
END
SUBROUTINE CONSTR1 (N, X, RESULT)
INTEGER N
REAL(KIND(1E0)) X(*), RESULT
RESULT = X(1) - 2.0E0*X(2) + 1.0E0
RETURN
END
```


## Output

The solution is
10.8229
$2 \quad 0.9114$

## CDGRD

Approximates the gradient using central differences.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
X - The point at which the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X C}$ - Vector of length N containing the point at which the gradient is to be estimated. (Input)
$G C$ - Vector of length N containing the estimated gradient at XC. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XC}, 1)$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
In the absence of other information, set all entries to 1.0.
Default: XSCALE = 1.0.

EPSFCN - Estimate for the relative noise in the function. (Input)
EPSFCN must be less than or equal to 0.1 . In the absence of other information, set EPSFCN to 0.0.
Default: $\operatorname{EPSFCN}=0.0$.

## FORTRAN 90 Interface

Generic: CALL CDGRD (FCN, XC, GC [, ...])

Specific: The specific interface names are S_CDGRD and D_CDGRD.

## FORTRAN 77 Interface

Single: CALL CDGRD (FCN, N, XC, XSCALE, EPSFCN, GC)
Double: The double precision name is DCDGRD.

## Description

The routine CDGRD uses the following finite-difference formula to estimate the gradient of a function of $n$ variables at $x$ :

$$
\frac{f\left(x+h_{i} e_{i}\right)-f\left(x-h_{i} e_{i}\right)}{2 h_{i}} \quad \text { for } i=1, \ldots, n
$$

where $h_{i}=\varepsilon^{1 / 2} \max \left\{\left|x_{i}\right|, 1 / s_{i}\right\} \operatorname{sign}\left(x_{i}\right), \varepsilon$ is the machine epsilon, $s_{i}$ is the scaling factor of the $i$-th variable, and $e_{i}$ is the $i$-th unit vector. For more details, see Dennis and Schnabel (1983).

Since the finite-difference method has truncation error, cancellation error, and rounding error, users should be aware of possible poor performance. When possible, high precision arithmetic is recommended.

## Comments

This is Description A5.6.4, Dennis and Schnabel, 1983, page 323.

## Example

In this example, the gradient of $f(x)=x_{1}-x_{1} x_{2}-2$ is estimated by the finite-difference method at the point (1.0, 1.0).

```
USE CDGRD_INT
USE UMACH_INT
```

```
IMPLICIT NONE
INTEGER I, N, NOUT
PARAMETER (N=2)
REAL EPSFCN, GC (N), XC (N)
EXTERNAL FCN
! Initialization.
DATA XC/2*1.0E0/
EPSFCN = 0.01
CALL CDGRD (FCN, XC, GC, EPSFCN=EPSFCN)
!
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) (GC(I), I=1,N)
99999 FORMAT (' The gradient is', 2F8.2, /)
!
END
!
SUBROUTINE FCN (N, X, F)
INTEGER N
REAL X(N), F
!
F = X(1) - X(1)*X(2) - 2.0EO
!
RETURN
END
```


## Output

The gradient is $0.00-1.00$

## FDGRD

Approximates the gradient using forward differences.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)
X - The point at which the function is evaluated. (Input) X should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X C}$ - Vector of length N containing the point at which the gradient is to be estimated. (Input)
$\boldsymbol{F C}$ - Scalar containing the value of the function at XC. (Input)
$\boldsymbol{G C}-$ Vector of length N containing the estimated gradient at XC. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XC}, 1)$.
XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
In the absence of other information, set all entries to 1.0.
Default: $\mathrm{XSCALE}=1.0$.
EPSFCN - Estimate of the relative noise in the function. (Input)
EPSFCN must be less than or equal to 0.1 . In the absence of other information, set EPSFCN to 0.0.
Default: $\operatorname{EPSFCN}=0.0$.

## FORTRAN 90 Interface

Generic: CALL FDGRD (FCN, XC, FC, GC [,...])
Specific: The specific interface names are S_FDGRD and D_FDGRD.

## FORTRAN 77 Interface

Single: CALL FDGRD (FCN, XC, FC, GC, N, XSCALE, EPSFCN)
Double: The double precision name is DFDGRD.

## Description

The routine FDGRD uses the following finite-difference formula to estimate the gradient of a function of $n$ variables at $x$ :

$$
\frac{f\left(x+h_{i} e_{i}\right)-f(x)}{h_{i}} \quad \text { for } i=1, \ldots, n
$$

where $h_{i}=\varepsilon^{1 / 2} \max \left\{\left|x_{i}\right|, 1 / s_{i}\right\} \operatorname{sign}\left(x_{i}\right), \varepsilon$ is the machine epsilon, $e_{i}$ is the $i$-th unit vector, and $s_{i}$ is the scaling factor of the $i$-th variable. For more details, see Dennis and Schnabel (1983).
Since the finite-difference method has truncation error, cancellation error, and rounding error, users should be aware of possible poor performance. When possible, high precision arithmetic is recommended. When accuracy of the gradient is important, IMSL routine CDGRD should be used.

## Comments

This is Description A5.6.3, Dennis and Schnabel, 1983, page 322.

## Example

In this example, the gradient of $f(x)=x_{1}-x_{1} x_{2}-2$ is estimated by the finite-difference method at the point (1.0, 1.0).

```
USE FDGRD INT
USE UMACH_INT
IMPLICIT NONE
INTEGER I, N, NOUT
PARAMETER (N=2)
REAL EPSFCN, FC, GC(N), XC(N)
EXTERNAL FCN
! DATA XC/2*1.0E0/
! Set function noise.
EPSFCN = 0.01
CALL FCN (N, XC, FC)
    CALL FDGRD (FCN, XC, FC, GC, EPSFCN=EPSFCN)
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) (GC(I),I=1,N)
99999 FORMAT (' The gradient is', 2F8.2, /)
    END
    SUBROUTINE FCN (N, X, F)
    INTEGER N
    REAL X(N), F
    F = X(1) - X(1)*X(2) - 2.0E0
RETURN
END
```

$!$
!
!
!
!

## Output

The gradient is $0.00-1.00$

## FDHES

Approximates the Hessian using forward differences and function values.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (N, X, F), where

N - Length of X . (Input)

X - The point at which the function is evaluated. (Input)
$x$ should not be changed by FCN.
F - The computed function value at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X C}$ - Vector of length N containing the point at which the Hessian is to be approximated. (Input)
$\boldsymbol{F C}$ - Function value at XC. (Input)
$\boldsymbol{H}-\mathrm{N}$ by N matrix containing the finite difference approximation to the Hessian in the lower triangle. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XC}, 1)$.

XSCALE - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
In the absence of other information, set all entries to 1.0.
Default: XSCALE = 1.0.
EPSFCN - Estimate of the relative noise in the function. (Input)
EPSFCN must be less than or equal to 0.1 . In the absence of other information, set EPSFCN to 0.0.
Default: $\operatorname{EPSFCN}=0.0$.
$\mathbf{L D H}$ - Row dimension of H exactly as specified in the dimension statement of the calling program. (Input)
Default: LDH = SIZE (H,1).

## FORTRAN 90 Interface

Generic: CALL FDHES (FCN, XC, FC, H [,...])

Specific: The specific interface names are S_FDHES and D_FDHES.

## FORTRAN 77 Interface

Single: CALL FDHES (FCN, N, XC, XSCALE, FC, EPSFCN, H, LDH)
Double: $\quad$ The double precision name is DFDHES.

## Description

The routine FDHES uses the following finite-difference formula to estimate the Hessian matrix of function $f$ at $x$ :

$$
\frac{f\left(x+h_{i} e_{i}+h_{j} e_{j}\right)-f\left(x+h_{i} e_{i}\right)-f\left(x+h_{j} e_{j}\right)+f(x)}{h_{i} h_{j}}
$$

where $h_{i}=\varepsilon^{1 / 3} \max \left\{\left|x_{i}\right|, 1 / s_{i}\right\} \operatorname{sign}\left(x_{i}\right), h_{j}=\varepsilon^{1 / 3} \max \left\{\left|x_{j}\right|, 1 / s_{i}\right\} \operatorname{sign}\left(x_{j}\right), \varepsilon$ is the machine epsilon or user-supplied estimate of the relative noise, $s_{i}$ and $s_{j}$ are the scaling factors of the $i$-th and $j$-th variables, and $e_{i}$ and $e_{j}$ are the $i$-th and $j$-th unit vectors, respectively. For more details, see Dennis and Schnabel (1983).

Since the finite-difference method has truncation error, cancellation error, and rounding error, users should be aware of possible poor performance. When possible, high precision arithmetic is recommended.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2HES/DF2HES. The reference is:

CALL F2HES (FCN, N, XC, XSCALE, FC, EPSFCN, H, LDH, WK1, WK2)
The additional arguments are as follows:

WK1 - Real work vector of length n.

WK2 - Real work vector of length N .
2. This is Description A5.6.2 from Dennis and Schnabel, 1983; page 321.

## Example

The Hessian is estimated for the following function at $(1,-1)$

$$
f(x)=x_{1}^{2}-x_{1} x_{2}-2
$$

```
USE FDHES_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N, LDHES, NOUT
```

```
PARAMETER (N=2, LDHES=2)
REAL XC(N), FVALUE, HES(LDHES,N), EPSFCN
EXTERNAL FCN
DATA XC/1.0E0,-1.0E0/
    EPSFCN = 0.001
CALL FCN (N, XC, FVALUE)
Get Hessian forward difference
approximation
CALL FDHES (FCN, XC, FVALUE, HES, EPSFCN=EPSFCN)
CALL UMACH (2, NOUT)
WRITE (NOUT, 99999) ((HES (I, J), J=1,I), I=1,N)
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ l o w e r ~ t r i a n g l e ~ o f ~ t h e ~ H e s s i a n ~ i s ' , ~ / , \& ~
    5X,F10.2,/,5X,2F10.2,/)
END
SUBROUTINE FCN (N, X, F)
INTEGER N
REAL X(N), F
F = X(1)*(X(1) - X(2)) - 2.0E0
RETURN
END
```

!
!

## Output

```
The lower triangle of the Hessian is
    2.00
-1.00 0.00
```


## GDHES

Approximates the Hessian using forward differences and a user-supplied gradient.

## Required Arguments

GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
X - The point at which the gradient is evaluated. (Input) $X$ should not be changed by GRAD.

G - The gradient evaluated at the point X . (Output)

GRAD must be declared EXTERNAL in the calling program.
$\boldsymbol{X C}$ - Vector of length N containing the point at which the Hessian is to be estimated. (Input)
$\boldsymbol{G C}$ - Vector of length N containing the gradient of the function at XC . (Input)
$\boldsymbol{H}-\mathrm{N}$ by N matrix containing the finite-difference approximation to the Hessian in the lower triangular part and diagonal. (Output)

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XC}, 1)$.
$\boldsymbol{X S C A L E}$ - Vector of length N containing the diagonal scaling matrix for the variables. (Input)
In the absence of other information, set all entries to 1.0.
Default: $\mathrm{XSCALE}=1.0$.
EPSFCN - Estimate of the relative noise in the function. (Input)
EPSFCN must be less than or equal to 0.1 . In the absence of other information, set EPSFCN to 0.0.
Default: EPSFCN $=0.0$.
$\boldsymbol{L D H}$ - Leading dimension of H exactly as specified in the dimension statement of the calling program. (Input)
Default: LDH $=\operatorname{SIZE}(H, 1)$.

## FORTRAN 90 Interface

Generic: CALL GDHES (GRAD, XC, GC, H [, ...])
Specific: The specific interface names are S_GDHES and D_GDHES.

## FORTRAN 77 Interface

Single: CALL GDHES (GRAD, N, XC, XSCALE, GC, EPSFCN, H, LDH)
Double: The double precision name is DGDHES.

## Description

The routine GDHES uses the following finite-difference formula to estimate the Hessian matrix of function $F$ at $x$ :

$$
\frac{g\left(x+h_{j} e_{j}\right)-g(x)}{h_{j}}
$$

where $h_{j}=\varepsilon^{1 / 2} \max \left\{\left|x_{j}\right|, 1 / s_{j}\right\} \operatorname{sign}\left(x_{j}\right), \varepsilon$ is the machine epsilon, $s_{j}$ is the scaling factor of the $j$-th variable, $g$ is the analytic gradient of $F$ at $x$, and $e_{j}$ is the $j$-th unit vector. For more details, see Dennis and Schnabel (1983).
Since the finite-difference method has truncation error, cancellation error, and rounding error, users should be aware of possible poor performance. When possible, high precision arithmetic is recommended.

## Comments

1. Workspace may be explicitly provided, if desired, by use of G2HES/DG2HES. The reference is:

CALL G2HES (GRAD, N, XC, XSCALE, GC, EPSFCN, H, LDH, WK)
The additional argument is
WK - Work vector of length N .
2. This is Description A5.6.1, Dennis and Schnabel, 1983; page 320.

## Example

The Hessian is estimated by the finite-difference method at point (1.0, 1.0) from the following gradient functions:

$$
\begin{aligned}
& g_{1}=2 x_{1} x_{2}-2 \\
& g_{2}=x_{1} x_{1}+1
\end{aligned}
$$

```
USE GDHES_INT
USE UMACH_INT
```

IMPLICIT NONE
N, LDHES, NOUT
PARAMETER ( $\mathrm{N}=2$, LDHES=2)
REAL $\quad \mathrm{XC}(\mathrm{N}), \mathrm{GC}(\mathrm{N}), \mathrm{HES}($ LDHES,N)
EXTERNAL GRAD
!
DATA XC/2*1.0E0/
CALL GRAD (N, XC, GC)
CALL GDHES (GRAD, XC, GC, HES)

```
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) ((HES(I,J),J=1,N),I=1,N)
99999 FORMAT (' THE HESSIAN IS', /, 2(5X,2F10.2,/),/)
!
    END
!
    SUBROUTINE GRAD (N, X, G)
    INTEGER N
    REAL X(N),G(N)
!
    G(1) = 2.0E0*X(1)*X(2) - 2.0E0
    G(2) = X(1)*X(1) + 1.0E0
!
    RETURN
    END
```


## Output

THE HESSIAN IS
2.002 .00
$2.00 \quad 0.00$

## FDJAC

Approximates the Jacobian of m functions in N unknowns using forward differences.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (M, N, X, F), where

M - Length of F . (Input)
N - Length of X . (Input)
X - The point at which the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function at the point X . (Output)

FCN must be declared EXTERNAL in the calling program.
$\boldsymbol{X C}$ - Vector of length N containing the point at which the gradient is to be estimated.
(Input)
FC - Vector of length M containing the function values at XC. (Input)
$\boldsymbol{F} \boldsymbol{J} A \boldsymbol{C}-\mathrm{m}$ by N matrix containing the estimated Jacobian at XC. (Output)

## Optional Arguments

$\boldsymbol{M}$ - The number of functions. (Input)
Default: M = SIZE (FC, 1 ).
$N$ - The number of variables. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{XC}, 1)$.
XSCALE - Vector of length N containing the diagonal scaling matrix for the variables.
(Input)
In the absence of other information, set all entries to 1.0 .
Default: $\mathrm{XSCALE}=1.0$.
EPSFCN - Estimate for the relative noise in the function. (Input)
EPSFCN must be less than or equal to 0.1 . In the absence of other information, set EPSFCN to 0.0.
Default: EPSFCN $=0.0$.
LDFJAC - Leading dimension of FJAC exactly as specified in the dimension statement of the calling program. (Input)
Default: LDFJAC $=\operatorname{SIZE}(F J A C, 1)$.

## FORTRAN 90 Interface

Generic: CALL FDJAC (FCN, XC, FC, FJAC [, ...])
Specific: The specific interface names are S_FDJAC and D_FDJAC.

## FORTRAN 77 Interface

Single: CALL FDJAC (FCN, M, N, XC, XSCALE, FC, EPSFCN, FJAC, LDFJAC)

Double: The double precision name is DFDJAC.

## Description

The routine FDJAC uses the following finite-difference formula to estimate the Jacobian matrix of function $f$ at $x$ :

$$
\frac{f\left(x+h_{j} e_{j}\right)-f(x)}{h_{j}}
$$

where $e_{j}$ is the $j$-th unit vector, $h_{j}=\varepsilon^{1 / 2} \max \left\{\left|x_{j}\right|, 1 / s_{j}\right\} \operatorname{sign}\left(x_{j}\right), \varepsilon$ is the machine epsilon, and $s_{j}$ is the scaling factor of the $j$-th variable. For more details, see Dennis and Schnabel (1983).

Since the finite-difference method has truncation error, cancellation error, and rounding error, users should be aware of possible poor performance. When possible, high precision arithmetic is recommended.

## Comments

1. Workspace may be explicitly provided, if desired, by use of F2 JAC/DF2JAC. The reference is:

CALL F2JAC (FCN, M, N, XC, XSCALE, FC, EPSFCN, FJAC, LDFJAC, WK)
The additional argument is:
WK - Work vector of length M.
2. This is Description A5.4.1, Dennis and Schnabel, 1983, page 314.

## Example

In this example, the Jacobian matrix of

$$
\begin{aligned}
& f_{1}(x)=x_{1} x_{2}-2 \\
& f_{2}(x)=x_{1}-x_{1} x_{2}+1
\end{aligned}
$$

is estimated by the finite-difference method at the point (1.0, 1.0).

```
    USE FDJAC INT
    USE UMACH_INT
    IMPLICIT NONE
! Declaration of variables
    INTEGER N, M, LDFJAC, NOUT
    PARAMETER (N=2, M=2, LDFJAC=2)
    REAL FJAC (LDFJAC,N), XC (N), FC (M), EPSFCN
    EXTERNAL FCN
!
    DATA XC/2*1.0E0/
    EPSFCN = 0.01
! Evaluate the function at the
    CALL FCN (M, N, XC, FC)
        Get Jacobian forward-difference
        approximation
    CALL FDJAC (FCN, XC, FC, FJAC, EPSFCN=EPFSCN)
        Print results
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) ((FJAC (I,J), J=1,N), I=1,M)
99999 FORMAT (' The Jacobian is', /, 2(5X,2F10.2,/),/)
!
END
!
    SUBROUTINE FCN (M, N, X, F)
```

    INTEGER M, N
    REAL \(\quad X(N), F(M)\)
    $!$
$F(1)=X(1) * X(2)-2.0 E 0$
$F(2)=X(1)-X(1) * X(2)+1.0 E 0$
!
RETURN
END

## Output

The Jacobian is
$1.00 \quad 1.00$
$0.00-1.00$

## CHGRD

Checks a user-supplied gradient of a function.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function of which the gradient will be checked. The usage is
CALL FCN (N, X, F), where

N - Length of X . (Input)

X - The point at which the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.

GRAD - Vector of length N containing the estimated gradient at X . (Input)
$\boldsymbol{X}$ - Vector of length N containing the point at which the gradient is to be checked. (Input)

INFO - Integer vector of length N. (Output)
$\operatorname{INFO}(I)=0$ means the user-supplied gradient is a poor estimate of the numerical gradient at the point $\mathrm{X}(\mathrm{I})$.
$\operatorname{INFO}(\mathrm{I})=1$ means the user-supplied gradient is a good estimate of the numerical gradient at the point $X(I)$.
$\operatorname{INFO}(\mathrm{I})=2$ means the user-supplied gradient disagrees with the numerical gradient at the point $X(I)$, but it might be impossible to calculate the numerical gradient.
$\operatorname{INFO}(\mathrm{I})=3$ means the user-supplied gradient and the numerical gradient are both zero at $X(I)$, and, therefore, the gradient should be rechecked at a different point.

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CHGRD (FCN, GRAD, X, INFO [, ...])
Specific: The specific interface names are S_CHGRD and D_CHGRD.

## FORTRAN 77 Interface

Single: CALL CHGRD (FCN, GRAD, N, X, INFO)
Double: The double precision name is DCHGRD.

## Description

The routine CHGRD uses the following finite-difference formula to estimate the gradient of a function of $n$ variables at $x$ :

$$
g_{i}(x)=\frac{f\left(x+h_{i} e_{i}\right)-f(x)}{h_{i}} \quad \text { for } i=1, \ldots, n
$$

where $h_{i}=\varepsilon^{1 / 2} \max \left\{\left|x_{i}\right|, 1 / s_{i}\right\} \operatorname{sign}\left(x_{i}\right), \varepsilon$ is the machine epsilon, $e_{i}$ is the $i$-th unit vector, and $s_{i}$ is the scaling factor of the $i$-th variable.

The routine CHGRD checks the user-supplied gradient $\nabla f(x)$ by comparing it with the finitedifference gradient $g(x)$. If

$$
\left|g_{i}(x)-(\nabla f(x))_{i}\right|<\tau\left|(\nabla f(x))_{i}\right|
$$

where $\tau=\varepsilon^{1 / 4}$, then $(\nabla f(x))_{i}$, which is the $i$-th element of $\nabla f(x)$, is declared correct; otherwise, CHGRD computes the bounds of calculation error and approximation error. When both bounds are too small to account for the difference, $(\nabla f(x))_{i}$ is reported as incorrect. In the case of a large error bound, CHGRD uses a nearly optimal stepsize to recompute $g_{i}(x)$ and reports that $(\nabla f(x))_{i}$ is correct if

$$
\left|g_{i}(x)-(\nabla f(x))_{i}\right|<2 \tau\left|(\nabla f(x))_{i}\right|
$$

Otherwise, $(\nabla f(x))_{i}$ is considered incorrect unless the error bound for the optimal step is greater than $\tau\left|(\nabla f(x))_{i}\right|$. In this case, the numeric gradient may be impossible to compute correctly. For more details, see Schnabel (1985).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2GRD/DC2GRD. The reference is:

CALL C2GRD (FCN, GRAD, N, X, INFO, FX, XSCALE, EPSFCN, XNEW)
The additional arguments are as follows:
$\boldsymbol{F X}$ - The functional value at X .

XSCALE - Real vector of length N containing the diagonal scaling matrix.
EPSFCN - The relative "noise" of the function FCN.
$\boldsymbol{X N E W}$ - Real work vector of length N.
2. Informational errors

Type Code
41 The user-supplied gradient is a poor estimate of the numerical gradient.

## Example

The user-supplied gradient of

$$
f(x)=x_{i}+x_{2} e^{-\left(t-x_{3}\right) 2 / x_{4}}
$$

at $(625,1,3.125,0.25)$ is checked where $t=2.125$.
USE CHGRD_INT
USE WRIRN INT

IMPLICIT NONE

INTEGER N
PARAMETER ( $\mathrm{N}=4$ )

INTEGER INFO (N)
REAL GRAD (N), X(N)
EXTERNAL DRIV, FCN
$!$
$!$
$!$
DATA X/625.0E0, 1.0E0, 3.125E0, 0.25E0/
CALL DRIV (N, X, GRAD)
CALL CHGRD (FCN, GRAD, X, INFO)
CALL WRIRN ('The information vector', INFO, 1, N, 1)

```
    END
!
    SUBROUTINE FCN (N, X, FX)
    INTEGER N
    REAL X(N), FX
!
    REAL EXP
    INTRINSIC EXP
!
    FX = X(1) + X(2)*EXP(-1.0E0*(2.125E0-X(3))**2/X(4))
    RETURN
    END
!
    SUBROUTINE DRIV (N, X, GRAD)
    INTEGER N
    REAL X(N), GRAD(N)
    REAL EXP
    INTRINSIC EXP
!
GRAD(1) = 1.0E0
GRAD(2) = EXP(-1.0E0* (2.125E0-X(3))**2/X(4))
GRAD(3) = X(2)*EXP(-1.0E0*(2.125E0-X(3))**2/X(4))*2.0E0/X(4)* &
    (2.125-X(3))
GRAD(4) = X(2)*EXP(-1.0E0* (2.125E0-X(3))**2/X(4))* &
        (2.125E0-X(3))**2/(X(4)*X(4))
    RETURN
    END
```


## Output

The information vector
1324
$\begin{array}{llll}1 & 1 & 1 & 1\end{array}$

## CHHES

Checks a user-supplied Hessian of an analytic function.

## Required Arguments

GRAD - User-supplied subroutine to compute the gradient at the point x . The usage is CALL GRAD ( $\mathrm{N}, \mathrm{X}, \mathrm{G}$ ), where

N - Length of X and G . (Input)
$x-$ The point at which the gradient is evaluated. $x$ should not be changed by GRAD. (Input)

G - The gradient evaluated at the point X . (Output)
GRAD must be declared EXTERNAL in the calling program.

HESS - User-supplied subroutine to compute the Hessian at the point X . The usage is CALL HESS ( $\mathrm{N}, \mathrm{X}, \mathrm{H}, \mathrm{LDH}$ ), where

N - Length of X. (Input)

X - The point at which the Hessian is evaluated. (Input)
X should not be changed by HESS.
H - The Hessian evaluated at the point X . (Output)
LDH - Leading dimension of $H$ exactly as specified in in the dimension statement of the calling program. (Input)

HESS must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length N containing the point at which the Hessian is to be checked. (Input)
INFO - Integer matrix of dimension N by N . (Output)
$\operatorname{INFO}(I, J)=0$ means the Hessian is a poor estimate for function $I$ at the point $X(J)$.
$\operatorname{INFO}(I, J)=1$ means the Hessian is a good estimate for function $I$ at the point $X(J)$.
$\operatorname{INFO}(I, J)=2$ means the Hessian disagrees with the numerical Hessian for function I at the point $\mathrm{X}(\mathrm{J})$, but it might be impossible to calculate the numerical Hessian.
$\operatorname{INFO}(I, J)=3$ means the Hessian for function $I$ at the point $X(J)$ and the numerical Hessian are both zero, and, therefore, the gradient should be rechecked at a different point.

## Optional Arguments

$N$ - Dimension of the problem. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
LDINFO - Leading dimension of INFO exactly as specified in the dimension statement of the calling program. (Input)
Default: LDINFO $=\operatorname{SIZE}($ INFO, 1$)$.

## FORTRAN 90 Interface

Generic: CALL CHHES (GRAD, HESS, X, INFO [,...])
Specific: The specific interface names are S_CHHES and D_CHHES.

## FORTRAN 77 Interface

Single: CALL CHHES (GRAD, HESS, N, X, INFO, LDINFO)
Double: $\quad$ The double precision name is DCHHES.

## Description

The routine CHHES uses the following finite-difference formula to estimate the Hessian of a function of $n$ variables at $x$ :

$$
B_{i j}(x)=\left(g_{i}\left(x+h_{j} e_{j}\right)-g_{i}(x)\right) / h_{j} \quad \text { for } j=1, \ldots, n
$$

where $h_{j}=\varepsilon^{1 / 2} \max \left\{\left|x_{j}\right|, 1 / s_{j}\right\} \operatorname{sign}\left(x_{j}\right), \varepsilon$ is the machine epsilon, $e_{j}$ is the $j$-th unit vector, $s_{j}$ is the scaling factor of the $j$-th variable, and $g_{i}(x)$ is the gradient of the function with respect to the $i$-th variable.

Next, CHHES checks the user-supplied Hessian $H(x)$ by comparing it with the finite difference approximation $B(x)$. If

$$
|B \operatorname{ij}(x)-H \mathrm{ij}(x)|<\tau|H \mathrm{ij}(x)|
$$

where $\tau=\varepsilon^{1 / 4}$, then $H_{i j}(x)$ is declared correct; otherwise, CHHES computes the bounds of calculation error and approximation error. When both bounds are too small to account for the difference, $H_{i j}(x)$ is reported as incorrect. In the case of a large error bound, CHHES uses a nearly optimal stepsize to recompute $B_{i j}(x)$ and reports that $B_{i j}(x)$ is correct if

$$
|B \mathrm{ij}(x)-H \mathrm{ij}(x)|<2 \tau|H \mathrm{ij}(x)|
$$

Otherwise, $H_{i j}(x)$ is considered incorrect unless the error bound for the optimal step is greater than $\tau\left|H_{i j}(x)\right|$. In this case, the numeric approximation may be impossible to compute correctly. For more details, see Schnabel (1985).

## Comments

Workspace may be explicitly provided, if desired, by use of C2HES/DC2HES. The reference is

```
CALL C2HES (GRAD, HESS, N, X, INFO, LDINFO, G, HX, HS, XSCALE, EPSFCN, INFT,
NEWX)
```

The additional arguments are as follows:
$\boldsymbol{G}$ - Vector of length N containing the value of the gradient GRD at X .
$\boldsymbol{H} \boldsymbol{X}$ - Real matrix of dimension N by N containing the Hessian evaluated at X .
HS - Real work vector of length N .
XSCALE - Vector of length N used to store the diagonal scaling matrix for the variables.

EPSFCN - Estimate of the relative noise in the function.

INFT - Vector of length N. For $\mathrm{I}=1$ through N, INFT contains information about the Jacobian.
$\boldsymbol{N E W X}$ - Real work array of length N.

## Example

The user-supplied Hessian of

$$
f(x)=100\left(x_{2}-x_{1}^{2}\right)^{2}+\left(1-x_{1}\right)^{2}
$$

at $(-1.2,1.0)$ is checked, and the error is found.

```
    USE CHHES_INT
```

    IMPLICIT NONE
    INTEGER LDINFO, N
    PARAMETER ( \(\mathrm{N}=2\), LDINFO=N)
    !
INTEGER INFO(LDINFO,N)
REAL $\quad X(N)$
EXTERNAL GRD, HES
Input values for $X$
$\mathrm{X}=(-1.2,1.0)$
DATA $\mathrm{X} /-1.2,1.0 /$
CALL CHHES (GRD, HES, X, INFO)
END
SUBROUTINE GRD (N, X, UG)
INTEGER N
REAL $\quad X(N), U G(N)$
$\mathrm{UG}(1)=-400.0 * X(1) *(X(2)-X(1) * X(1))+2.0 * X(1)-2.0$
$\mathrm{UG}(2)=200.0 * \mathrm{X}(2)-200.0 * \mathrm{X}(1) * \mathrm{X}(1)$
RETURN
END
SUBROUTINE HES (N, X, HX, LDHS)
INTEGER N , LDHS
REAL $X(N), H X(L D H S, N)$
$H X(1,1)=-400.0 * X(2)+1200.0 * X(1) * X(1)+2.0$
$H X(1,2)=-400.0 * X(1)$
$H X(2,1)=-400.0 * X(1)$
A sign change is made to $H X(2,2)$
$H X(2,2)=-200.0$

## Output

```
*** FATAL ERROR 1 from CHHES. The Hessian evaluation with respect to
*** X(2) and X(2) is a poor estimate.
```


## CHJAC

Checks a user-supplied Jacobian of a system of equations with $M$ functions in $N$ unknowns.

## Required Arguments

FCN - User-supplied subroutine to evaluate the function to be minimized. The usage is CALL FCN (M, N, X, F), where
$M$ - Length of $F$. (Input)
N - Length of X . (Input)

X - The point at which the function is evaluated. (Input) $X$ should not be changed by FCN.

F - The computed function value at the point X . (Output)
FCN must be declared EXTERNAL in the calling program.
JAC - User-supplied subroutine to evaluate the Jacobian at a point X . The usage is CALL JAC (M, N, X, FJAC, LDFJAC), where

M - Length of F. (Input)
N - Length of X. (Input)
x - The point at which the function is evaluated. (Input)
$x$ should not be changed by FCN.
FJAC - The computed m by N Jacobian at the point X . (Output)
LDFJAC - Leading dimension of FJAC. (Input)
JAC must be declared EXTERNAL in the calling program.
$\boldsymbol{X}$ - Vector of length N containing the point at which the Jacobian is to be checked. (Input)
INFO - Integer matrix of dimension M by N . (Output)
$\operatorname{INFO}(I, J)=0$ means the user-supplied Jacobian is a poor estimate for function $I$ at the point $X(J)$.
$\operatorname{INFO}(I, J)=1$ means the user-supplied Jacobian is a good estimate for function $I$ at the point $\mathrm{X}(\mathrm{J})$.
$\operatorname{INFO}(I, J)=2$ means the user-supplied Jacobian disagrees with the numerical Jacobian for function I at the point $X(J)$, but it might be impossible to calculate the numerical Jacobian.
$\operatorname{INFO}(I, J)=3$ means the user-supplied Jacobian for function $I$ at the point $X(J)$ and the numerical Jacobian are both zero. Therefore, the gradient should be rechecked at a different point.

## Optional Arguments

$\boldsymbol{M}$ - The number of functions in the system of equations. (Input) Default: M = SIZE (INFO, 1).
$N$ - The number of unknowns in the system of equations. (Input) Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.

LDINFO - Leading dimension of INFO exactly as specified in the dimension statement of the calling program. (Input)
Default: LDINFO $=\operatorname{SIZE}($ INFO, 1$)$.

## FORTRAN 90 Interface

Generic: CALL CHJAC (FCN, JAC, X, INFO [,...])

Specific: The specific interface names are S_CHJAC and D_CHJAC.

## FORTRAN 77 Interface

Single: CALL CHJAC (FCN, JAC, M, N, X, INFO, LDINFO)
Double: The double precision name is DCHJAC.

## Description

The routine CHJAC uses the following finite-difference formula to estimate the gradient of the $i$-th function of $n$ variables at $x$ :

$$
g \mathrm{ij}(x)=(f \mathrm{i}(x+h \mathrm{j} e \mathrm{j})-f \mathrm{i}(x)) / h \mathrm{j} \quad \text { for } j=1, \ldots, n
$$

where $h_{j}=\varepsilon^{1 / 2} \max \left\{\left|x_{j}\right|, 1 / s_{j}\right\} \operatorname{sign}\left(x_{j}\right), \varepsilon$ is the machine epsilon, $e_{j}$ is the $j$-th unit vector, and $s_{j}$ is the scaling factor of the $j$-th variable.

Next, CHJAC checks the user-supplied Jacobian $J(x)$ by comparing it with the finite difference gradient $g_{i}(x)$. If

$$
\left|g \mathrm{ij}(x)-J_{\mathrm{ij}}(x)\right|<\tau\left|J_{\mathrm{ij}}(x)\right|
$$

where $\tau=\varepsilon^{1 / 4}$, then $\mathrm{J}_{i j}(x)$ is declared correct; otherwise, CHJAC computes the bounds of calculation error and approximation error. When both bounds are too small to account for the difference, $J_{i j}(x)$ is reported as incorrect. In the case of a large error bound, CHJAC uses a nearly optimal stepsize to recompute $g_{i j}(x)$ and reports that $J_{i j}(x)$ is correct if

$$
\left|g \mathrm{ij}(x)-J_{\mathrm{ij}}(x)\right|<2 \tau\left|J_{\mathrm{ij}}(x)\right|
$$

Otherwise, $J_{i j}(x)$ is considered incorrect unless the error bound for the optimal step is greater than $\tau\left|J_{i j}(x)\right|$. In this case, the numeric gradient may be impossible to compute correctly. For more details, see Schnabel (1985).

## Comments

1. Workspace may be explicitly provided, if desired, by use of C2 JAC/DC2JAC. The reference is:

CALL C2JAC (FCN, JAC, N, X, INFO, LDINFO, FX, FJAC, GRAD, XSCALE, EPSFCN, INFT, NEWX)

The additional arguments are as follows:
$\boldsymbol{F} \boldsymbol{X}$ - Vector of length M containing the value of each function in FCN at X .
$\boldsymbol{F J A C}$ - Real matrix of dimension M by n containing the Jacobian of FCN evaluated at X.

GRAD - Real work vector of length N used to store the gradient of each function in FCN.

XSCALE - Vector of length N used to store the diagonal scaling matrix for the variables.

EPSFCN - Estimate of the relative noise in the function.

INFT - Vector of length N. For I = 1 through N, INFT contains information about the Jacobian.
$\boldsymbol{N E W X}$ - Real work array of length N.
2. Informational errors

Type Code
41 The user-supplied Jacobian is a poor estimate of the numerical Jacobian.

## Example

The user-supplied Jacobian of

$$
\begin{aligned}
& f_{1}=1-x_{1} \\
& f_{2}=10\left(x_{2}-x_{1}^{2}\right)
\end{aligned}
$$

at $(-1.2,1.0)$ is checked.
USE CHJAC_INT
USE WRIRN_INT
IMPLICIT NONE
INTEGER LDINFO, N
PARAMETER ( $\mathrm{M}=2, \mathrm{~N}=2$, LDINFO $=\mathrm{M}$ )
!
INTEGER INFO(LDINFO,N)
REAL $\quad \mathrm{X}(\mathrm{N})$
EXTERNAL FCN, JAC
$!$

DATA X/-1.2, 1.0/
CALL CHJAC (FCN, JAC, X, INFO)
CALL WRIRN ('The information matrix', INFO)
END
!
SUBROUTINE FCN (M, N, X, F)
INTEGER M, N
REAL $\quad X(N), F(M)$
$F(1)=1.0-X(1)$
$\mathrm{F}(2)=10.0 *(\mathrm{X}(2)-\mathrm{X}(1) * \mathrm{X}(1))$
RETURN
END
!
SUBROUTINE JAC (M, N, X, FJAC, LDFJAC)
INTEGER M, N, LDFJAC
REAL $\quad \mathrm{X}(\mathrm{N})$, FJAC (LDFJAC, N$)$
!
$\operatorname{FJAC}(1,1)=-1.0$
$\operatorname{FJAC}(1,2)=0.0$
$\operatorname{FJAC}(2,1)=-20.0 * X(1)$
$\operatorname{FJAC}(2,2)=10.0$
RETURN
END

## Output

```
*** WARNING ERROR 2 from C2JAC. The numerical value of the Jacobian
*** evaluation for function 1 at the point X(2) = 1.000000E+00 and
```

```
*** the user-supplied value are both zero. The Jacobian for this
*** function should probably be re-checked at another value for
*** this point.
The information matrix
    1 2
1 1 3
2 1 1
```


## GGUES

Generates points in an N-dimensional space.

## Required Arguments

$\boldsymbol{A}-$ Vector of length N. (Input) See B.

B-Real vector of length N. (Input)
$A$ and $B$ define the rectangular region in which the points will be generated, i.e., $A(I)<S(I)<B(I)$ for $I=1,2, \ldots, N$. Note that if $B(I)<A(I)$, then $B(I)<S(I)<A(I)$.
$\boldsymbol{K}$ - The number of points to be generated. (Input)
IDO - Initialization parameter. (Input/Output) IDO must be set to zero for the first call. GGUES resets IDO to 1 and returns the first generated point in $S$. Subsequent calls should be made with $I D O=1$.
$\boldsymbol{S}$ - Vector of length N containing the generated point. (Output) Each call results in the next generated point being stored in S .

## Optional Arguments

$N$ - Dimension of the space. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{B}, 1)$.

## FORTRAN 90 Interface

Generic: CALL GGUES (A, B, K, IDO, S [,...])
Specific: The specific interface names are S_GGUES and D_GGUES.

## FORTRAN 77 Interface

Single: CALL GGUES (N, A, B, K, IDO, S)
Double: The double precision name is DGGUES.

## Description

The routine GGUES generates starting points for algorithms that optimize functions of several variables-or, almost equivalently-algorithms that solve simultaneous nonlinear equations.

The routine GGUES is based on systematic placement of points to optimize the dispersion of the set. For more details, see Aird and Rice (1977).

## Comments

1. Workspace may be explicitly provided, if desired, by use of G2UES / DG2UES. The reference is:

CALL G2UES (N, A, B, K, IDO, S, WK, IWK)
The additional arguments are:

WK - Work vector of length N. WK must be preserved between calls to G2UES.
IWK - Work vector of length 10 . IWK must be preserved between calls to G2UES.
2. Informational error

| Type | Code |  |
| :---: | :---: | :---: |
| 4 | 1 | Attempt to generate more than K points. |

3. The routine GGUES may be used with any nonlinear optimization routine that requires starting points. The rectangle to be searched (defined by A, B, and N) must be determined; and the number of starting points, K , must be chosen. One possible use for GGUES would be to call GGUES to generate a point in the chosen rectangle. Then, call the nonlinear optimization routine using this point as an initial guess for the solution. Repeat this process $k$ times. The number of iterations that the optimization routine is allowed to perform should be quite small ( 5 to 10) during this search process. The best (or best several) point(s) found during the search may be used as an initial guess to allow the optimization routine to determine the optimum more accurately. In this manner, an N dimensional rectangle may be effectively searched for a global optimum of a nonlinear function. The choice of $k$ depends upon the nonlinearity of the function being optimized. A function with many local optima requires a larger value than a function with only a few local optima.

## Example

We want to search the rectangle with vertices at coordinates $(1,1),(3,1),(3,2)$, and $(1,2)$ ten times for a global optimum of a nonlinear function. To do this, we need to generate starting points. The following example illustrates the use of GGUES in this process:

```
USE GGUES INT
USE UMACH INT
IMPLICIT NONE
! Variable Declarations
```

```
    INTEGER N
    PARAMETER (N=2)
!
    INTEGER IDO, J, K, NOUT
    REAL A(N), B(N), S(N)
                                Initializations
                                A = ( 1.0, 1.0)
                                B = (3.0, 2.0)
    DATA A/1.0, 1.0/
    DATA B/3.0, 2.0/
!
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99998)
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ P o i n t ~ N u m b e r ' , ~ 7 X , ~ ' G e n e r a t e d ~ P o i n t ' )
!
    K = 10
    IDO = 0
    DO 10 J=1, K
        CALL GGUES (A, B, K, IDO, S)
!
99999 FORMAT (1X, I7, 14X, '(', F4.1, ',', F6.3, ')')
!
    10 CONTINUE
!
    END
```


## Output

| Point Number | Generated Point |
| :---: | :---: |
| 1 | $(1.5,1.125)$ |
| 2 | $(2.0,1.500)$ |
| 3 | $(2.5,1.750)$ |
| 4 | $(1.5,1.375)$ |
| 5 | $(2.0,1.750)$ |
| 6 | $(1.5,1.625)$ |
| 7 | $(2.5,1.250)$ |
| 8 | $(1.5,1.875)$ |
| 9 | $(2.0,1.250)$ |
| 10 | $(2.5,1.500)$ |

## Chapter 9: Basic Matrix/Vector Operations

## Routines

9.1. Basic Linear Algebra Subprograms (BLAS)
Set a vector to a constant value, $x_{i} \leftarrow a$................................SSET ..... 1425
 ..... 1425
Scale a vector by a constant, $x_{i} \leftarrow a x_{i}$ SSCAL ..... 1425
Set a vector to a constant multiple of a vector, $y_{i} \leftarrow a x_{i} \ldots \ldots$. SVCAL ..... 1425
Add a constant to a vector, $x_{i} \leftarrow x_{i}+a$ ..... SADD ..... 1425
Subtract a vector from a constant, $x_{i} \leftarrow a-x_{i}$ ..... SSUB ..... 1426
Add a multiple of one vector to another, $y_{i} \leftarrow a x_{i}+y_{i}$ .SAXPY ..... 1426
Swap two vectors, $y_{i} \leftrightarrow x_{i}$ ..... SSWAP ..... 1426
Compute $x^{T} y$ or $x^{H} y$ SDOT ..... 1426
Compute extended precision $x^{T} y$ or $x^{H} y$ DSDOT ..... 1427
Compute extended precision $a+x^{T} y$ or $a+x^{H} y$ SDSDOT ..... 1427
Compute ACC $+b+x^{T} y$
with extended precision accumulator SDDOTI ..... 1428
Compute $z_{i} \leftarrow x_{i} y_{i}$ SHPROD ..... 1428
Compute $\Sigma x_{i} y_{i} z_{i}$. .SXYZ ..... 1428
Compute $\Sigma x_{i}$ ..... SSUM ..... 1428
Compute $\Sigma\left|x_{i}\right|$, SASUM ..... 1429
Compute $\|x\|_{2}$ SNRM2 ..... 1429
Compute $\Pi x_{i}$ SPRDCT ..... 1429
Find the index $i$ such that $x_{i}=\min _{j} x_{j}$ ..... ISMIN ..... 1429
Find the index $i$ such that $x_{i}=\max _{j} x_{j}$ ..... ISMAX ..... 1430
Find the index $i$ such that $\left|x_{i}\right|=\min _{j}\left|x_{j}\right|$ ..... ISAMIN ..... 1430
Find the index $i$ such that $\left|x_{i}\right|=\max _{j}\left|x_{j}\right|$ ..... ISAMAX ..... 1430
Construct a Givens rotation SROTG ..... 1430
Apply a Givens rotation. SROT ..... 1431
Construct a modified Givens rotation SROTMG ..... 1432
Apply a modified Givens rotation SROTM ..... 1433
Matrix-vector multiply, general ..... 1436
Matrix-vector multiply, banded ..... 1437
Matrix-vector multiply, Hermitian CHEMV ..... 1437
Matrix-vector multiply, Hermitian and banded CHBMV ..... 1437
Matrix-vector multiply, symmetric and real ..... SSYMV ..... 1437
Matrix-vector multiply, symmetric and banded ..... SSBMV ..... 1438
Matrix-vector multiply, triangular ..... 1438
Matrix-vector multiply, triangular and banded ..... 1438
Matrix-vector solve, triangular ..... 1438
Matrix-vector solve, triangular and banded ..... 1439
Rank-one matrix update, general and real ..... 1439
Rank-one matrix update, general, complex, and transpose. ..... CGERU ..... 1439
Rank-one matrix update, general, complex, and conjugate transpose CGERC ..... 1439
Rank-one matrix update, Hermitian and conjugate transpose CHER ..... 1439
Rank-two matrix update,
Hermitian and conjugate transpose CHER2 ..... 1440
Rank-one matrix update, symmetric and real SSYR ..... 1440
Rank-two matrix update, symmetric and real ..... 1440
Matrix-matrix multiply, general ..... 1440
Matrix-matrix multiply, symmetric ..... 1441
Matrix-matrix multiply, Hermitian ..... 1441
Rank- $k$ update, symmetric ..... 1441
Rank- $k$ update, Hermitian ..... 1441
Rank- $2 k$ update, symmetric ..... 1442
Rank- $2 k$ update, Hermitian. ..... 1442
Matrix-matrix multiply, triangular ..... 1442
Matrix-matrix solve, triangular ..... 1443
9.2. Other Matrix/Vector Operations
9.2.1 Matrix Copy
Real general CRGRG ..... 1444
Complex general CCGCG ..... 1445
Real band ..... CRBRB ..... 1447
Complex band ..... 1448
9.2.2 Matrix Conversion
Real general to real band CRGRB ..... 1450
Real band to real general CRBRG ..... 1452
Complex general to complex band CCGCB ..... 1453
Complex band to complex general. ..... CCBCG ..... 1455
Real general to complex general ..... CRGCG ..... 1457
Real rectangular to complex rectangular ..... 1458
Real band to complex band ..... 1460
Real symmetric to real general ..... 1462
Complex Hermitian to complex general ..... 1463
Real symmetric band to real band ..... 1465
Complex Hermitian band to complex band ..... 1467
Real rectangular matrix to its transpose ..... 1469
9.2.3 Matrix Multiplication
Compute $X^{T} X$. ..... MXTXF ..... 1470
Compute $X$ Т $Y$ MXTYF ..... 1472
Compute $X Y^{T}$ ..... MXYTF ..... 1474
Multiply two real rectangular matrices ..... 1476
Multiply two complex rectangular matrices MCRCR ..... 1479
Compute matrix Hadamard product HRRRR ..... 1481
Compute the bilinear form $x^{T} A y$ ..... 1483
Compute the matrix polynomial $p(A)$. ..... POLRG ..... 1485
9.2.4 Matrix-Vector Multiplication
Real rectangular matrix times a real vector MURRV ..... 1487
Real band matrix times a real vector MURBV ..... 1489
Complex rectangular matrix times a complex vector. MUCRV ..... 1491
Complex band matrix times a complex vector MUCBV ..... 1493
9.2.5 Matrix Addition
Real band matrix plus a real band matrix .ARBRB ..... 1495
Complex band matrix plus a complex band matrix .ACBCB ..... 1497
9.2.6 Matrix Norm
$\infty$-norm of a real rectangular matrix ..... 1499
1-norm of a real rectangular matrix ..... 1501
Frobenius norm of a real rectangular matrix ..... 1502
1-norm of a real band matrix ..... 1504
1-norm of a complex band matrix ..... 1505
9.2.7 Distance Between Two Points Euclidean distance ..... 1507
1-norm distance DISL1 ..... 1509
$\infty$-norm distance DISLI ..... 1510
9.2.8 Vector Convolutions Convolution of real vectors VCONR ..... 1512
Convolution of complex vectors VCONC ..... 1514
9.3. Extended Precision Arithmetic
Initialize a real accumulator, ACC $\leftarrow a$ DQINI ..... 1516
Store a real accumulator, $a \leftarrow \mathrm{ACC}$ ..... DQSTO ..... 1516
Add to a real accumulator, $\mathrm{ACC} \leftarrow \mathrm{ACC}+a$ ..... DQADD ..... 1516
Add a product to a real accumulator, $\mathrm{ACC} \leftarrow \mathrm{ACC}+a b$ DQMUL ..... 1516
Initialize a complex accumulator, ACC $\leftarrow a$ ZQINI ..... 1516
Store a complex accumulator, $a \leftarrow$ ACC ..... ZQSTO ..... 1516
Add to a complex accumulator, $\mathrm{ACC} \leftarrow \mathrm{ACC}+a$ ZQADD ..... 1516
Add a product to a complex accumulator,
$\mathrm{ACC} \leftarrow \mathrm{ACC}+a b$ ZQMUL ..... 1516

## Basic Linear Algebra Subprograms

The basic linear algebra subprograms, normally referred to as the BLAS, are routines for low-level operations such as dot products, matrix times vector, and matrix times matrix. Lawson et al. (1979) published the original set of 38 BLAS. The IMSL BLAS collection includes these 38 subprograms plus additional ones that extend their functionality. Since Dongarra et al. (1988 and 1990) published extensions to this set, it is customary to refer to the original 38 as Level 1 BLAS. The Level 1 operations are performed on one or two vectors of data. An extended set of subprograms perform operations involving a matrix and one or two vectors. These are called the Level 2 BLAS (see Specification of the Level 2 BLAS). An additional extended set of operations on matrices is called the Level 3 BLAS (see Specification of the Level 3 BLAS).

Users of the BLAS will often benefit from using versions of the BLAS supplied by hardware vendors, if available. This can provide for more efficient execution of many application programs. The BLAS provided by IMSL are written in FORTRAN. Those supplied by vendors may be written in other languages, such as assembler. The documentation given below for the BLAS is compatible with a vendor's version of the BLAS that conforms to the published specifications.

## Programming Notes for Level 1 BLAS

The Level 1 BLAS do not follow the usual IMSL naming conventions. Instead, the names consist of a prefix of one or more of the letters " $I$," "S," "D," "C" and " $z$;" a root name; and sometimes a suffix. For subprograms involving a mixture of data types, the output type is indicated by the first prefix letter. The suffix denotes a variant algorithm. The prefix denotes the type of the operation according to the following table:

| I | Integer |  |  |
| :--- | :--- | :--- | :--- |
| S | Real | C | Complex |
| D | Double | Z | Double complex |
| SD | Single and Double | CZ | Single and double complex |
| DQ | Double and Quadruple | ZQ | Double and quadruple complex |

Vector arguments have an increment parameter that specifies the storage space or stride between elements. The correspondence between the vectors $x$ and $y$ and the arguments SX and SY, and INCX and INCY is

$$
\begin{aligned}
& x_{i}= \begin{cases}\mathrm{SX}((\mathrm{I}-1) * \mathrm{INCX}+1) & \text { if INCX } \geq 0 \\
\mathrm{SX}((\mathrm{I}-\mathrm{N}) * \mathrm{INCX}+1) & \text { if INCX }<0\end{cases} \\
& y_{i}= \begin{cases}\mathrm{SY}((\mathrm{I}-1) * \mathrm{INCY}+1) & \text { if INCY } \geq 0 \\
\mathrm{SY}((\mathrm{I}-\mathrm{N}) * \mathrm{INCY}+1) & \text { if INCY }<0\end{cases}
\end{aligned}
$$

Function subprograms SXYZ and DXYZ refer to a third vector argument $z$. The storage increment INCZ for $z$ is defined like INCX, INCY. In the Level 1 BLAS, only positive values of INCX are allowed for operations that have a single vector argument. The loops in all of the Level 1 BLAS process the vector arguments in order of increasing $i$. For INCX, INCY, INCZ $<0$, this implies processing in reverse storage order.

The function subprograms in the Level 1 BLAS are all illustrated by means of an assignment statement. For example, see SDOT. Any value of a function subprogram can be used in an expression or as a parameter passed to a subprogram as long as the data types agree.

## Descriptions of the Level 1 BLAS Subprograms

The set of Level 1 BLAS are summarized in Table 9.1. This table also lists the page numbers where the subprograms are described in more detail.

## Specification of the Level 1 BLAS

With the definitions,

$$
\begin{aligned}
& \mathrm{MX}=\max \{1,1+(N-1) \mid \text { INCX } \mid\} \\
& \mathrm{MY}=\max \{1,1+(N-1) \mid \text { INCY } \mid\} \\
& \mathrm{MZ}=\max \{1,1+(N-1) \mid \text { INCZ } \mid\}
\end{aligned}
$$

the subprogram descriptions assume the following FORTRAN declarations:

```
IMPLICIT INTEGER (I-N)
IMPLICIT REAL S
IMPLICIT DOUBLE PRECISION D
IMPLICIT COMPLEX
IMPLICIT DOUBLE COMPLEX Z
INTEGER IX(MX)
REAL SX(MX), SY(MY), SZ(MZ),
DOUBLE PRECISION
DOUBLE PRECISION DACC(2), DZACC(4)
COMPLEX CX(MX), CY(MY)
DOUBLE COMPLEX
SPARAM(5)
DX(MX), DY(MY), DZ(MZ),
DPARAM (5)
ZX(MX), ZY(MY)
```

Since FORTRAN 77 does not include the type DOUBLE COMPLEX, subprograms with DOUBLE COMPLEX arguments are not available for all systems. Some systems use the declaration COMPLEX * 16 instead of DOUBLE COMPLEX.

In the following descriptions, the original BLAS are marked with an * in the left column.
Table 9.1: Level 1 Basic Linear Algebra Subprograms

| Operation | Integer | Real | Double | Complex | Double <br> Complex | Pg. |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $x_{i} \leftarrow a$ | ISET | SSET | DSET | CSET | ZSET | 1425 |
| $y_{i} \leftarrow x_{i}$ | ICOPY | SCOPY | DCOPY | CCOPY | ZCOPY | 1425 |
| $x_{i} \leftarrow a x_{i}$ <br> $a \in \mathbf{R}$ | SSCAL | DSCAL | CSCAL <br> CSSCAL | ZSCAL <br> ZDSCAL | 1425 |  |
| $y_{i} \leftarrow a x_{i}$ <br> $a \in \mathbf{R}$ | SVCAL | DVCAL | CVCAL <br> CSVCAL | ZVCAL <br> ZDVCAL | 1425 |  |


| Operation | Integer | Real | Double | Complex | Double Complex | Pg. |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $x_{i} \leftarrow x_{i}+a$ | IADD | SADD | DADD | CADD | ZADD | 1425 |
| $x_{i} \leftarrow a-x_{i}$ | ISUB | SSUB | DSUB | CSUB | ZSUB | 1426 |
| $y_{i} \leftarrow a x_{i}+y_{i}$ |  | SAXPY | DAXPY | CAXPY | ZAXPY | 1426 |
| $y_{i} \leftrightarrow x_{i}$ | ISWAP | SSWAP | DSWAP | CSWAP | ZSWAP | 1426 |
| $\left\lvert\, \begin{aligned} & x \cdot y \\ & \bar{x} \cdot \mathrm{y} \end{aligned}\right.$ |  | SDOT | DDOT | $\begin{aligned} & \text { CDOTU } \\ & \text { CDOTC } \end{aligned}$ | $\begin{aligned} & \text { ZDOTU } \\ & \text { ZDOTC } \end{aligned}$ | 1426 |
| $\begin{aligned} & x \cdot y \dagger \\ & \bar{x} \cdot y \dagger \end{aligned}$ |  | DSDOT |  | $\begin{aligned} & \text { CZDOTU } \\ & \text { CZDOTC } \end{aligned}$ | ZQDOTU ZQDOTC | 1427 |
| $\begin{aligned} & a+x \cdot y \dagger \\ & a+\bar{x} \cdot y \dagger \end{aligned}$ |  | SDSDOT | DQDDOT | $\begin{aligned} & \text { CZUDOT } \\ & \text { CZCDOT } \end{aligned}$ | $\begin{aligned} & \text { ZQUDOT } \\ & \text { ZQCDOT } \end{aligned}$ | 1427 |
| $\begin{aligned} & b+x \cdot y \dagger \\ & \text { ACC }+b+x \cdot y \dagger \end{aligned}$ |  | $\begin{aligned} & \text { SDDOTI } \\ & \text { SDDOTA } \end{aligned}$ | DQDOTI DQDOTA | $\begin{aligned} & \text { CZDOTI } \\ & \text { CZDOTA } \end{aligned}$ | ZQDOTI <br> ZQDOTA | 1428 |
| $z_{i} \leftarrow x_{i} y_{i}$ |  | SHPROD | DHPROD |  |  | 1428 |
| $\sum x_{i} y_{i} z_{i}$ |  | SXYZ | DXYZ |  |  | 1428 |
| $\sum x_{i}$ | ISUM | SSUM | DSUM |  |  | 1428 |
| $\sum\left\|x_{i}\right\|$ |  | SASUM | DASUM | SCASUM | DZASUM | 1429 |
| $\\|x\\|_{2}$ |  | SNRM2 | DNRM2 | SCNRM2 | DZNRM2 | 1429 |
| $\Pi x_{i}$ |  | SPRDCT | DPRDCT |  |  | 1429 |
| $i: x_{i}=\min _{j} x_{j}$ | IIMIN | ISMIN | IDMIN |  |  | 1429 |
| $i: x_{i}=\max _{j} x_{j}$ | IIMAX | I SMAX | I DMAX |  |  | 1430 |
| $i:\left\|x_{i}\right\|=\min _{j}\left\|x_{j}\right\|$ |  | ISAMIN | IDAMIN | ICAMIN | IZAMIN | 1430 |
| $i:\left\|x_{i}\right\|=\max _{j}\left\|x_{j}\right\|$ |  | ISAMAX | IDAMAX | ICAMAX | IZAMAX | 1430 |
| Construct Givens rotation |  | SROTG | DROTG |  |  | 1430 |
| Apply Givens rotation |  | SROT | DROT | CSROT | ZDROT | 1431 |
| Construct modified Givens transform |  | SROTMG | DROTMG |  |  | 1432 |
| Apply modified Givens transform |  | SROTM | DROTM | CSROTM | ZDROTM | 1433 |

$\dagger$ Higher precision accumulation used

## Set a Vector to a Constant Value

| CALL $\operatorname{ISET}$ | (N, IA, IX, INCX) |
| :--- | :--- |
| CALL $\operatorname{SSET}$ | $(N, S A, ~ S X, ~ I N C X)$ |
| CALL DSET | (N, DA, DX, INCX) |
| CALL CSET | (N, CA, CX, INCX) |
| CALL | ZSET |

These subprograms set $x_{i} \leftarrow a$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Copy a Vector

> CALL ICOPY (N, IX, INCX, IY, INCY)
> *CALL SCOPY (N, SX, INCX, SY, INCY)
> *CALL DCOPY (N, DX, INCX, DY, INCY)
> *CALL CCOPY (N, CX, INCX, CY, INCY)
> CALL ZCOPY (N, ZX, INCX, ZY, INCY)

These subprograms set $y_{i} \leftarrow x_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Scale a Vector

$$
\begin{aligned}
& \text { *CALL SSCAL (N, SA, SX, INCX) } \\
& \text { *CALL DSCAL (N, DA, DX, INCX) } \\
& \text { *CALL CSCAL (N, CA, CX, INCX) } \\
& \text { CALL ZSCAL (N, ZA, ZX, INCX) } \\
& \text { *CALL CSSCAL (N, SA, CX, INCX) } \\
& \text { CALL ZDSCAL (N, DA, ZX, INCX) }
\end{aligned}
$$

These subprograms set $x_{i} \leftarrow a x_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately. CAUTION: For CSSCAL and ZDSCAL, the scalar quantity $a$ is real and the vector $x$ is complex.

## Multiply a Vector by a Constant

```
CALL SVCAL (N, SA, SX, INCX, SY, INCY)
CALL DVCAL (N, DA, DX, INCX, DY, INCY)
CALL CVCAL (N, CA, CX, INCX, CY, INCY)
CALL ZVCAL (N, ZA, ZX, INCX, ZY, INCY)
CALL CSVCAL (N, SA, CX, INCX, CY, INCY)
CALL ZDVCAL (N, DA, ZX, INCX, ZY, INCY)
```

These subprograms set $y_{i} \leftarrow a x_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately. CAUTION: For CSVCAL and ZDVCAL, the scalar quantity $a$ is real and the vector $x$ is complex.

## Add a Constant to a Vector

```
CALL IADD (N, IA, IX, INCX)
CALL SADD (N, SA, SX, INCX)
```

```
CALL DADD (N, DA, DX, INCX)
CALL CADD (N, CA, CX, INCX)
CALL ZADD (N, ZA, ZX, INCX)
```

These subprograms set $x_{i} \leftarrow x_{i}+a$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Subtract a Vector from a Constant

```
CALL ISUB (N,IA,IX,INCX)
CALL SSUB (N, SA, SX, INCX)
CALL DSUB (N, DA, DX, INCX)
CALL CSUB (N, CA, CX, INCX)
CALL ZSUB (N, ZA, ZX,INCX)
```

These subprograms set $x_{i} \leftarrow a-x_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Constant Times a Vector Plus a Vector

```
*CALL SAXPY (N, SA, SX, INCX, SY, INCY)
*CALL DAXPY (N, DA, DX, INCX, DY, INCY)
*CALL CAXPY (N, CA, CX, INCX, CY, INCY)
CALL ZAXPY (N, ZA, ZX, INCX, ZY, INCY)
```

These subprograms set $y_{i} \leftarrow a x_{i}+y_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Swap Two Vectors

$$
\begin{array}{r}
\text { CALL ISWAP }(N, I X, \text { INCX, IY, INCY) } \\
\text { *CALL SSWAP }(N, S X, ~ I N C X, ~ S Y, ~ I N C Y) ~ \\
\text { *CALL DSWAP }(N, D X, ~ I N C X, ~ D Y, ~ I N C Y) ~ \\
\text { *CALL CSWAP }(N, C X, ~ I N C X, ~ C Y, ~ I N C Y) ~ \\
\text { CALL ZSWAP }(N, Z X, ~ I N C X, ~ Z Y, ~ I N C Y) ~
\end{array}
$$

These subprograms perform the exchange $y_{i} \leftrightarrow x_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Dot Product

$$
\begin{aligned}
* \mathrm{SW} & =\operatorname{SDOT} \quad(\mathrm{N}, \mathrm{SX}, \text { INCX, SY, INCY) } \\
* \mathrm{DW} & =\operatorname{DDOT}(\mathrm{N}, \mathrm{DX}, \text { INCX, DY, INCY) } \\
*_{\mathrm{CW}} & =\mathrm{CDOTU}(\mathrm{~N}, \mathrm{CX}, \text { INCX, CY, INCY) } \\
{ }^{\mathrm{C}} \mathrm{CW} & =\mathrm{CDOTC}(\mathrm{~N}, \mathrm{CX}, \text { INCX, CY, INCY) } \\
\text { ZW } & =\text { ZDOTU (N, ZX, INCX, ZY, INCY) } \\
\text { ZW } & =\text { ZDOTC }(\mathrm{N}, \mathrm{ZX}, \text { INCX, ZY, INCY) }
\end{aligned}
$$

The function subprograms SDOT, DDOT, CDOTU, and ZDOTU compute

$$
\sum_{i=1}^{N} x_{i} y_{i}
$$

The function subprograms CDOTC and ZDOTC compute

$$
\sum_{i=1}^{N} \bar{x}_{i} y_{i}
$$

The suffix C indicates that the complex conjugates of $x_{i}$ are used. The suffix U indicates that the unconjugated values of $x_{i}$ are used. If $N \leq 0$, then the subprograms return zero.

## Dot Product with Higher Precision Accumulation

| * DW | DSDOT | (N, SX, INCX, SY, INCY) |
| :---: | :---: | :---: |
| CW | CZDOTC | ( $\mathrm{N}, \mathrm{CX}, \mathrm{INCX}, \mathrm{CY}, \mathrm{INCY}$ ) |
| CW | CZDOTU | ( $\mathrm{N}, \mathrm{CX}, \mathrm{INCX}, \mathrm{CY}, \mathrm{INCY}$ ) |
| ZW | ZQDOTC | ( N, ZX, INCX, ZY, INCY) |
| ZW | ZQDOTU | ( N, ZX, INCX, ZY, INCY) |

The function subprogram DSDOT computes

$$
\sum_{i=1}^{N} x_{i} y_{i}
$$

using double precision accumulation. The function subprograms CZDOTU and ZQDOTU compute

$$
\sum_{i=1}^{N} x_{i} y_{i}
$$

using double and quadruple complex accumulation, respectively. The function subprograms CZDOTC and ZQDOTC compute

$$
\sum_{i=1}^{N} \bar{x}_{i} y_{i}
$$

using double and quadruple complex accumulation, respectively. If $N \leq 0$, then the subprograms return zero.

## Constant Plus Dot Product with Higher Precision Accumulation

```
*SW = SDSDOT (N, SA, SX, INCX, SY, INCY)
    DW = DQDDOT (N, DA, DX, INCX, DY, INCY)
    CW = CZCDOT (N, CA, CX, INCX, CY, INCY)
    CW = CZUDOT (N, CA, CX, INCX, CY, INCY)
    ZW = ZQCDOT (N, ZA, ZX, INCX, ZY, INCY)
    ZW = ZQUDOT (N, ZA, ZX, INCX, ZY, INCY)
```

The function subprograms SDSDOT, $\operatorname{DQDDOT,~CZUDOT,~and~ZQUDOT~compute~}$

$$
a+\sum_{i=1}^{N} x_{i} y_{i}
$$

using higher precision accumulation where SDSDOT uses double precision accumulation, DQDDOT uses quadruple precision accumulation, CZUDOT uses double complex accumulation, and ZQUDOT uses quadruple complex accumulation. The function subprograms CZCDOT and ZQCDOT compute

$$
a+\sum_{i=1}^{N} \bar{x}_{i} y_{i}
$$

using double complex and quadruple complex accumulation, respectively. If $N \leq 0$, then the subprograms return zero.

## Dot Product Using the Accumulator

```
SW = SDDOTI (N, SB, DACC, SX, INCX, SY, INCY)
SW = SDDOTA (N, SB, DACC,SX, INCX, SY, INCY)
CW = CZDOTI (N, CB, DACC, CX, INCX, CY, INCY)
CW = CZDOTA (N, CB, DACC, CX, INCX, CY, INCY)
*DW = DQDOTI (N, DB, DACC, DX, INCX, DY, INCY)
*DW = DQDOTA (N,DB, DACC, DX, INCX, DY, INCY)
ZW = ZQDOTI (N, ZB, DZACC, ZX, INCX, ZY, INCY)
ZW = ZQDOTA (N, ZB, DZACC, ZX, INCX, ZY, INCY)
```

The variable DACC, a double precision array of length two, is used as a quadruple precision accumulator. DZACC, a double precision array of length four, is its complex analog. The function subprograms, with a name ending in I, initialize DACC to zero. All of the function subprograms then compute

$$
\mathrm{DACC}+b+\sum_{i=1}^{N} x_{i} y_{i}
$$

and store the result in DACC. The result, converted to the precision of the function, is also returned as the function value. If $N \leq 0$, then the function subprograms return zero.

## Hadamard Product

```
CALL SHPROD (N, SX, INCX, SY, INCY, SZ, INCZ)
CALL DHPROD (N, DX, INCX, DY, INCY, DZ, INCZ)
```

These subprograms set $z_{i} \leftarrow x_{i} y_{i}$ for $i=1,2, \ldots, N$. If $N \leq 0$, then the subprograms return immediately.

## Triple Inner Product

$$
\begin{aligned}
& S W=S X Y Z(N, S X, I N C X, S Y, I N C Y, S Z, I N C Z) \\
& D W=D X Y Z(N, D X, I N C X, D Y, ~ I N C Y, D Z, I N C Z)
\end{aligned}
$$

These function subprograms compute

$$
\sum_{i=1}^{N} x_{i} y_{i} z_{i}
$$

If $N \leq 0$ then the subprograms return zero.

## Sum of the Elements of a Vector

$$
\begin{aligned}
& \text { IW }=\text { ISUM } \quad(\mathrm{N}, \text { IX, INCX }) \\
& \text { SW }=\text { SSUM }(\mathrm{N}, \text { SX, INCX) } \\
& \text { DW }=\mathrm{DSUM} \quad(\mathrm{~N}, \mathrm{DX}, \text { INCX) }
\end{aligned}
$$

These function subprograms compute

$$
\sum_{i=1}^{N} x_{i}
$$

If $N \leq 0$, then the subprograms return zero.

## Sum of the Absolute Values of the Elements of a Vector

$$
\begin{aligned}
* S W & =\text { SASUM }(N, S X, \text { INCX }) \\
* D W & =\text { DASUM }(N, D X, \text { INCX }) \\
* S W & =\text { SCASUM }(N, C X, ~ I N C X) \\
\text { DW } & =\text { DZASUM }(N, Z X, ~ I N C X)
\end{aligned}
$$

The function subprograms SASUM and DASUM compute

$$
\sum_{i=1}^{N}\left|x_{i}\right|
$$

The function subprograms SCASUM and DZASUM compute

$$
\sum_{i=1}^{N}\left[\left|\mathfrak{R} x_{i}\right|+\left|\mathfrak{J} x_{i}\right|\right]
$$

If $N \leq 0$, then the subprograms return zero. CAUTION: For SCASUM and DZASUM, the function subprogram returns a real value.

## Euclidean or $\ell_{2}$ Norm of a Vector

$$
\begin{aligned}
& \text { *SW }=\text { SNRM2 (N, SX, INCX) } \\
& \text { *DW = DNRM2 (N, DX, INCX) } \\
& \text { *SW }=\text { SCNRM2 ( } \mathrm{N}, \mathrm{CX}, \text { INCX) } \\
& \text { DW = DZNRM2 (N, ZX, INCX) }
\end{aligned}
$$

These function subprograms compute

$$
\left[\sum_{i=1}^{N}\left|x_{i}\right|^{2}\right]^{1 / 2}
$$

If $N \leq 0$, then the subprograms return zero. CAUTION: For SCNRM2 and DZNRM2, the function subprogram returns a real value.

## Product of the Elements of a Vector

$$
\begin{aligned}
& S W=\operatorname{SPRDCT} \quad(N, S X, \text { INCX }) \\
& D W=\operatorname{DPRDCT} \quad(N, D X, \text { INCX })
\end{aligned}
$$

These function subprograms compute

$$
\prod_{n=1}^{v} x_{i}
$$

If $N \leq 0$, then the subprograms return zero.
Index of Element Having Minimum Value

$$
\begin{aligned}
\text { IW } & =\text { IIMIN } \quad(N, I X, \text { INCX }) \\
\text { IW } & =\text { ISMIN } \quad(N, S X, ~ I N C X) \\
\text { IW } & =\text { IDMIN } \quad(N, D X, ~ I N C X)
\end{aligned}
$$

These function subprograms compute the smallest index $i$ such that $x_{i}=\min _{1 \leq j \leq N} x_{j}$. If $N \leq 0$, then the subprograms return zero.

## Index of Element Having Maximum Value

$$
\begin{aligned}
& \text { IW }=\text { IIMAX } \quad(\mathrm{N}, \text { IX, INCX) } \\
& \text { IW }=\text { ISMAX }(N, S X, \text { INCX }) \\
& \text { IW }=\text { IDMAX }(N, D X, \text { INCX })
\end{aligned}
$$

These function subprograms compute the smallest index $i$ such that $x_{i}=\max _{1 \leq j \leq N} x_{j}$. If $N \leq 0$, then the subprograms return zero.

## Index of Element Having Minimum Absolute Value

$$
\begin{aligned}
& \text { IW }=\text { ISAMIN } \quad(\mathrm{N}, \mathrm{SX}, \text { INCX) } \\
& \text { IW }=\text { IDAMIN } \quad(\mathrm{N}, \text { DX, INCX) } \\
& \text { IW }=\text { ICAMIN } \quad(\mathrm{N}, \mathrm{CX}, \text { INCX) } \\
& \text { IW }=\text { IZAMIN } \quad(\mathrm{N}, \mathrm{ZX}, \text { INCX) }
\end{aligned}
$$

The function subprograms ISAMIN and IDAMIN compute the smallest index $i$ such that $\left|x_{i}\right|=\min _{1 \leq j \leq N}\left|x_{j}\right|$. The function subprograms ICAMIN and IZAMIN compute the smallest index $i$ such that

$$
\left|\mathfrak{R} x_{i}\right|+\left|\mathfrak{I} x_{i}\right|=\min _{1 \leq j \leq N}\left[\left|\mathfrak{R} x_{j}\right|+\left|\mathfrak{I} x_{j}\right|\right]
$$

If $N \leq 0$, then the subprograms return zero.

## Index of Element Having Maximum Absolute Value

$$
\begin{aligned}
& \text { *IW }=\text { ISAMAX ( } \mathrm{N}, \mathrm{SX}, \text { INCX) } \\
& \text { *IW }=\text { IDAMAX ( } \mathrm{N}, \mathrm{DX}, \text { INCX) } \\
& \text { *IW }=\text { ICAMAX ( } \mathrm{N}, \mathrm{CX}, \text { INCX) } \\
& I W=\text { IZAMAX ( } \mathrm{N}, \mathrm{ZX}, \text { INCX) }
\end{aligned}
$$

The function subprograms ISAMAX and IDAMAX compute the smallest index $i$ such that $\left|x_{i}\right|=\max _{1 \leq j \leq N}\left|x_{j}\right|$. The function subprograms ICAMAX and IZAMAX compute the smallest index $i$ such that

$$
\left|\mathfrak{R} x_{i}\right|+\left|\mathfrak{J} x_{i}\right|=\max _{1 \leq j \leq N}\left[\left|\mathfrak{R} x_{j}\right|+\left|\mathfrak{I} x_{j}\right|\right]
$$

If $N \leq 0$, then the subprograms return zero.

## Construct a Givens Plane Rotation

```
*CALL SROTG (SA, SB, SC, SS)
*CALL DROTG (SA, SB, SC, SS)
```

Given the values $a$ and $b$, these subprograms compute

$$
c= \begin{cases}a / r & \text { if } r \neq 0 \\ 1 & \text { if } r=0\end{cases}
$$

and

$$
s= \begin{cases}b / r & \text { if } r \neq 0 \\ 1 & \text { if } r=0\end{cases}
$$

where $r=\sigma\left(a^{2}+b^{2}\right)^{1 / 2}$ and

$$
\sigma= \begin{cases}\operatorname{sign}(a) & \text { if }|a|>|b| \\ \operatorname{sign}(b) & \text { otherwise }\end{cases}
$$

Then, the values $c, s$ and $r$ satisfy the matrix equation

$$
\left[\begin{array}{rr}
c & s \\
-s & c
\end{array}\right]\left[\begin{array}{l}
a \\
b
\end{array}\right]=\left[\begin{array}{l}
r \\
0
\end{array}\right]
$$

The introduction of $\sigma$ is not essential to the computation of the Givens rotation matrix; but its use permits later stable reconstruction of $c$ and $s$ from just one stored number, an idea due to Stewart (1976). For this purpose, the subprogram also computes

$$
z= \begin{cases}s & \text { if }|s|<c \text { or } c=0 \\ 1 / c & \text { if } 0<|c| \leq s\end{cases}
$$

In addition to returning $c$ and $s$, the subprograms return $r$ overwriting $a$, and $z$ overwriting $b$.
Reconstruction of $c$ and $s$ from $z$ can be done as follows:
If $z=1$, then set $c=0$ and $s=1$
If $|z|<1$, then set

$$
c=\sqrt{1-z^{2}} \text { and } s=z
$$

If $|z|>1$, then set

$$
c=1 / z \text { and } s=\sqrt{1-c^{2}}
$$

## Apply a Plane Rotation

$$
\begin{aligned}
& \text { *CALL SROT (N, SX, INCX, SY, INCY, SC, SS ) } \\
& \text { *CALL DROT (N, DX, INCX, DY, INCY, DC, DS ) } \\
& \text { CALL CSROT (N, CX, INCX, CY, INCY, SC, SS) } \\
& \text { CALL ZDROT (N, ZX, INCX, ZY, INCY, DC, DS) }
\end{aligned}
$$

These subprograms compute

$$
\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \leftarrow\left[\begin{array}{rr}
c & s \\
-s & c
\end{array}\right]=\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \text { for } i=1, \ldots, N
$$

If $N \leq 0$, then the subprograms return immediately. CAUTION: For CSROT and zDROT, the scalar quantities $c$ and $s$ are real, and $x$ and $y$ are complex.

## Construct a Modified Givens Transformation

```
*CALL SROTMG (SD1, SD2, SX1, SY1, SPARAM)
*CALL DROTMG (DD1, DD2, DX1, DY1, DPARAM)
```

The input quantities $d_{1}, d_{2}, x_{1}$ and $y_{1}$ define a 2 -vector $\left[w_{1}, z_{1}\right]^{T}$ by the following:

$$
\left[\begin{array}{l}
w_{i} \\
z_{i}
\end{array}\right]=\left[\begin{array}{ll}
\sqrt{d_{1}} & 0 \\
0 & \sqrt{d_{2}}
\end{array}\right]\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right]
$$

The subprograms determine the modified Givens rotation matrix $H$ that transforms $y_{1}$, and thus, $z_{1}$ to zero. They also replace $d_{1}, d_{2}$ and $x_{1}$ with

$$
\tilde{d}_{1}, \tilde{d}_{2} \text { and } \tilde{x}_{1}
$$

respectively. That is,

$$
\left[\begin{array}{l}
\tilde{w}_{1} \\
0
\end{array}\right]=\left[\begin{array}{ll}
\sqrt{\tilde{d}_{1}} & 0 \\
0 & \sqrt{\tilde{d}_{2}}
\end{array}\right] H\left[\begin{array}{l}
x_{1} \\
y_{1}
\end{array}\right]=\left[\begin{array}{ll}
\sqrt{\tilde{d}_{1}} & 0 \\
0 & \sqrt{\tilde{d}_{2}}
\end{array}\right]\left[\begin{array}{l}
x_{1} \\
y_{1}
\end{array}\right]=\left[\begin{array}{l}
\tilde{x}_{1} \\
0
\end{array}\right]
$$

A representation of this matrix is stored in the array SPARAM or DPARAM. The form of the matrix $H$ is flagged by PARAM(1).
$\operatorname{PARAM}(1)=1$. In this case,

$$
\left|d_{1} x_{1}^{2}\right| \leq\left|d_{2} y_{1}^{2}\right|
$$

and

$$
H=\left[\begin{array}{cc}
\operatorname{PARAM}(2) & 1 \\
-1 & \operatorname{PARAM}(5)
\end{array}\right]
$$

The elements $\operatorname{PARAM}(3)$ and $\operatorname{PARAM}(4)$ are not changed.
$\operatorname{PARAM}(1)=0$. In this case,

$$
\left|d_{1} x_{1}^{2}\right|>\left|d_{2} y_{1}^{2}\right|
$$

and

$$
H=\left[\begin{array}{cc}
1 & \text { PARAM(4) } \\
\text { PARAM(3) } & 1
\end{array}\right]
$$

The elements $\operatorname{PARAM}(2)$ and $\operatorname{PARAM}(5)$ are not changed. $\operatorname{PARAM}(1)=-1$. In this case, rescaling was done and

$$
H=\left[\begin{array}{l}
\operatorname{PARAM}(2) \operatorname{PARAM}(4) \\
\operatorname{PARAM}(3) \operatorname{PARAM}(5)
\end{array}\right]
$$

$\operatorname{PARAM}(1)=-2$. In this case, $H=I$ where $I$ is the identity matrix. The elements $\operatorname{PARAM}(2)$, $\operatorname{PARAM}(3), \operatorname{PARAM}(4)$ and $\operatorname{PARAM}(5)$ are not changed.

The values of $d_{1}, d_{2}$ and $x_{1}$ are changed to represent the effect of the transformation. The quantity $y_{1}$, which would be zeroed by the transformation, is left unchanged.

The input value of $d_{1}$ should be nonnegative, but $d_{2}$ can be negative for the purpose of removing data from a least-squares problem.
See Lawson et al. (1979) for further details.

## Apply a Modified Givens Transformation

```
*CALL SROTM (N, SX, INCX, SY, INCY, SPARAM)
*CALL DROTM (N, DX, INCX, DY, INCY, DPARAM)
CALL CSROTM (N, CX, INCX, CY, INCY, SPARAM)
CALL ZDROTM (N, ZX, INCX, ZY, INCY, DPARAM)
```

If $\operatorname{PARAM}(1)=1.0$, then these subprograms compute

$$
\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \leftarrow\left[\begin{array}{cc}
\operatorname{PARAM}(2) & 1 \\
-1 & \operatorname{PARAM}(5)
\end{array}\right]\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \text { for } i=1, \ldots, N
$$

If $\operatorname{PARAM}(1)=0.0$, then the subprograms compute

$$
\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \leftarrow\left[\begin{array}{cc}
1 & \operatorname{PARAM}(4) \\
\operatorname{PARAM}(3) & 1
\end{array}\right]\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \text { for } i=1, \ldots, N
$$

If $\operatorname{PARAM}(1)=-1.0$, then the subprograms compute

$$
\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \leftarrow\left[\begin{array}{l}
\operatorname{PARAM}(2) \operatorname{PARAM}(4) \\
\operatorname{PARAM}(3) \operatorname{PARAM}(5)
\end{array}\right]\left[\begin{array}{l}
x_{i} \\
y_{i}
\end{array}\right] \text { for } i=1, \ldots, N
$$

If $N \leq 0$ or if $\operatorname{PARAM}(1)=-2.0$, then the subprograms return immediately. CAUTION: For CSROTM and ZDROTM, the scalar quantities $\operatorname{PARAM}(*)$ are real and $x$ and $y$ are complex.

## Programming Notes for Level 2 and Level 3 BLAS

For definitions of the matrix data structures used in the discussion below, see Reference Material. The Level 2 and Level 3 BLAS, like the Level 1 BLAS, do not follow the IMSL naming conventions. Instead, the names consist of a prefix of one of the letters "S," "D," "C" or "z." Next is a root name denoting the kind of matrix. This is followed by a suffix indicating the type of the operation. ${ }^{1}$ The prefix denotes the type of operation according to the following table:

| S | Real | C | Complex |  |
| :--- | :--- | :--- | :--- | :--- |
| D | Double | Z | Double | Complex |

The root names for the kind of matrix:

| GE | General | GB | General Band |
| :--- | :--- | :--- | :--- |
| SY | Symmetric | SB | Symmetric Band |
| HE | Hermitian | HB | Hermitian Band |

TR Triangular TB Triangular Band
The suffixes for the type of operation:

| MV | Matrix-Vector Product | SV | Solve for vector |
| :--- | :--- | :--- | :--- |
| R | Rank-One Update |  |  |
| RU | Rank-One Update, <br> Unconjugated | RC | Rank-One Update, <br> Conjugated |
| R2 | Rank-Two Update |  |  |
| MM | Matrix-Multiply | SM | Symmetric Matrix Multiply |
| RK | Rank-K Update | R2K | Rank 2K Update |

${ }^{1}$ IMSL does not support the Packed Symmetric, Packed-Hermitian, or Packed-Triangular data structures, with respective root names SP, HP or TP, nor any extended precision versions of the Level 2 BLAS.

The specifications of the operations are provided by subprogram arguments of CHARACTER* 1 data type. Both lower and upper case of the letter have the same meaning:

| TRANS, TRANSA, TRANSB | ' N ' | No Transpose |
| :---: | :---: | :---: |
|  | 'T' | Transpose |
|  | 'C' | Conjugage and Transpose |
| UPLO | 'L' | Lower Triangular |
|  | 'U' | Upper Triangular |
| DIAGNL | 'N' | Non-unit Triangular |
|  | 'U' | Unit Triangular |
| SIDE | 'L' | Multiply " $A$ " Matrix on Left side or |
|  | 'R' | Right side of the " $B$ " matrix |

Note: See the "Triangular Mode" section in the Reference Material for definitions of these terms.

## Descriptions of the Level 2 and Level 3 BLAS

The subprograms for Level 2 and Level 3 BLAS that perform operations involving the expression $\beta y$ or $\beta C$ do not require that the contents of $y$ or $C$ be defined when $\beta=0$. In that case, the expression $\beta y$ or $\beta C$ is defined to be zero. Note that for the _GEMV and _GBMV subprograms, the dimensions of the vectors $x$ and $y$ are implied by the specification of the operation. If TRANS $={ }^{\prime} \mathrm{N}^{\prime}$, the dimension of $y$ is $m$; if TRANS $={ }^{\prime} \mathrm{T}^{\prime}$ or $={ }^{\prime} \mathrm{C}^{\prime}$, the dimension of $y$ is $n$. The Level 2 and Level 3 BLAS are summarized in Table 9.2. This table also lists the page numbers where the subprograms are described in more detail.

## Specification of the Level 2 BLAS

Type and dimension for variables occurring in the subprogram specifications are

```
INTEGER
INCX, INCY, NCODA, NLCA, NUCA, LDA, M, N
CHARACTER*1 DIAGNL, TRANS, UPLO
```

| REAL | SALPHA, SBETA, SX (*), SY (*), SA (LDA, *) |
| :--- | :--- |
| DOUBLE PRECISION | DALPHA, DBETA, DX (*), DY (*), DA (LDA, *) |
| COMPLEX | CALPHA, CBETA, CX (*), CY (*), CA (LDA, *) |
| DOUBLE COMPLEX | ZALPHA, ZBETA, $\mathrm{ZX}(*), \mathrm{ZY}(*), \mathrm{ZA}(L D A, *)$ |

There is a lower bound on the leading dimension LDA. It must be $\geq$ the number of rows in the matrix that is contained in this array. Vector arguments have an increment parameter that specifies the storage space or stride between elements. The correspondence between the vector $x, y$ and the arguments SX, SY and INCX, INCY is

$$
\begin{aligned}
& x_{i}= \begin{cases}\mathrm{SX}((\mathrm{I}-1) * \mathrm{INCX}+1) & \text { if INCX }>0 \\
\mathrm{SX}((\mathrm{I}-\mathrm{N}) * \mathrm{INCX}+1) & \text { if INCX }<0\end{cases} \\
& y_{i}= \begin{cases}\mathrm{SY}((\mathrm{I}-1) * \mathrm{INCY}+1) & \text { if INCY }>0 \\
\mathrm{SY}((\mathrm{I}-\mathrm{N}) * \mathrm{INCY}+1) & \text { if INCY }<0\end{cases}
\end{aligned}
$$

In the Level 2 BLAS, only nonzero values of INCX, INCY are allowed for operations that have vector arguments. The Level 3 BLAS do not refer to INCX, INCY.

## Specification of the Level 3 BLAS

Type and dimension for variables occurring in the subprogram specifications are

```
INTEGER
CHARACTER*1
REAL
```

DOUBLE PRECISION
COMPLEX

DOUBLE COMPLEX

```
K, LDA, LDB, LDC, M, N
DIAGNL, TRANS, TRANSA, TRANSB, SIDE, UPLO
SALPHA, SBETA, SA(LDA,*), SB(LDB,*),
SC(LDC,*)
DALPHA, DBETA, DA(LDA,*), DB(LDB,*),
DC(LDC,*)
CALPHA, CBETA, CA(LDA,*), CB(LDB,*),
CC (LDC, *)
ZALPHA, ZBETA, ZA(LDA,*), ZB(LDB,*),
ZC(LDC,*)
```

Each of the integers $K, M, N$ must be $\geq 0$. It is an error if any of them are $<0$. If any of them are $=0$, the subprograms return immediately. There are lower bounds on the leading dimensions LDA, LDB, LDC. Each must be $\geq$ the number of rows in the matrix that is contained in this array.

Table 9.2: Level 2 and 3 Basic Linear Algebra Subprograms

| Operation | Real | Double | Complex | Double <br> Complex | Pg. |
| :--- | :--- | :--- | :--- | :--- | :--- |
| Matrix-Vector Multiply, General | SGEMV | DGEMV | CGEMV | ZGEMV | 1436 |
| Matrix-Vector Multiply, Banded | SGBMV | DGBMV | CGBMV | ZGBMV | 1437 |
| Matrix-Vector Multiply, Hermitian |  |  | CHEMV | ZHEMV | 1437 |
| Matrix-Vector Multiply, <br> Hermitian and Banded |  |  | CHBMV | ZHBMV | 1437 |
| Matrix-Vector Multiply <br> Symmetric and Real | SSYMV | DSYMV |  |  | 1437 |
| Matrix-Vector Multiply, <br> Symmetric and Banded | SSBMV | DSBMV |  |  | 1438 |


| Operation | Real | Double | Complex | Double <br> Complex | Pg. |
| :--- | :--- | :--- | :--- | :--- | :--- |
| Matrix-Vector Multiply, Triangular | STRMV | DTRMV | CTRMV | ZTRMV | 1438 |
| Matrix-Vector Multiply, <br> Triangular and Banded | STBMV | DTBMV | CTBMV | ZTBMV | 1438 |
| Matrix-Vector Solve, Triangular | STRSV | DTRSV | CTRSV | ZTRSV | 1438 |
| Matrix-Vector Solve, <br> Triangular and Banded | STBSV | DTBSV | CTBSV | ZTBSV | 1439 |
| Rank-One Matrix Update, <br> General and Real | SGER | DGER |  |  | 1439 |
| Rank-One Matrix Update, <br> General, Complex and Transpose |  |  | CGERU | ZGERU | 1439 |
| Rank-One Matrix Update, <br> General, Complex, and Conjugate <br> Transpose |  |  | CGERC | ZGERC | 1439 |
| Rank-One Matrix Update, <br> Hermitian and Conjugate Transpose |  |  | CHER | ZHER | 1439 |
| Rank-Two Matrix Update, <br> Hermitian and Conjugate Transpose |  |  | CHER2 | ZHER2 | 1440 |
| Rank-One Matrix Update, <br> Symmetric and Real | SSYR | DSYR |  | 1440 |  |
| Rank-Two Matrix Update, <br> Symmetric and Real | SSYR2 | DSYR2 |  | 1440 |  |
| Matrix--Matrix Multiply, General | SGEMM | DGEMM | CGEMM | ZGEMM | 1440 |
| Matrix-Matrix Multiply, Symmetric | SSYMM | DSYMM | CSYMM | ZSYMM | 1441 |
| Matrix-Matrix Multiply, Hermitian |  |  | CHEMM | ZHEMM | 1441 |
| Rank - $k$ Update, Symmetric | SSYRK | DSYRK | CSYRK | ZSYRK | 1441 |
| Rank - $k$ Update, Hermitian |  |  | CHERK | ZHERK | 1441 |
| Rank -2 $k$ Update, Symmetric | SSYR2K | DSYR2K | CSYR2K | ZSYR2K | 1442 |
| Rank -2 $k$ Update, Hermitian |  |  | CHER2K | ZHER2K | 1442 |
| Matrix-Matrix Multiply, Triangular | STRMM | DTRMM | CTRMM | ZTRMM | 1442 |
| Matrix-Matrix solve, Triangular | STRSM | DTRSM | CTRSM | ZTRSM | 1443 |

## Matrix-Vector Multiply, General

CALL SGEMV (TRANS, M, N, SALPHA, SA, LDA, SX, INCX, SBETA, SY, INCY)
CALL DGEMV (TRANS, M, N, DALPHA, DA, LDA, DX, INCX, DBETA, DY, INCY)
CALL CGEMV (TRANS, M, N, CALPHA, CA, LDA, CX, INCX, CBETA, CY, INCY)
CALL ZGEMV (TRANS, M, N, ZALPHA, ZA, LDA, ZX, INCX, ZBETA, ZY, INCY)
For all data types, $A$ is an $M \times N$ matrix. These subprograms set $y$ to one of the expressions: $y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y$, or for complex data,

$$
y \leftarrow \alpha \bar{A}^{T}+\beta y
$$

The character flag TRANS determines the operation.

## Matrix-Vector Multiply, Banded

```
CALL SGBMV (TRANS,M, N, NLCA, NUCA SALPHA, SA, LDA, SX, INCX, SBETA,
    SY, INCY)
CALL DGBMV (TRANS,M, N, NLCA, NUCA DALPHA, DA, LDA, DX, INCX, DBETA,
    DY, INCY)
CALL CGBMV (TRANS, M, N, NLCA, NUCA CALPHA, CA, LDA, CX, INCX, BETA,
    CY, INCY)
CALL ZGBMV (TRANS,M,N,NLCA, NUCA ZALPHA, ZA, LDA, ZX, INCX, ZBETA,
    ZY, INCY)
```

For all data types, $A$ is an $M \times N$ matrix with NLCA lower codiagonals and NUCA upper codiagonals. The matrix is stored in band storage mode. These subprograms set $y$ to one of the expressions: $y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y$, or for complex data,

$$
y \leftarrow \alpha \bar{A}^{T} x+\beta y
$$

The character flag TRANS determines the operation.

## Matrix-Vector Multiply, Hermitian

> CALL CHEMV (UPLO, N, CALPHA, CA, LDA, CX, INCX, CBETA, CY, INCY)
> CALL ZHEMV (UPLO, N, ZALPHA, ZA, LDA, ZX, INCX, ZBETA, ZY, INCY)

For all data types, $A$ is an $N \times N$ matrix. These subprograms set $y \leftarrow \alpha A x+\beta y$ where $A$ is an Hermitian matrix. The matrix $A$ is either referenced using its upper or lower triangular part. The character flag UPLO determines the part used.

## Matrix-Vector Multiply, Hermitian and Banded

```
CALL CHBMV (UPLO, N, NCODA, CALPHA, CA, LDA, CX, INCX, CBETA, CY, INCY)
CALL ZHBMV (UPLO,N, NCODA, ZALPHA, ZA, LDA, ZX, INCX, ZBETA, ZY, INCY)
```

For all data types, $A$ is an $N \times N$ matrix with NCODA codiagonals. The matrix is stored in band Hermitian storage mode. These subprograms set $y \leftarrow \alpha A x+\beta y$. The matrix $A$ is either referenced using its upper or lower triangular part. The character flag UPLO determines the part used.

## Matrix-Vector Multiply, Symmetric and Real

```
CALL SSYMV (UPLO, N, SALPHA, SA, LDA, SX, INCX, SBETA, SY, INCY)
CALL DSYMV (UPLO, N, DALPHA, DA, LDA, DX, INCX, DBETA, DY, INCY)
```

For all data types, $A$ is an $N \times N$ matrix. These subprograms set $y \leftarrow \alpha A x+\beta y$ where $A$ is a symmetric matrix. The matrix $A$ is either referenced using its upper or lower triangular part. The character flag UPLO determines the part used.

## Matrix-Vector Multiply, Symmetric and Banded

```
CALL SSBMV (UPLO, N, NCODA, SALPHA, SA, LDA, SX, INCX, SBETA, SY, INCY)
CALL DSBMV (UPLO, N, NCODA, DALPHA, DA, LDA, DX, INCX, DBETA, DY, INCY)
```

For all data types, $A$ is an $N \times N$ matrix with NCODA codiagonals. The matrix is stored in band symmetric storage mode. These subprograms set $y \leftarrow \alpha A x+\beta y$. The matrix $A$ is either referenced using its upper or lower triangular part. The character flag UPLO determines the part used.

## Matrix-Vector Multiply, Triangular

```
CALL STRMV (UPLO, TRANS, DIAGNL, N, SA, LDA, SX, INCX)
CALL DTRMV (UPLO,TRANS, DIAGNL, N, DA, LDA, DX, INCX)
CALL CTRMV (UPLO, TRANS, DIAGNL, N, CA, LDA, CX, INCX)
CALL ZTRMV (UPLO, TRANS, DIAGNL, N, ZA, LDA, ZX, INCX)
```

For all data types, $A$ is an $N \times N$ triangular matrix. These subprograms set $x$ to one of the expressions: $x \leftarrow A x, x \leftarrow A^{T} x$, or for complex data,

$$
x \leftarrow \bar{A}^{T} x
$$

The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags UPLO, TRANS, and DIAGNL determine the part of the matrix used and the operation performed.

## Matrix-Vector Multiply, Triangular and Banded

```
CALL STBMV (UPLO, TRANS, DIAGNL, N, NCODA, SA, LDA, SX, INCX)
CALL DTBMV (UPLO, TRANS, DIAGNL, N, NCODA, DA, LDA, DX, INCX)
CALL CTBMV (UPLO, TRANS, DIAGNL, N, NCODA, CA, LDA, CX, INCX)
CALL ZTBMV (UPLO, TRANS, DIAGNL, N, NCODA, ZA, LDA, ZX, INCX)
```

For all data types, $A$ is an $N \times N$ matrix with NCODA codiagonals. The matrix is stored in band triangular storage mode. These subprograms set $x$ to one of the expressions: $x \leftarrow A x, x \leftarrow A^{T} x$, or for complex data,

$$
x \leftarrow \bar{A}^{T} x
$$

The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags UPLO, TRANS, and DIAGNL determine the part of the matrix used and the operation performed.

## Matrix-Vector Solve, Triangular

```
CALL STRSV (UPLO,TRANS, DIAGNL, N, SA, LDA, SX, INCX)
CALL DTRSV (UPLO,TRANS, DIAGNL, N, DA, LDA, DX, INCX)
CALL CTRSV (UPLO, TRANS, DIAGNL, N, CA, LDA, CX, INCX)
CALL ZTRSV (UPLO, TRANS, DIAGNL, N, ZA, LDA, ZX, INCX)
```

For all data types, $A$ is an $N \times N$ triangular matrix. These subprograms solve $x$ for one of the expressions: $x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x$, or for complex data,

$$
x \leftarrow\left(\bar{A}^{T}\right)^{-1} x
$$

The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags UPLO, TRANS, and DIAGNL determine the part of the matrix used and the operation performed.

## Matrix-Vector Solve, Triangular and Banded

```
CALL STBSV (UPLO, TRANS, DIAGNL, N, NCODA, SA, LDA, SX, INCX)
CALL DTBSV (UPLO, TRANS, DIAGNL, N, NCODA, DA, LDA, DX, INCX)
CALL CTBSV (UPLO, TRANS, DIAGNL, N, NCODA, CA, LDA, CX, INCX)
CALL ZTBSV (UPLO, TRANS, DIAGNL, N, NCODA, ZA, LDA, ZX, INCX)
```

For all data types, $A$ is an $N \times N$ triangular matrix with NCODA codiagonals. The matrix is stored in band triangular storage mode. These subprograms solve $x$ for one of the expressions: $x \leftarrow A^{-1} x$, $x \leftarrow\left(A^{T}\right)^{-1} x$, or for complex data,

$$
x \leftarrow\left(\bar{A}^{T}\right)^{-1} x
$$

The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags UPLO, TRANS, and DIAGNL determine the part of the matrix used and the operation performed.

## Rank-One Matrix Update, General and Real

```
CALL SGER (M, N, SALPHA, SX, INCX, SY, INCY, SA, LDA)
CALL DGER (M, N, DALPHA, DX, INCX, DY, INCY, DA, LDA)
```

For all data types, $A$ is an $M \times N$ matrix. These subprograms set $A \leftarrow A+\alpha x y^{T}$.

## Rank-One Matrix Update, General, Complex, and Transpose

```
CALL CGERU (M, N, CALPHA, CX, INCX, CY, INCY, CA, LDA)
CALL ZGERU (M, N, ZALPHA, ZX, INCX, ZY, INCY, ZA, LDA)
```

For all data types, $A$ is an $M \times N$ matrix. These subprograms set $A \leftarrow A+\alpha x y^{T}$.

## Rank-One Matrix Update, General, Complex, and Conjugate Transpose

```
CALL CGERC (M, N, CALPHA, CX, INCX, CY, INCY, CA, LDA)
CALL ZGERC (M, N, ZALPHA, ZX, INCX, ZY, INCY, ZA, LDA)
```

For all data types, $A$ is an $M \times N$ matrix. These subprograms set

$$
A \leftarrow A+\alpha x \bar{y}^{T}
$$

## Rank-One Matrix Update, Hermitian and Conjugate Transpose

```
CALL CHER (UPLO,N, SALPHA, CX, INCX, CA, LDA)
CALL ZHER (UPLO,N, DALPHA, ZX, INCX, ZA, LDA)
```

For all data types, $A$ is an $N \times N$ matrix. These subprograms set

$$
A \leftarrow A+\alpha x \bar{x}^{T}
$$

where $A$ is Hermitian. The matrix $A$ is either referenced by its upper or lower triangular part. The character flag UPLO determines the part used. CAUTION: Notice the scalar parameter $\alpha$ is real, and the data in the matrix and vector are complex.

## Rank-Two Matrix Update, Hermitian and Conjugate Transpose

```
CALL CHER2 (UPLO, N, CALPHA, CX, INCX, CY, INCY, CA, LDA)
CALL ZHER2 (UPLO, N, ZALPHA, ZX, INCX, ZY, INCY, ZA, LDA)
```

For all data types, $A$ is an $N \times N$ matrix. These subprograms set

$$
A \leftarrow A+\alpha x \bar{y}^{T}+\bar{\alpha} y \bar{x}^{T}
$$

where $A$ is an Hermitian matrix. The matrix $A$ is either referenced by its upper or lower triangular part. The character flag UPLO determines the part used.

## Rank-One Matrix Update, Symmetric and Real

```
CALL SSYR (UPLO, N, SALPHA, SX, INCX, SA, LDA)
CALL DSYR (UPLO, N, DALPHA, DX, INCX, DA, LDA)
```

For all data types, $A$ is an $N \times N$ matrix. These subprograms set $A \leftarrow A+\alpha x x^{T}$ where $A$ is a symmetric matrix. The matrix $A$ is either referenced by its upper or lower triangular part. The character flag UPLO determines the part used.

## Rank-Two Matrix Update, Symmetric and Real

```
CALL SSYR2 (UPLO, N, SALPHA, SX, INCX, SY, INCY, SA, LDA)
CALL DSYR2 (UPLO, N, DALPHA, DX, INCX, DY, INCY, DA, LDA)
```

For all data types, $A$ is an $N \times N$ matrix. These subprograms set $A \leftarrow A+\alpha x y^{T}+\alpha y x^{T}$ where $A$ is a symmetric matrix. The matrix $A$ is referenced by its upper or lower triangular part. The character flag UPLO determines the part used.

## Matrix-Matrix Multiply, General

```
CALL SGEMM (TRANSA, TRANSB, M, N, K, SALPHA, SA, LDA, SB, LDB, SBETA, SC,
    LDC)
CALL DGEMM (TRANSA, TRANSB, M, N, K, DALPHA, DA, LDA, DB, LDB, DBETA, DC,
    LDC)
CALL CGEMM (TRANSA, TRANSB, M, N, K, CALPHA, CA, LDA, CB, LDB, CBETA, CC,
    LDC)
CALL ZGEMM (TRANSA, TRANSB, M, N, K, ZALPHA, ZA, LDA, ZB, LDB, ZBETA, ZC,
    LDC)
```

For all data types, these subprograms set $C_{M \times N}$ to one of the expressions:

$$
C \leftarrow \alpha A B+\beta C, C \leftarrow \alpha A^{T} B+\beta C, C \leftarrow \alpha A B^{T}+\beta C, C \leftarrow \alpha A^{T} B^{T}+\beta C
$$

or for complex data, $C \leftarrow \alpha A \bar{B}^{T}+\beta C, C \leftarrow \alpha \bar{A}^{T} B+\beta C, C \leftarrow \alpha A^{T} \bar{B}^{T}+\beta C$,

$$
C \leftarrow \alpha \bar{A}^{T} B^{T}+\beta C, C \leftarrow \alpha \bar{A}^{T} \bar{B}^{T}+\beta C
$$

The character flags TRANSA and TRANSB determine the operation to be performed. Each matrix product has dimensions that follow from the fact that $C$ has dimension $M \times N$.

## Matrix-Matrix Multiply, Symmetric

```
CALL SSYMM (SIDE, UPLO,M, N, SALPHA, SA, LDA, SB, LDB, SBETA, SC, LDC)
CALL DSYMM (SIDE, UPLO, M, N, DALPHA, DA, LDA, DB, LDB, DBETA, DC, LDC)
CALL CSYMM (SIDE, UPLO, M, N, CALPHA, CA, LDA, CB, LDB, CBETA, CC, LDC)
CALL ZSYMM (SIDE, UPLO, M, N, ZALPHA, ZA, LDA, ZB, LDB, ZBETA, ZC, LDC)
```

For all data types, these subprograms set $C_{M \times N}$ to one of the expressions: $C \leftarrow \alpha A B+\beta C$ or $C \leftarrow \alpha B A+\beta C$, where $A$ is a symmetric matrix. The matrix $A$ is referenced either by its upper or lower triangular part. The character flags SIDE and UPLO determine the part of the matrix used and the operation performed.

## Matrix-Matrix Multiply, Hermitian

```
CALL CHEMM (SIDE, UPLO,M, N, CALPHA, CA, LDA, CB, LDB, CBETA, CC, LDC)
CALL ZHEMM (SIDE, UPLO,M, N, ZALPHA, ZA, LDA, ZB, LDB, ZBETA, ZC, LDC)
```

For all data types, these subprograms set $C_{M \times N}$ to one of the expressions: $C \leftarrow \alpha A B+\beta C$ or $C \leftarrow \alpha B A+\beta C$, where $A$ is an Hermitian matrix. The matrix $A$ is referenced either by its upper or lower triangular part. The character flags SIDE and UPLO determine the part of the matrix used and the operation performed.

## Rank-k Update, Symmetric

```
CALL SSYRK (UPLO, TRANS, N, K, SALPHA, SA, LDA, SBETA, SC, LDC)
CALL DSYRK (UPLO, TRANS, N, K, DALPHA, DA, LDA, DBETA, DC, LDC)
CALL CSYRK (UPLO, TRANS, N, K, CALPHA, CA, LDA, CBETA, CC, LDC)
CALL ZSYRK (UPLO, TRANS,N, K, ZALPHA, ZA, LDA, ZBETA, ZC, LDC)
```

For all data types, these subprograms set $C_{M \times N}$ to one of the expressions: $C \leftarrow \alpha A A^{T}+\beta C$ or $C \leftarrow \alpha A^{T} A+\beta C$. The matrix $C$ is referenced either by its upper or lower triangular part. The character flags UPLO and TRANS determine the part of the matrix used and the operation performed. In subprogram CSYRK and ZSYRK, only values ' $\mathrm{N}^{\prime}$ or ' $\mathrm{T}^{\prime}$ are allowed for TRANS; ${ }^{\prime} \mathrm{C}^{\prime}$ is not acceptable.

## Rank-k Update, Hermitian

$$
\begin{aligned}
& \text { CALL CHERK (UPLO, TRANS, N, K, SALPHA, CA, LDA, SBETA, CC, LDC) } \\
& \text { CALL ZHERK (UPLO, TRANS, N, K, DALPHA, ZA, LDA, DBETA, ZC, LDC) }
\end{aligned}
$$

For all data types, these subprograms set $C_{N \times N}$ to one of the expressions:

$$
C \leftarrow \alpha A \bar{A}^{T}+\beta C \text { or } C \leftarrow \alpha \bar{A}^{T} A+\beta C
$$

The matrix $C$ is referenced either by its upper or lower triangular part. The character flags UPLO and TRANS determine the part of the matrix used and the operation performed. CAUTION: Notice the scalar parameters $\alpha$ and $\beta$ are real, and the data in the matrices are complex. Only values ' $\mathrm{N}^{\prime}$ or ' C ' are allowed for TRANS; ' T ' is not acceptable.

## Rank-2k Update, Symmetric

```
CALL SSYR2K (UPLO, TRANS, N, K, SALPHA, SA, LDA, SB, LDB, SBETA, SC,
    LDC)
CALL DSYR2K (UPLO,TRANS, N, K, DALPHA, DA, LDA, DB, LDB, DBETA, DC,
    LDC)
CALL CSYR2K (UPLO, TRANS, N, K, CALPHA, CA, LDA, CB, LDB, CBETA, CC,
    LDC)
CALL ZSYR2K (UPLO, TRANS, N, K, ZALPHA, ZA, LDA, ZB, LDB, ZBETA, ZC,
    LDC)
```

For all data types, these subprograms set $C_{N \times N}$ to one of the expressions

$$
C \leftarrow \alpha A B^{T}+\alpha \beta A^{T}+\beta \mathrm{C} \text { or } C \leftarrow \alpha A^{T} B+\alpha B^{T} A+\beta C
$$

The matrix $C$ is referenced either by its upper or lower triangular part. The character flags UpLO and TRANS determine the part of the matrix used and the operation performed. In subprogram CSYR2K and ZSYR2K, only values ' $N$ ' or ' $T$ ' are allowed for TRANS; ' C ' is not acceptable.

## Rank-2k Update, Hermitian

```
CALL CHER2K (UPLO, TRANS, N, K, CALPHA, CA, LDA, CB, LDB, SBETA, CC,
                    LDC)
CALL ZHER2K (UPLO, TRANS, N, K, ZALPHA, ZA, LDA, ZB, LDB, DBETA, ZC,
LDC)
```

For all data types, these subprograms set $C_{N \times N}$ to one of the expressions

$$
C \leftarrow \alpha A \bar{B}^{T}+\bar{\alpha} B \bar{A}^{T}+\beta \mathrm{C} \text { or } C \leftarrow \alpha \bar{A}^{T} B+\bar{\alpha} \bar{B}^{T} A+\beta C
$$

The matrix $C$ is referenced either by its upper or lower triangular part. The character flags UPLO and TRANS determine the part of the matrix used and the operation performed. CAUTION: Notice the scalar parameter $\beta$ is real, and the data in the matrices are complex. In subprogram CHER2K and ZHER2K, only values ' $N$ ' or ' C ' are allowed for TRANS; ' T ' is not acceptable.

## Matrix-Matrix Multiply, Triangular

```
CALL STRMM (SIDE, UPLO, TRANSA, DIAGNL, M, N, SALPHA, SA, LDA, SB, LDB)
CALL DTRMM (SIDE, UPLO, TRANSA, DIAGNL, M, N, DALPHA, DA, LDA, DB, LDB)
CALL CTRMM (SIDE, UPLO, TRANSA, DIAGNL, M, N, CALPHA, CA, LDA, CB, LDB)
CALL ZTRMM (SIDE, UPLO, TRANSA, DIAGNL, M, N, ZALPHA, ZA, LDA, ZB, LDB)
```

For all data types, these subprograms set $B_{M \times N}$ to one of the_expressions:

$$
\begin{aligned}
& B \leftarrow \alpha A B, B \leftarrow \alpha A^{T} B, B \leftarrow \alpha B A, B \leftarrow \alpha B A^{T}, \\
& \text { or for complex data, } B \leftarrow \alpha \bar{A}^{T} B \text {, or } B \leftarrow \alpha B \bar{A}^{T}
\end{aligned}
$$

where $A$ is a triangular matrix. The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags SIDE, UPLO, TRANSA, and DIAGNL determine the part of the matrix used and the operation performed.

## Matrix-Matrix Solve, Triangular

```
CALL STRSM (SIDE, UPLO, TRANSA, DIAGNL, M, N, SALPHA, SA, LDA, SB, LDB)
CALL DTRSM (SIDE, UPLO, TRANSA, DIAGNL, M, N, DALPHA, DA, LDA, DB, LDB)
CALL CTRSM (SIDE, UPLO, TRANSA, DIAGNL, M, N, CALPHA, CA, LDA, CB, LDB)
CALL ZTRSM (SIDE, UPLO, TRANSA, DIAGNL, M, N, ZALPHA, ZA, LDA, ZB, LDB)
```

For all data types, these subprograms set $B_{M \times N}$ to one of the expressions:

$$
\begin{aligned}
& B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B, B \leftarrow \alpha B\left(A^{-1}\right)^{T}, \\
& \text { or for complex data, } B \leftarrow \alpha\left(\bar{A}^{T}\right)^{-1} B \text {, or } B \leftarrow \alpha B\left(\bar{A}^{T}\right)^{-1}
\end{aligned}
$$

where $A$ is a triangular matrix. The matrix $A$ is either referenced using its upper or lower triangular part and is unit or nonunit triangular. The character flags SIDE, UPLO, TRANSA, and DIAGNL determine the part of the matrix used and the operation performed.

## Other Matrix/Vector Operations

This section describes a set of routines for matrix/vector operations. The matrix copy and conversion routines are summarized by the following table:

| From | To |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  | Real <br> General | Complex <br> General | Real <br> Band | Complex <br> Band |
| Real General | CRGRG | CRGCG | CRGRB |  |
| Complex General |  | CCGCG |  | CCGCB |
| Real Band | CRBRG |  | CRBRB | CRBCB |
| Complex Band |  | CCBCG |  | CCBCB |
| Symmetric Full | CSFRG |  |  |  |
| Hermitian Full |  | CHFCG |  |  |
| Symmetric Band |  |  | CSBRB |  |
| Hermitian Band |  |  |  | CHBCB |

The matrix multiplication routines are summarized as follows:

| $\boldsymbol{A B}$ | $\boldsymbol{A}$ |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
| $\boldsymbol{B}$ | Real <br> Rect. | Complex <br> Rect. | Real <br> Band | Complex <br> Band |
| Real Rectangular | MRRRR |  |  |  |
| Complex Rect. |  | MCRCR |  |  |
| Vector | MURRV | MUCRV | MURBV | MUCBV |

The matrix norm routines are summarized as follows:

| $\\|\boldsymbol{A}\\|$ | Real <br> Rectangular | Real <br> Band | Complex <br> Band |
| :--- | :--- | :--- | :--- |
| $\infty-$ norm | NRIRR |  |  |
| 1-norm | NR1RR | NR1RB | NR1CB |
| Frobenius | NR2RR |  |  |

## CRGRG

Copies a real general matrix.

## Required Arguments

$\boldsymbol{A}$ - Matrix of order N. (Input)
B - Matrix of order N containing a copy of A. (Output)

## Optional Arguments

$N$ - Order of the matrices. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRGRG (A, B $[, \ldots]$ )
Specific: The specific interface names are S_CRGRG and D_CRGRG.

## FORTRAN 77 Interface

Single: CALL CRGRG (N, A, LDA, B, LDB)
Double: The double precision name is DCRGRG.

## Description

The routine CRGRG copies the real $N \times N$ general matrix $A$ into the real $N \times N$ general matrix $B$.

## Example

A real $3 \times 3$ general matrix is copied into another real $3 \times 3$ general matrix.

```
    USE CRGRG INT
    USE WRRRN_INT
    IMPLICIT NONE
    INTEGER LDA, LDB, N
    PARAMETER (LDA=3, LDB=3, N=3)
    REAL A(LDA,N), B(LDB,N)
                                    Set values for A
                                    A =( 0.0 1.0 1.0 )
```



```
    DATA A/0.0, 2* - 1.0, 1.0, 0.0, -1.0, 2*1.0, 0.0/
                                    Copy real matrix A to real matrix B
    CALL CRGRG (A, B)
    CALL WRRRN ('B', B)
    END
```


## Output

B

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 0.000 | 1.000 | 1.000 |
| 2 | -1.000 | 0.000 | 1.000 |
| 3 | -1.000 | -1.000 | 0.000 |

## CCGCG

Copies a complex general matrix.

## Required Arguments

$\boldsymbol{A}$ - Complex matrix of order N. (Input)
B - Complex matrix of order N containing a copy of A. (Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input) Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CCGCG (A, B $[, \ldots]$ )
Specific: The specific interface names are S_CCGCG and D_CCGCG.

## FORTRAN 77 Interface

Single: CALL CCGCG (N, A, LDA, B, LDB)
Double: The double precision name is DCCGCG.

## Description

The routine CCGCG copies the complex $N \times N$ general matrix $A$ into the complex $N \times N$ general matrix $B$.

## Example

A complex $3 \times 3$ general matrix is copied into another complex $3 \times 3$ general matrix.

```
USE CCGCG_INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N
PARAMETER (LDA=3, LDB=3, N=3)
    COMPLEX A(LDA,N), B(LDB,N)
        Set values for A
        A = ( 0.0+0.0i 1.0+1.0i 1.0+1.0i )
                            ( -1.0-1.0i 0.0+0.0i 1.0+1.0i )
                            ( -1.0-1.0i -1.0-1.0i 0.0+0.0i )
    DATA A/ (0.0,0.0), 2* (-1.0,-1.0), (1.0,1.0), (0.0,0.0), &
        (-1.0,-1.0), 2*(1.0,1.0), (0.0,0.0)/
        Copy matrix A to matrix B
    CALL CCGCG (A, B)
CALL WRCRN ('B', B)
END
```


## Output

B
23

```
1 ( 0.000, 0.000) ( 1.000, 1.000) ( 1.000, 1.000)
2 (-1.000,-1.000) ( 0.000, 0.000) ( 1.000, 1.000)
3 (-1.000,-1.000) (-1.000,-1.000) (0.000, 0.000)
```


## CRBRB

Copies a real band matrix stored in band storage mode.

## Required Arguments

$\boldsymbol{A}$ - Real band matrix of order N. (Input)

NLCA - Number of lower codiagonals in A. (Input)
NUCA - Number of upper codiagonals in A. (Input)
B - Real band matrix of order N containing a copy of A. (Output)
NLCB - Number of lower codiagonals in B. (Input) NLCB must be at least as large as NLCA.

NUCB - Number of upper codiagonals in B. (Input) nUCB must be at least as large as NUCA.

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRBRB (A, NLCA, NUCA, B, NLCB, NUCB [, ...])
Specific: The specific interface names are $S_{\text {_ }}$ CRBRB and $D_{-}$CRBRB.

## FORTRAN 77 Interface

Single: CALL CRBRB (N, A, LDA, NLCA, NUCA, B, LDB, NLCB, NUCB)

Double: The double precision name is DCRBRB.

## Description

The routine CRBRB copies the real band matrix $A$ in band storage mode into the real band matrix $B$ in band storage mode.

## Example

A real band matrix of order 3, in band storage mode with one upper codiagonal, and one lower codiagonal is copied into another real band matrix also in band storage mode.

```
USE CRBRB_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N, NLCA, NLCB, NUCA, NUCB
PARAMETER (LDA=3, LDB=3, N=3, NLCA=1, NLCB=1, NUCA=1, NUCB=1)
REAL A(LDA,N), B (LDB,N)
    Set values for A (in band mode)
    A =( ( 0.0 1.0 1.0 )
                                    ( 1.0 1.0 1.0)
                                    (1.0 1.0 0.0 )
DATA A/0.0, 7*1.0, 0.0/
CALL CRBRB (A, NLCA, NUCA, B, NLCB, NUCB)
CALL WRRRN ('B', B)
END
```


## Output

## B

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 0.000 | 1.000 | 1.000 |
| 2 | 1.000 | 1.000 | 1.000 |
| 3 | 1.000 | 1.000 | 0.000 |

## CCBCB

Copies a complex band matrix stored in complex band storage mode.

## Required Arguments

$\boldsymbol{A}$ - Complex band matrix of order N. (Input)
NLCA - Number of lower codiagonals in A. (Input)

NUCA - Number of upper codiagonals in A. (Input)

B - Complex matrix of order N containing a copy of A. (Output)
NLCB - Number of lower codiagonals in B. (Input)
NLCB must be at least as large as NLCA.
$N U C B$ - Number of upper codiagonals in B. (Input)
nUCB must be at least as large as nUCA.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB = SIZE (B, 1 ).

## FORTRAN 90 Interface

Generic: CALL CCBCB (A, NLCA, NUCA, B, NLCB, NUCB [, ...])
Specific: The specific interface names are $S_{-} C C B C B$ and $D_{-} C C B C B$.

## FORTRAN 77 Interface

Single: CALL CCBCB (N, A, LDA, NLCA, NUCA, B, LDB, NLCB, NUCB)
Double: The double precision name is DCCBCB.

## Description

The routine $C C B C B$ copies the complex band matrix $A$ in band storage mode into the complex band matrix $B$ in band storage mode.

## Example

A complex band matrix of order 3 in band storage mode with one upper codiagonal and one lower codiagonal is copied into another complex band matrix in band storage mode.

```
USE CCBCB INT
USE WRCRN__INT
```

```
    IMPLICIT NONE
    INTEGER LDA, LDB, N, NLCA, NLCB, NUCA, NUCB
    PARAMETER (LDA=3, LDB=3, N=3, NLCA=1, NLCB=1, NUCA=1, NUCB=1)
    COMPLEX A(LDA,N), B(LDB,N)
        Set values for A (in band mode)
        A = ( 0.0+0.0i 1.0+1.0i 1.0+1.0i )
        ( 1.0+1.0i 1.0+1.0i 1.0+1.0i )
        ( 1.0+1.0i 1.0+1.0i 0.0+0.0i )
    DATA A/(0.0,0.0), 7* (1.0,1.0), (0.0,0.0)/
        Copy A to B
    CALL CCBCB (A, NLCA, NUCA, B, NLCB, NUCB)
    CALL WRCRN ('B', B)
    END
```


## Output

B

|  | 2 |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
| 1 | $(0.000,0.000)$ | $(1.000,1.000)$ | $(1.000,1.000)$ |  |
| 2 | $(1.000,1.000)$ | $(1.000,1.000)$ | $(1.000,1.000)$ |  |
| 3 | $(1.000,1.000)$ | $(1.000,1.000)$ | $(0.000,0.000)$ |  |

## CRGRB

Converts a real general matrix to a matrix in band storage mode.

## Required Arguments

$\boldsymbol{A}-$ Real N by N matrix. (Input)
$N L C$ - Number of lower codiagonals in B. (Input)
NUC - Number of upper codiagonals in B. (Input)
B-Real (NUC $+1+\mathrm{NLC})$ by N array containing the band matrix in band storage mode.
(Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRGRB (A, NLC, NUC, B [,...])
Specific: The specific interface names are S_CRGRB and D_CRGRB.

## FORTRAN 77 Interface

Single: CALL CRGRB (N, A, LDA, NLC, NUC, B, LDB)
Double: The double precision name is DCRGRB.

## Description

The routine CRGRB converts the real general $N \times N$ matrix $A$ with $m_{u}=$ NUC upper codiagonals and $m_{l}=$ NLC lower codiagonals into the real band matrix $B$ of order $N$. The first $m_{u}$ rows of $B$ then contain the upper codiagonals of $A$, the next row contains the main diagonal of $A$, and the last $m_{l}$ rows of $B$ contain the lower codiagonals of $A$.

## Example

A real $4 \times 4$ matrix with one upper codiagonal and three lower codiagonals is copied to a real band matrix of order 4 in band storage mode.

```
USE CRGRB INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N, NLC, NUC
PARAMETER (LDA=4, LDB=5, N=4, NLC=3, NUC=1)
REAL A(LDA,N), B(LDB,N)
                                    Set values for A
                                    A = (\begin{array}{llll}{1.0}&{2.0}&{0.0}&{0.0}\end{array})
                                    ( -2.0 1.0 3.0 0.0)
                                    ( 0.0 -3.0 1.0 4.0)
                                    ( -7.0 0.0 -4.0 1.0)
DATA A/1.0, -2.0, 0.0, -7.0, 2.0, 1.0, -3.0, 0.0, 0.0, 3.0, 1.0, &
    -4.0, 0.0, 0.0, 4.0, 1.0/
                                    Convert A to band matrix B
CALL CRGRB (A, NLC, NUC, B)
CALL WRRRN ('B', B)
END
```


## Output

|  |  | B |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 |
| 1 | 0.000 | 2.000 | 3.000 | 4.000 |
| 2 | 1.000 | 1.000 | 1.000 | 1.000 |
| 3 | -2.000 | -3.000 | -4.000 | 0.000 |
| 4 | 0.000 | 0.000 | 0.000 | 0.000 |
| 5 | -7.000 | 0.000 | 0.000 | 0.000 |

## CRBRG

Converts a real matrix in band storage mode to a real general matrix.

## Required Arguments

A-Real (NUC $+1+\mathrm{NLC})$ by N array containing the band matrix in band storage mode.
(Input)
$N L C$ - Number of lower codiagonals in A. (Input)
$N U C$ - Number of upper codiagonals in A. (Input)
B-Real n by n array containing the matrix. (Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRBRG (A, NLC, NUC, B [, ...])

Specific: The specific interface names are S_CRBRG and D_CRBRG.

## FORTRAN 77 Interface

Single: CALL CRBRG (N, A, LDA, NLC, NUC, B, LDB)

Double: The double precision name is DCRBRG.

## Description

The routine CRBRG converts the real band matrix $A$ of order $N$ in band storage mode into the real $N \times N$ general matrix $B$ with $m_{u}=$ NUC upper codiagonals and $m_{l}=$ NLC lower codiagonals. The first $m_{u}$ rows of $A$ are copied to the upper codiagonals of $B$, the next row of $A$ is copied to the diagonal of $B$, and the last $m_{l}$ rows of $A$ are copied to the lower codiagonals of $B$.

## Example

A real band matrix of order 3 in band storage mode with one upper codiagonal and one lower codiagonal is copied to a $3 \times 3$ real general matrix.

```
USE CRBRG_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N, NLC, NUC
PARAMETER (LDA=3, LDB=3, N=3, NLC=1, NUC=1)
REAL A(LDA,N), B(LDB,N)
                                    Set values for A (in band mode)
                                    A = ( 0.0 1.0 1.0)
                                    ( 4.0 3.0 2.0)
                                    (2.0 2.0 0.0)
DATA A/0.0, 4.0, 2.0, 1.0, 3.0, 2.0, 1.0, 2.0, 0.0/
                                    Convert band matrix A to matrix B
CALL CRBRG (A, NLC, NUC, B)
CALL WRRRN ('B', B)
END
```


## Output

```
            B
```

|  | 1 | 2 | 3 |
| ---: | ---: | ---: | ---: |
| 1 | 4.000 | 1.000 | 0.000 |
| 2 | 2.000 | 3.000 | 1.000 |
| 3 | 0.000 | 2.000 | 2.000 |

## CCGCB

Converts a complex general matrix to a matrix in complex band storage mode.

## Required Arguments

A-Complex n by n array containing the matrix. (Input)
$N L C$ - Number of lower codiagonals in B. (Input)
NUC - Number of upper codiagonals in B. (Input)
$\boldsymbol{B}$ - Complex (NUC $+1+\mathrm{NLC}$ ) by N array containing the band matrix in band storage mode. (Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CCGCB (A, NLC, NUC, B [,...])
Specific: The specific interface names are S_CCGCB and D_CCGCB.

## FORTRAN 77 Interface

Single: CALL CCGCB (N, A, LDA, NLC, NUC, B, LDB)
Double: The double precision name is DCCGCB.

## Description

The routine CCGCB converts the complex general matrix $A$ of order $N$ with $m_{u}=$ NUC upper codiagonals and $m_{l}=$ NLC lower codiagonals into the complex band matrix $B$ of order $N$ in band storage mode. The first $m_{u}$ rows of $B$ then contain the upper codiagonals of $A$, the next row contains the main diagonal of $A$, and the last $m_{l}$ rows of $B$ contain the lower codiagonals of $A$.

## Example

A complex general matrix of order 4 with one upper codiagonal and three lower codiagonals is copied to a complex band matrix of order 4 in band storage mode.

```
USE CCGCB_INT
USE WRCRN_INT
IMPLICIT NONE
```

```
!
    INTEGER LDA, LDB, N, NLC, NUC
    PARAMETER (LDA=4, LDB=5, N=4, NLC=3, NUC=1)
    COMPLEX A(LDA,N), B(LDB,N)
        Set values for A
        A = ( 1.0+0.0i 2.0+1.0i 0.0+0.0i 0.0+0.0i )
                            ( -2.0+1.0i 1.0+0.0i 3.0+2.0i 0.0+0.0i )
                            ( 0.0+0.0i -3.0+2.0i 1.0+0.0i 4.0+3.0i )
                            ( -7.0+1.0i 0.0+0.0i -4.0+3.0i 1.0+0.0i )
    DATA A/ (1.0,0.0), (-2.0,1.0), (0.0,0.0), (-7.0,1.0), (2.0,1.0), &
        (1.0,0.0), (-3.0,2.0), (0.0,0.0), (0.0,0.0), (3.0,2.0), &
        (1.0,0.0), (-4.0,3.0), (0.0,0.0), (0.0,0.0), (4.0,3.0), &
        (1.0,0.0)/
        Convert A to band matrix B
    CALL CCGCB (A, NLC, NUC, B)
    CALL WRCRN ('B', B)
    END
```


## Output

| B |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 |  | 2 |  | 3 |  | 4 |
| ( 0.000, 0.000) | ( 2.000, | 1.000) | ( 3.000, | $2.000)$ | ( 4.000, | $3.000)$ |
| ( 1.000, 0.000) | ( 1.000, | $0.000)$ | ( 1.000, | $0.000)$ | ( 1.000, | $0.000)$ |
| (-2.000, 1.000) | (-3.000, | $2.000)$ | (-4.000, | $3.000)$ | ( 0.000, | $0.000)$ |
| ( 0.000, 0.000) | ( 0.000, | $0.000)$ | ( 0.000, | $0.000)$ | ( 0.000, | $0.000)$ |
| (-7.000, 1.000) | ( 0.000, | $0.000)$ | ( 0.000, | $0.000)$ | ( 0.000, | $0.000)$ |

## CCBCG

Converts a complex matrix in band storage mode to a complex matrix in full storage mode.

## Required Arguments

$\boldsymbol{A}-$ Complex (NUC $+1+\mathrm{NLC}$ ) by N matrix containing the band matrix in band mode. (Input)
$N L C$ - Number of lower codiagonals in A. (Input)
NUC - Number of upper codiagonals in A. (Input)
$\boldsymbol{B}$ - Complex N by N matrix containing the band matrix in full mode. (Output)

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.

LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CCBCG (A, NLC, NUC, B [, ...])
Specific: The specific interface names are $S_{-} C C B C G$ and $D_{-} C C B C G$.

## FORTRAN 77 Interface

Single: CALL CCBCG (N, A, LDA, NLC, NUC, B, LDB)
Double: The double precision name is DCCBCG.

## Description

The routine CCBCG converts the complex band matrix $A$ of order $N$ with $m_{u}=$ NUC upper codiagonals and $m_{l}=$ NLC lower codiagonals into the $N \times N$ complex general matrix $B$. The first $m_{u}$ rows of $A$ are copied to the upper codiagonals of $B$, the next row of $A$ is copied to the diagonal of $B$, and the last $m_{l}$ rows of $A$ are copied to the lower codiagonals of $B$.

## Example

A complex band matrix of order 4 in band storage mode with one upper codiagonal and three lower codiagonals is copied into a $4 \times 4$ complex general matrix.

```
USE CCBCG_INT
USE WRCRN INT
IMPLICIT NONE
INTEGER LDA, LDB, N, NLC, NUC
PARAMETER (LDA=5, LDB=4, N=4, NLC=3, NUC=1)
COMPLEX A(LDA,N), B(LDB,N)
    Set values for A (in band mode)
    A = ( 0.0+0.0i 2.0+1.0i 3.0+2.0i 4.0+3.0i )
        ( 1.0+0.0i 1.0+0.0i 1.0+0.0i 1.0+0.0i )
        ( -2.0+1.0i -3.0+2.0i -4.0+3.0i 0.0+0.0i )
            ( 0.0+0.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
            ( -7.0+1.0i 0.0+0.0i 0.0+0.0i 0.0+0.0i )
DATA A/(0.0,0.0), (1.0,0.0), (-2.0,1.0), (0.0,0.0), (-7.0,1.0), &
    (2.0,1.0), (1.0,0.0), (-3.0,2.0), 2*(0.0,0.0), (3.0,2.0), &
```

```
(1.0,0.0), (-4.0,3.0), 2* (0.0,0.0), (4.0,3.0), (1.0,0.0), &
3*(0.0,0.0)/
CALL CCBCG (A, NLC, NUC, B)
CALL WRCRN ('B', B)
```

!
END

## Output

|  | B |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 1 |  | 2 |  | 3 |  | 4 |
| 1 | ( 1.000, 0.000) | ( 2.000, | $1.000)$ | ( 0.000, | $0.000)$ | ( 0.000, | $0.000)$ |
| 2 | (-2.000, 1.000) | ( 1.000, | $0.000)$ | ( 3.000, | $2.000)$ | ( 0.000, | $0.000)$ |
| 3 | ( 0.000, 0.000) | (-3.000, | $2.000)$ | ( 1.000, | $0.000)$ | ( 4.000, | $3.000)$ |
| 4 | (-7.000, 1.000) | ( 0.000, | $0.000)$ | (-4.000, | $3.000)$ | ( 1.000, | $0.000)$ |

## CRGCG

Copies a real general matrix to a complex general matrix.

## Required Arguments

A-Real matrix of order N. (Input)
B - Complex matrix of order N containing a copy of A. (Output)

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRGCG (A, B [, ...])

Specific: The specific interface names are $S_{-}$CRGCG and D_CRGCG.

## FORTRAN 77 Interface

Single: CALL CRGCG (N, A, LDA, B, LDB)

Double: The double precision name is DCRGCG.

## Description

The routine CRGCG copies a real $N \times N$ matrix to a complex $N \times N$ matrix.

## Example

A $3 \times 3$ real matrix is copied to a $3 \times 3$ complex matrix.

```
USE CRGCG_INT
USE WRCRN__INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDB, N
PARAMETER (LDA=3, LDB=3, N=3)
    REAL A (LDA,N)
    COMPLEX B(LDB,N)
```

```
! Set values for A
```

! Set values for A
A = ( 2.0 1.0 3.0 )
A = ( 2.0 1.0 3.0 )
( 4.0 1.0 0.0 )
( 4.0 1.0 0.0 )
( -1.0 2.0 0.0 )
( -1.0 2.0 0.0 )
DATA A/2.0, 4.0, -1.0, 1.0, 1.0, 2.0, 3.0, 0.0, 0.0/
DATA A/2.0, 4.0, -1.0, 1.0, 1.0, 2.0, 3.0, 0.0, 0.0/
Convert real A to complex B
Convert real A to complex B
CALL CRGCG (A, B)
CALL CRGCG (A, B)
CALL WRCRN ('B', B)
CALL WRCRN ('B', B)
END

```
    END
```

$!$

## Output

B

| 1 | $(2.000,0.000)$ | $(1.000,0.000)$ | $(3.000,0.000)$ |
| :--- | ---: | ---: | ---: | :--- |
| 2 | $(4.000,0.000)$ | $(1.000,0.000)$ | $(0.000,0.000)$ |
| 3 | $(-1.000,0.000)$ | $(2.000,0.000)$ | $(0.000,0.000)$ |

## CRRCR

Copies a real rectangular matrix to a complex rectangular matrix.

## Required Arguments

A-Real NRA by NCA rectangular matrix. (Input)
B - Complex NRB by NCB rectangular matrix containing a copy of A. (Output)

## Optional Arguments

NRA - Number of rows in A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.

NCA - Number of columns in A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N R B$ - Number of rows in B. (Input)
It must be the same as NRA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.

NCB - Number of columns in B. (Input)
It must be the same as NCA.
Default: NCB $=$ SIZE $(B, 2)$.
$\mathbf{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB = SIZE (B, 1 ).

## FORTRAN 90 Interface

Generic: CALL CRRCR (A, B [, ...])

Specific: The specific interface names are S_CRRCR and D_CRRCR.

## FORTRAN 77 Interface

Single: CALL CRRCR (NRA, NCA, A, LDA, NRB, NCB, B, LDB)

Double: The double precision name is DCRRCR.

## Description

The routine CRRCR copies a real rectangular matrix to a complex rectangular matrix.

## Example

A $3 \times 2$ real matrix is copied to a $3 \times 2$ complex matrix.

```
USE CRRCR_INT
USE WRCRN_INT
IMPLICIT NONE
```

```
INTEGER LDA, LDB, NCA, NCB, NRA, NRB
PARAMETER (LDA=3, LDB=3, NCA=2, NCB=2, NRA=3, NRB=3)
!
    REAL A (LDA,NCA)
    COMPLEX B (LDB,NCB)
!
!
!
    DATA A/1.0, 2.0, 3.0, 4.0, 5.0, 6.0/
    CALL CRRCR (A, B)
    CALL WRCRN ('B', B)
    END
```


## Output

B
12
$(1.000,0.000)(4.000,0.000)$
$(2.000,0.000)(5.000,0.000)$
$(3.000,0.000)(6.000,0.000)$

## CRBCB

Converts a real matrix in band storage mode to a complex matrix in band storage mode.

## Required Arguments

A-Real band matrix of order N. (Input)
NLCA - Number of lower codiagonals in A. (Input)

NUCA - Number of upper codiagonals in A. (Input)
B - Complex matrix of order N containing a copy of A. (Output)
NLCB - Number of lower codiagonals in B. (Input)
NLCB must be at least as large as NLCA.

NUCB - Number of upper codiagonals in B. (Input) NUCB must be at least as large as NUCA.

## Optional Arguments

$N$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\mathbf{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(\mathrm{B}, 1)$.

## FORTRAN 90 Interface

Generic: CALL CRBCB (A, NLCA, NUCA, B, NLCB, NUCB [, ...])
Specific: The specific interface names are $S_{-}$CRBCB and D_CRBCB.

## FORTRAN 77 Interface

Single: CALL CRBCB (N, A, LDA, NLCA, NUCA, B, LDB, NLCB, NUCB)
Double: The double precision name is DCRBCB.

## Description

The routine CRBCB converts a real band matrix in band storage mode with NUCA upper codiagonals and NLCA lower codiagonals into a complex band matrix in band storage mode with NUCB upper codiagonals and NLCB lower codiagonals.

## Example

A real band matrix of order 3 in band storage mode with one upper codiagonal and one lower codiagonal is copied into another complex band matrix in band storage mode.

```
USE CRBCB INT
USE WRCRN_INT
IMPLICIT NONE
Declare variables
PARAMETER (LDA=3, LDB=3, N=3, NLCA=1, NLCB=1, NUCA=1, NUCB=1)
REAL A(LDA,N)
COMPLEX B(LDB,N)
    Set values for A (in band mode)
    A =( 0.0 1.0 1.0)
    (1.0 1.0 1.0)
    (1.0 1.0 0.0)
DATA A/0.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 0.0/
    Convert real band matrix A
    to complex band matrix B
CALL CRBCB (A, NLCA, NUCA, B, NLCB, NUCB)
    Print results
```


## Output

|  | B |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| 1 | $(0.000,0.000)$ | $(1.000,0.000)$ | $(1.000,0.000)$ |  |
| 2 | $(1.000,0.000)$ | $(1.000,0.000)$ | $(1.000,0.000)$ |  |
| 3 | $(1.000,0.000)$ | $(1.000,0.000)$ | $(0.000,0.000)$ |  |

## CSFRG

Extends a real symmetric matrix defined in its upper triangle to its lower triangle.

## Required Arguments

$\boldsymbol{A}-\mathrm{N}$ by N symmetric matrix of order N to be filled out. (Input/Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL CSFRG (A [,...])
Specific: The specific interface names are S_CSFRG and D_CSFRG.

## FORTRAN 77 Interface

Single: CALL CSFRG ( $\mathrm{N}, \mathrm{A}, \mathrm{LDA}$ )
Double: The double precision name is DCSFRG.

## Description

The routine CSFRG converts an $N \times N$ matrix $A$ in symmetric mode into a general matrix by filling in the lower triangular portion of $A$ using the values defined in its upper triangular portion.

## Example

The lower triangular portion of a real $3 \times 3$ symmetric matrix is filled with the values defined in its upper triangular portion.

```
USE CSFRG INT
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
REAL A(LDA,N)
    Set values for A
    A =( (0.0
    ( 1.0 5.0 )
DATA A/3*0.0, 3.0, 1.0, 0.0, 4.0, 5.0, 2.0/
CALL CSFRG (A)
CALL WRRRN ('A', A)
END
```


## Output

|  | A |  |  |
| ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 |
| 1 | 0.000 | 3.000 | 4.000 |
| 2 | 3.000 | 1.000 | 5.000 |
| 3 | 4.000 | 5.000 | 2.000 |

## CHFCG

Extends a complex Hermitian matrix defined in its upper triangle to its lower triangle.

## Required Arguments

A-Complex Hermitian matrix of order N. (Input/Output)
On input, the upper triangle of A defines a Hermitian matrix. On output, the lower triangle of A is defined so that A is Hermitian.

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL CHFCG (A $[, \ldots]$ )

Specific: The specific interface names are __CHFCG and D_CHFCG. $^{\text {C }}$

## FORTRAN 77 Interface

Single: CALL CHFCG ( $\mathrm{N}, \mathrm{A}$, LDA)
Double: The double precision name is DCHFCG.

## Description

The routine CHFCG converts an $N \times N$ complex matrix $A$ in Hermitian mode into a complex general matrix by filling in the lower triangular portion of $A$ using the values defined in its upper triangular portion.

## Comments

Informational errors
Type Code
31 The matrix is not Hermitian. It has a diagonal entry with a small imaginary part.
42 The matrix is not Hermitian. It has a diagonal entry with an imaginary part.

## Example

A complex $3 \times 3$ Hermitian matrix defined in its upper triangle is extended to its lower triangle.

```
USE CHFCG INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER LDA, N
PARAMETER (LDA=3, N=3)
COMPLEX A(LDA,N)
    A=(\begin{array}{clll}{\mathrm{ Set values for A}}\\{1.0+0.0i 1.0+1.0i 1.0+2.0i ( )}\end{array})
DATA A/(1.0,0.0), 2* (0.0,0.0), (1.0,1.0), (2.0,0.0), (0.0,0.0), &
    (1.0,2.0), (2.0.2.0), (3.0,0.0)/
```

```
! Fill in lower Hermitian matrix
    CALL CHFCG (A)
    CALL WRCRN ('A', A)
    END
```


## Output

A
1
$)^{(1.000,1.000)}$
$(2.000,0.000)$
$(2.000,-2.000)$$\left(\begin{array}{l}(2.000,2.000) \\ (2.000,2.000) \\ (2.000,0.000)\end{array}\right.$

## CSBRB

Copies a real symmetric band matrix stored in band symmetric storage mode to a real band matrix stored in band storage mode.

## Required Arguments

A-Real band symmetric matrix of order n. (Input)

NUCA - Number of codiagonals in A. (Input)

B - Real band matrix of order N containing a copy of A. (Output)
NLCB - Number of lower codiagonals in B. (Input) NLCB must be at least as large as NUCA.

NUCB - Number of upper codiagonals in B. (Input) nUCB must be at least as large as NUCA.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: $\quad$ CALL $\operatorname{CSBRB}(A, N U C A, B, N L C B, N U C B[, \ldots])$
Specific: The specific interface names are S_CSBRB and D_CSBRB.

## FORTRAN 77 Interface

Single: CALL CSBRB (N, A, LDA, NUCA, B, LDB, NLCB, NUCB)
Double: The double precision name is DCSBRB.

## Description

The routine CSBRB copies a real matrix $A$ stored in symmetric band mode to a matrix B stored in band mode. The lower codiagonals of $B$ are set using the values from the upper codiagonals of $A$.

## Example

A real matrix of order 4 in band symmetric storage mode with 2 upper codiagonals is copied to a real matrix in band storage mode with 2 upper codiagonals and 2 lower codiagonals.

```
USE CSBRB INT
USE WRRRN_INT
IMPLICIT NONE
Declare variables
INTEGER LDA, LDB, N, NLCB, NUCA, NUCB
PARAMETER (N=4, NUCA=2, LDA=NUCA+1, NLCB=NUCA, NUCB=NUCA, &
        LDB=NLCB+NUCB+1)
REAL A(LDA,N), B(LDB,N)
                            Set values for A, in band mode
                            A =( 0.0 0.0 2.0 1.0 )
                            ( 0.0 2.0 3.0 1.0 )
                            ( 1.0 2.0 3.0 4.0 )
DATA A/2*0.0, 1.0, 0.0, 2.0, 2.0, 2.0, 3.0, 3.0, 1.0, 1.0, 4.0/
                Copy A to B
CALL CSBRB (A, NUCA, B, NLCB, NUCB)
CALL WRRRN ('B', B)
END
```


## Output

|  | B |  |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 |
| 1 | 0.000 | 0.000 | 2.000 | 1.000 |
| 2 | 0.000 | 2.000 | 3.000 | 1.000 |
| 3 | 1.000 | 2.000 | 3.000 | 4.000 |


| 4 | 2.000 | 3.000 | 1.000 | 0.000 |
| :--- | :--- | :--- | :--- | :--- |
| 5 | 2.000 | 1.000 | 0.000 | 0.000 |

## CHBCB

Copies a complex Hermitian band matrix stored in band Hermitian storage mode to a complex band matrix stored in band storage mode.

## Required Arguments

$\boldsymbol{A}$ - Complex band Hermitian matrix of order N. (Input)
$N U C A$ - Number of codiagonals in A. (Input)
$\boldsymbol{B}$ - Complex band matrix of order N containing a copy of A. (Output)
NLCB - Number of lower codiagonals in B. (Input)
NLCB must be at least as large as nUCA.
NUCB - Number of upper codiagonals in B. (Input) nUCB must be at least as large as NUCA.

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A and B. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL CHBCB (A, NUCA, B, NLCB, NUCB [, ...])
Specific: The specific interface names are $S_{\text {_ }}$ CHBCB and D_CHBCB.

## FORTRAN 77 Interface

Single: CALL CHBCB (N, A, LDA, NUCA, B, LDB, NLCB, NUCB)
Double: The double precision name is $\operatorname{DCHBCB}$.

## Description

The routine CSBRB copies a complex matrix $A$ stored in Hermitian band mode to a matrix $B$ stored in complex band mode. The lower codiagonals of $B$ are filled using the values in the upper codiagonals of $A$.

## Comments

Informational errors
Type Code
31 An element on the diagonal has a complex part that is near zero, the complex part is set to zero.
41 An element on the diagonal has a complex part that is not zero.

## Example

A complex Hermitian matrix of order 3 in band Hermitian storage mode with one upper codiagonal is copied to a complex matrix in band storage mode.

```
USE CHBCB INT
USE WRCRN_INT
IMPLICIT NONE
LDA, LDB, N, NLCB, NUCA, NUCB
PARAMETER (N=3, NUCA=1, LDA=NUCA+1, NLCB=NUCA, NUCB=NUCA, &
            LDB=NLCB+NUCB+1)
COMPLEX A (LDA,N), B(LDB,N)
                                    Set values for A (in band mode)
                                    A = ( 0.0+0.0i -1.0+1.0i -2.0+2.0i )
DATA A/ (0.0,0.0), (1.0,0.0), (-1.0,1.0), (1.0,0.0), (-2.0,2.0), &
        (1.0,0.0)/
                            Copy a complex Hermitian band matrix
                            to a complex band matrix
CALL CHBCB (A, NUCA, B, NLCB, NUCB)
CALL WRCRN ('B', B)
END
```


## Output

|  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | B |  |  |  | 2 |
| 2 | $(0.000,0.000)$ | $(-1.000,1.000)$ | $(-2.000,2.000)$ |  |  |
| 3 | $(1.000,0.000)$ | $(1.000,0.000)$ | $(1.000,0.000)$ |  |  |
| $(-1.000,-1.000)$ | $(-2.000,-2.000)$ | $(0.000,0.000)$ |  |  |  |

## TRNRR

Transposes a rectangular matrix.

## Required Arguments

A-Real NRA by NCA matrix in full storage mode. (Input)
$\boldsymbol{B}$ - Real NRB by NCB matrix in full storage mode containing the transpose of A. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
$N C A-$ Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N R B$ - Number of rows of B. (Input)
NRB must be equal to NCA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.
$N C B-$ Number of columns of B. (Input)
NCB must be equal to NRA.
Default: $\mathrm{NCB}=\operatorname{SIZE}(\mathrm{B}, 2)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL TRNRR (A, B [,...])
Specific: The specific interface names are S_TRNRR and D_TRNRR.

## FORTRAN 77 Interface

Single: CALL TRNRR (NRA, NCA, A, LDA, NRB, NCB, B, LDB)

Double: The double precision name is DTRNRR.

## Description

The routine TRNRR computes the transpose $B=A^{T}$ of a real rectangular matrix $A$.

## Comments

If $\operatorname{LDA}=\operatorname{LDB}$ and $N R A=N C A$, then $A$ and $B$ can occupy the same storage locations; otherwise, $A$ and $B$ must be stored separately.

## Example

Transpose the $5 \times 3$ real rectangular matrix $A$ into the $3 \times 5$ real rectangular matrix $B$.

```
USE TRNRR INT
USE WRRRN_INT
IMPLICIT NONE
! Declare variables
INTEGER NCA, NCB, NRA, NRB
PARAMETER (NCA=3, NCB=5, NRA=5, NRB=3)
REAL A(NRA,NCA), B(NRB,NCB)
                                    Set values for A
                                    A =( 11.0 12.0 13.0 )
                                    ( 21.0 22.0 23.0 )
                                    ( 31.0 32.0 33.0 )
                                    (41.0 42.0 43.0 )
                                    ( 51.0 52.0 53.0 )
DATA A/11.0, 21.0, 31.0, 41.0, 51.0, 12.0, 22.0, 32.0, 42.0,&
    52.0, 13.0, 23.0, 33.0, 43.0, 53.0/
                            B = transpose(A)
CALL TRNRR (A, B)
CALL WRRRN ('B = trans(A)', B)
END
```


## Output

|  | $\mathrm{B}=\operatorname{trans}(\mathrm{A})$ |  |  |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | :---: |
|  | 1 | 2 | 3 | 4 | 5 |  |
| 1 | 11.00 | 21.00 | 31.00 | 41.00 | 51.00 |  |
| 2 | 12.00 | 22.00 | 32.00 | 42.00 | 52.00 |  |
| 3 | 13.00 | 23.00 | 33.00 | 43.00 | 53.00 |  |

## MXTXF

Computes the transpose product of a matrix, $A^{T} A$.

## Required Arguments

A-Real NRA by NCA rectangular matrix. (Input)
The transpose product of $A$ is to be computed.

B-Real NB by NB symmetric matrix containing the transpose product $A^{T} A$. (Output)

## Optional Arguments

$N R A$ - Number of rows in A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.

NCA - Number of columns in A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N B$ - Order of the matrix B. (Input)
NB must be equal to NCA.
Default: NB $=\operatorname{SIZE}(B, 1)$.

LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB = SIZE (B, 1 ).

## FORTRAN 90 Interface

Generic: CALL MXTXF (A, B [, ...])

Specific: The specific interface names are S_MXTXF and D_MXTXF.

## FORTRAN 77 Interface

Single: CALL MXTXF (NRA, NCA, A, LDA, NB, B, LDB)
Double: The double precision name is DMXTXF.

## Description

The routine MXTXF computes the real general matrix $B=A^{T} A$ given the real rectangular matrix $A$.

## Example

Multiply the transpose of a $3 \times 4$ real matrix by itself. The output matrix will be a $4 \times 4$ real symmetric matrix.

```
        USE MXTXF_INT
        USE WRRRN_INT
        IMPLICIT NONE
        INTEGER NB, NCA, NRA
        PARAMETER (NB=4, NCA=4, NRA=3)
        REAL A(NRA,NCA), B (NB,NB)
        Set values for A
        A =( 3.0 1.0 4.0 2.0 )
        ( 0.0 2.0 1.0 -1.0 )
        ( 6.0 1.0 3.0 2.0 )
        DATA A/3.0, 0.0, 6.0, 1.0, 2.0, 1.0, 4.0, 1.0, 3.0, 2.0, -1.0, &
        2.0/
        CALL MXTXF (A, B)
        CALL WRRRN ('B = trans(A)*A', B)
        END
```


## Output

            \(B=\operatorname{trans}(A) * A\)
    \(\begin{array}{llrr}1 & 2 & 3 & 4\end{array}\)
    $\begin{array}{rrrrr}1 & 45.00 & 9.00 & 30.00 & 18.00 \\ 2 & 9.00 & 6.00 & 9.00 & 2.00\end{array}$
$3 \quad 30.00 \quad 9.00 \quad 26.00 \quad 13.00$

| 4 | 18.00 | 2.00 | 13.00 | 9.00 |
| :--- | :--- | :--- | :--- | :--- |

## MXTYF

Multiplies the transpose of matrix $A$ by matrix $B, A^{T} B$.

## Required Arguments

A-Real NRA by NCA matrix. (Input)
$\boldsymbol{B}$ - Real NRB by NCB matrix. (Input)
$C$ - Real NCA by NCB matrix containing the transpose product $A^{T} B$. (Output)

## Optional Arguments

$\boldsymbol{N R A}$ - Number of rows in A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
$N C A$ - Number of columns in A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N R B}$ - Number of rows in B. (Input)
NRB must be the same as NRA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.
NCB - Number of columns in B. (Input)
Default: $\mathrm{NCB}=\operatorname{SIZE}(\mathrm{B}, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(\mathrm{B}, 1)$.
$N R C$ - Number of rows of C. (Input)
NRC must be equal to NCA.
Default: NRC $=\operatorname{SIZE}(C, 1)$.
NCC - Number of columns of C. (Input)
NCC must be equal to NCB.
Default: NCC $=\operatorname{SIZE}(C, 2)$.
$\boldsymbol{L D C}$ - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC = SIZE (C,1).

## FORTRAN 90 Interface

Generic: CALL MXTYF (A, B, C [,...])
Specific: The specific interface names are S_MXTYF and D_MXTYF.

## FORTRAN 77 Interface

Single: CALL MXTYF (NRA, NCA, A, LDA, NRB, NCB, B, LDB, NRC, NCC, C, LDC)

Double: The double precision name is DMXTYF.

## Description

The routine MXTYF computes the real general matrix $C=A^{T} B$ given the real rectangular matrices $A$ and $B$.

## Example

Multiply the transpose of a $3 \times 4$ real matrix by a $3 \times 3$ real matrix. The output matrix will be a $4 \times 3$ real matrix.

```
    USE MXTYF INT
    USE WRRRN_INT
    IMPLICIT NONE
    INTEGER NCA, NCB, NCC, NRA, NRB, NRC
    PARAMETER (NCA=4, NCB=3, NCC=3, NRA=3, NRB=3, NRC=4)
    REAL A(NRA,NCA), B(NRB,NCB), C(NRC,NCC)
        Set values for A
        A =( (1.0 0.0 2.0 0.0
            ( 3.0 4.0 -1.0 0.0 )
            (2.0 1.0 2.0 1.0)
            Set values for B
            B =( (-1.0 2.0 0.0
            ( 3.0 0.0 -1.0 )
            ( 0.0 5.0 2.0 )
    DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
        1.0/
    DATA B/-1.0, 3.0, 0.0, 2.0, 0.0, 5.0, 0.0, -1.0, 2.0/
    CALL MXTYF (A, B, C)
        Compute C = trans(A)*B
    CALL WRRRN ('C = trans(A)*B', C)
    END
```


## Output

$C=\operatorname{trans}(A) * B$

| 1 | 2 | 3 |
| ---: | ---: | ---: |
| 8.00 | 12.00 | 1.00 |
| 12.00 | 5.00 | -2.00 |
| -5.00 | 14.00 | 5.00 |
| 0.00 | 5.00 | 2.00 |

## MXYTF

Multiplies a matrix $A$ by the transpose of a matrix $B, A B^{T}$.

## Required Arguments

$\boldsymbol{A}$ - Real NRA by NCA rectangular matrix. (Input)
$\boldsymbol{B}-$ Real NRB by NCB rectangular matrix. (Input)
$C$ - Real NRC by NCC rectangular matrix containing the transpose product $A B^{T}$. (Output)

## Optional Arguments

$N R A$ - Number of rows in A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns in A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N R B}$ - Number of rows in B. (Input)
Default: NRB $=\operatorname{SIZE}(\mathrm{B}, 1)$.
$N C B$ - Number of columns in B. (Input)
NCB must be the same as NCA.
Default: NCB $=\operatorname{SIZE}(\mathrm{B}, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
$N R C$ - Number of rows of $C$. (Input)
NRC must be equal to NRA.
Default: NRC $=\operatorname{SIZE}(\mathrm{C}, 1)$.
NCC - Number of columns of C. (Input)
NCC must be equal to NRB.
Default: NCC $=\operatorname{SIZE}(C, 2)$.
$\boldsymbol{L D C}$ - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC $=\operatorname{SIZE}(C, 1)$.

## FORTRAN 90 Interface

Generic: CALL MXYTF (A, B, C $[, \ldots]$ )
Specific: The specific interface names are S_MXYTF and D_MXYTF.

## FORTRAN 77 Interface

Single: CALL MXYTF (NRA, NCA, A, LDA, NRB, NCB, B, LDB, NRC, NCC, C, LDC)

Double: $\quad$ The double precision name is DMXYTF.

## Description

The routine MXYTF computes the real general matrix $C=A B^{T}$ given the real rectangular matrices $A$ and $B$.

## Example

Multiply a $3 \times 4$ real matrix by the transpose of a $3 \times 4$ real matrix. The output matrix will be a $3 \times 3$ real matrix.

```
USE MXYTF INT
USE WRRRN_INT
IMPLICIT NONE
NNTEGER Declare variables
NTEGER NCA, NCB, NCC, NRA, NRB, NRC
PARAMETER (NCA=4, NCB=4, NCC=3, NRA=3, NRB=3, NRC=3)
REAL A(NRA,NCA), B(NRB,NCB), C(NRC,NCC)
    Set values for A
    A = ( 1.0 0.0 2.0 0.0 )
                                ( 3.0 4.0 -1.0 0.0 )
    (2.0 1.0 2.0 1.0)
    Set values for B
    B = ( -1.0 2.0 0.0 2.0 )
        ( 3.0 0.0 -1.0 -1.0 )
        (0.0 5.0 2.0 5.0)
DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
        1.0/
DATA B/-1.0, 3.0, 0.0, 2.0, 0.0, 5.0, 0.0, -1.0, 2.0, 2.0, -1.0, &
        5.0/
CALL MXYTF (A, B, C)
    Compute C = A*trans(B)
    Print results
CALL WRRRN ('C = A*trans(B)', C)
END
```


## Output

| $C=A * t r a n s(B)$ |  |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -1.00 | 1.00 | 4.00 |
| 5.00 | 10.00 | 18.00 |
| 2.00 | 3.00 | 14.00 |

## MRRRR

Multiplies two real rectangular matrices, $A B$.

## Required Arguments

$\boldsymbol{A}$ - Real NRA by NCA matrix in full storage mode. (Input)
$\boldsymbol{B}-$ Real NRB by NCB matrix in full storage mode. (Input)
$C$ - Real NRC by NCC matrix containing the product $A B$ in full storage mode. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.

NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N R B$ - Number of rows of B. (Input)
NRB must be equal to NCA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.
NCB - Number of columns of B. (Input)
Default: NCB $=\operatorname{SIZE}(\mathrm{B}, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
$N R C$ - Number of rows of C. (Input)
NRC must be equal to NRA.
Default: NRC $=\operatorname{SIZE}(C, 1)$.
NCC - Number of columns of C. (Input)
NCC must be equal to NCB.
Default: NCC $=\operatorname{SIZE}(C, 2)$.
LDC - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC $=\operatorname{SIZE}(C, 1)$.

## FORTRAN 90 Interface

Generic: CALL MRRRR (A, B, C [, ...])

Specific: The specific interface names are S_MRRRR and D_MRRRR.

## FORTRAN 77 Interface

Single: $\quad$ CALL MRRRR (NRA, NCA, A, LDA, NRB, NCB, B, LDB, NRC, NCC, C, LDC)

Double: The double precision name is DMRRRR.

## Description

Given the real rectangular matrices $A$ and $B, \operatorname{MRRRR}$ computes the real rectangular matrix $C=A B$.

## Example

Multiply a $3 \times 4$ real matrix by a $4 \times 3$ real matrix. The output matrix will be a $3 \times 3$ real matrix. USE MRRRR_INT

USE WRRRN_INT

```
IMPLICIT NONE
! Declare variables
    INTEGER NCA, NCB, NCC, NRA, NRB, NRC
    PARAMETER (NCA=4, NCB=3, NCC=3, NRA=3, NRB=4, NRC=3)
    REAL A(NRA,NCA), B (NRB,NCB), C (NRC,NCC)
    Set values for A
    A =( (1.0 0.0 2.0 0.0)
        ( 3.0 4.0 -1.0 0.0)
        ( 2.0 1.0 2.0 1.0 )
    Set values for B
    B =( -1.0 0.0 2.0 )
        ( 3.0 5.0 2.0 )
        ( 0.0 0.0 -1.0 )
        (2.0-1.0 5.0)
    DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
        1.0/
    DATA B/-1.0, 3.0, 0.0, 2.0, 0.0, 5.0, 0.0, -1.0, 2.0, 2.0, -1.0, &
        5.0/
    Compute C = A*B
    CALL MRRRR (A, B, C)
    CALL WRRRN ('C = A*B', C)
    END
```


## Output

|  | $C=A * B$ |  |
| ---: | ---: | ---: |
| 1 | 2 | 3 |
| -1.00 | 0.00 | 0.00 |
| 9.00 | 20.00 | 15.00 |
| 3.00 | 4.00 | 9.00 |

## MCRCR

Multiplies two complex rectangular matrices, $A B$.

## Required Arguments

$\boldsymbol{A}$ - Complex NRA by NCA rectangular matrix. (Input)
$\boldsymbol{B}$ - Complex NRB by NCB rectangular matrix. (Input)
$\boldsymbol{C}$ - Complex NRC by NCC rectangular matrix containing the product A * B. (Output)

## Optional Arguments

$N R A$ - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N R B$ - Number of rows of B. (Input)
NRB must be equal to NCA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.

NCB - Number of columns of B. (Input)
Default: NCB $=\operatorname{SIZE}(B, 2)$.

LDB - Leading dimension of B exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
NRC - Number of rows of C. (Input)
NRC must be equal to NRA.
Default: NRC $=\operatorname{SIZE}(C, 1)$.

NCC - Number of columns of C. (Input)
NCC must be equal to NCB.
Default: NCC = SIZE (C,2).
$\boldsymbol{L D C}$ - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC = SIZE (C,1).

## FORTRAN 90 Interface

Generic: CALL MCRCR (A, B, C [, ...])

Specific: The specific interface names are S_MCRCR and D_MCRCR.

## FORTRAN 77 Interface

Single: CALL MCRCR (NRA, NCA, A, LDA, NRB, NCB, B, LDB, NRC, NCC, C, LDC)

Double: The double precision name is DMCRCR.

## Description

Given the complex rectangular matrices $A$ and $B, \operatorname{MCRCR}$ computes the complex rectangular matrix $C=A B$.

## Example

Multiply a $3 \times 4$ complex matrix by a $4 \times 3$ complex matrix. The output matrix will be a $3 \times 3$ complex matrix.

```
USE MCRCR INT
USE WRCRN INT
IMPLICIT NONE
INTEGER NCA, NCB, NCC, NRA, NRB, NRC
PARAMETER (NCA=4, NCB=3, NCC=3, NRA=3, NRB=4, NRC=3)
COMPLEX A(NRA,NCA), B(NRB,NCB), C (NRC,NCC)
    A = ( 1.0 + 1.0i -1.0+ 2.0i 0.0 + 1.0i 0.0 - 2.0i
        ( 3.0 + 7.0i 6.0-4.0i 2.0-1.0i 0.0 + 1.0i )
        ( 1.0 + 0.0i 1.0-2.0i -2.0+ 0.0i 0.0 + 0.0i )
    Set values for B
    B = ( 2.0 + 1.0i 3.0 + 2.0i 3.0 + 1.0i )
        ( 2.0 - 1.0i 4.0 - 2.0i 5.0 - 3.0i )
        ( 1.0 + 0.0i 0.0 - 1.0i 0.0 + 1.0i )
        ( 2.0 + 1.0i 1.0 + 2.0i 0.0 - 1.0i )
DATA A/ (1.0,1.0), (3.0,7.0), (1.0,0.0), (-1.0,2.0), (6.0,-4.0), &
    (1.0,-2.0), (0.0,1.0), (2.0,-1.0), (-2.0,0.0), (0.0,-2.0), &
        (0.0,1.0), (0.0,0.0)/
DATA B/(2.0,1.0), (2.0,-1.0), (1.0,0.0), (2.0,1.0), (3.0,2.0), &
    (4.0,-2.0), (0.0,-1.0), (1.0,2.0), (3.0,1.0), (5.0,-3.0), &
    (0.0,1.0), (0.0,-1.0)/
                            Compute C = A*B
CALL MCRCR (A, B, C)
    Print results
CALL WRCRN ('C = A*B', C)
```


## Output



## HRRRR

Computes the Hadamard product of two real rectangular matrices.

## Required Arguments

A-Real NRA by NCA rectangular matrix. (Input)
B-Real NRB by NCB rectangular matrix. (Input)
C-Real NRC by NCC rectangular matrix containing the Hadamard product of A and B. (Output)
If $A$ is not needed, then $C$ can share the same storage locations as A. Similarly, if $B$ is not needed, then $C$ can share the same storage locations as $B$.

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$N R B$ - Number of rows of B. (Input)
NRB must be equal to NRA.
Default: NRB $=\operatorname{SIZE}(B, 1)$.
$N C B$ - Number of columns of B. (Input)
NCB must be equal to NCA.
Default: $\mathrm{NCB}=\operatorname{SIZE}(\mathrm{B}, 2)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.
$N R C$ - Number of rows of C. (Input)
NRC must be equal to NRA.
Default: NRC $=\operatorname{SIZE}(C, 1)$.
NCC - Number of columns of C. (Input)
NCC must be equal to NCA.
Default: NCC = SIZE (C,2).
$\mathbf{L D C}$ - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC $=\operatorname{SIZE}(C, 1)$.

## FORTRAN 90 Interface

Generic: CALL $\operatorname{HRRRR}(A, B, C[, \ldots])$

Specific: The specific interface names are S_HRRRR and D_HRRRR.

## FORTRAN 77 Interface

## Single: CALL HRRRR (NRA, NCA, A, LDA, NRB, NCB, B, LDB, NRC, NCC,

 C, LDC)Double: The double precision name is DHRRRR.

## Description

The routine HRRRR computes the Hadamard product of two real matrices $A$ and $B$ and returns a real matrix $C$, where $C_{i j}=A_{i j} B_{i j}$.

## Example

Compute the Hadamard product of two $4 \times 4$ real matrices. The output matrix will be a $4 \times 4$ real matrix.

```
USE HRRRR INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER NCA, NCB, NCC, NRA, NRB, NRC
PARAMETER (NCA=4, NCB=4, NCC=4, NRA=4, NRB=4, NRC=4)
REAL A(NRA,NCA), B(NRB,NCB), C(NRC,NCC)
    Set values for A
    A = ( -1.0 0.0 -3.0 8.0 )
    ( 2.0 1.0 7.0 2.0 )
    ( 3.0 -2.0 2.0 -6.0 )
    (4.0 1.0 -5.0 -8.0 )
```

```
!
    Set values for B
    B = ( 2.0 3.0 0.0 -10.0 )
    ( 1.0 -1.0 4.0 2.0 )
    ( -1.0 -2.0 7.0 1.0 )
    ( 2.0 1.0 9.0 0.0 )
DATA A/-1.0, 2.0, 3.0, 4.0, 0.0, 1.0, -2.0, 1.0, -3.0, 7.0, 2.0, &
    -5.0, 8.0, 2.0, -6.0, -8.0/
DATA B/2.0, 1.0, -1.0, 2.0, 3.0, -1.0, -2.0, 1.0, 0.0, 4.0, 7.0, &
    9.0, -10.0, 2.0, 1.0, 0.0/
        Compute Hadamard product of A and B
CALL HRRRR (A, B, C)
CALL WRRRN ('C = A (*) B', C)
END
```


## Output

|  | $C=A$ | (*) $B$ |  |
| ---: | ---: | ---: | ---: |
| 1 | 2 | 3 | 4 |
| -2.00 | 0.00 | 0.00 | -80.00 |
| 2.00 | -1.00 | 28.00 | 4.00 |
| -3.00 | 4.00 | 14.00 | -6.00 |
| 8.00 | 1.00 | -45.00 | 0.00 |

## BLINF

This function computes the bilinear form $x^{T} A y$.

## Function Return Value

BLINF - The value of $x^{T} A y$ is returned in BLINF. (Output)

## Required Arguments

$\boldsymbol{A}-$ Real NRA by NCA matrix. (Input)
$\boldsymbol{X}$ - Real vector of length NRA. (Input)
$\boldsymbol{Y}$ - Real vector of length NCA. (Input)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: $\operatorname{BLINF}(A, X, Y[, \ldots])$
Specific: The specific interface names are S_BLINF and D_BLINF.

## FORTRAN 77 Interface

Single: BLINF (NRA, NCA, A, LDA, X, Y)

Double: The double precision name is DBLINF.

## Description

Given the real rectangular matrix $A$ and two vectors $x$ and $y$, BLINF computes the bilinear form $x^{T} A y$.

## Comments

The quadratic form can be computed by calling BLINF with the vector X in place of the vector Y .

## Example

Compute the bilinear form $x^{T} A y$, where $x$ is a vector of length $5, A$ is a $5 \times 2$ matrix and $y$ is a vector of length 2 .

```
USE BLINF INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NCA, NRA
PARAMETER (NCA=2, NRA=5)
    INTEGER NOUT
REAL A(NRA,NCA), VALUE, X(NRA), Y (NCA)
    Set values for A
    A =( -2.0 2.0 )
        ( 3.0 -6.0)
        (-4.0 7.0 )
        ( 1.0 -8.0)
        (0.0 10.0 )
        Set values for X
    X = ( 1.0 -2.0 3.0 -4.0 -5.0 )
    Set values for Y
    Y =( -6.0 3.0 )
```

!
!

```
DATA A/-2.0, 3.0, -4.0, 1.0, 0.0, 2.0, -6.0, 7.0, -8.0, 10.0/
DATA X/1.0, -2.0, 3.0, -4.0, -5.0/
DATA Y/-6.0, 3.0/
VALUE = BLINF (A,X,Y)
    Compute bilinear form
                            Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The bilinear form trans(x)*A*y = ', VALUE
END
```


## Output

The bilinear form trans(x)*A*y $=195.000$

## POLRG

Evaluates a real general matrix polynomial.

## Required Arguments

A -N by N matrix for which the polynomial is to be computed. (Input)

COEF - Vector of length NCOEF containing the coefficients of the polynomial in order of increasing power. (Input)
$\boldsymbol{B}-\mathrm{N}$ by N matrix containing the value of the polynomial evaluated at A. (Output)

## Optional Arguments

$N$ - Order of the matrix A. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 1)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
NCOEF - Number of coefficients. (Input)
Default: NCOEF = SIZE (COEF,1).
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

## FORTRAN 90 Interface

Generic: CALL POLRG (A, COEF, B [, ...])

Specific: The specific interface names are S_POLRG and D_POLRG.

## FORTRAN 77 Interface

Single: CALL POLRG (N, A, LDA, NCOEF, COEF, B, LDB)
Double: The double precision name is DPOLRG.

## Description

Let $m=$ NCOEF and $c=$ COEF.
The routine POLRG computes the matrix polynomial

$$
B=\sum_{k=1}^{m} c_{k} A^{k-1}
$$

using Horner's scheme

$$
B=\left(\ldots\left(\left(c_{m} A+c_{m-1} I\right) A+c_{m-2} I\right) A+\ldots+c_{1} I\right)
$$

where $I$ is the $N \times N$ identity matrix.

## Comments

Workspace may be explicitly provided, if desired, by use of P2LRG/DP2LRG. The reference is
CALL P2LRG (N, A, LDA, NCOEF, COEF, B, LDB, WORK)
The additional argument is
WORK - Work vector of length $\mathrm{N} * \mathrm{~N}$.

## Example

This example evaluates the matrix polynomial $3 I+A+2 A^{2}$, where $A$ is a $3 \times 3$ matrix.

```
USE POLRG INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, N, NCOEF
PARAMETER (N=3, NCOEF=3, LDA=N, LDB=N)
REAL A(LDA,N), B(LDB,N), COEF (NCOEF)
    Set values of A and COEF
    A = (\begin{array}{lll}{1.0}&{3.0}&{2.0}\end{array})
    ( -5.0 1.0 7.0 )
    ( 1.0 5.0 -4.0 )
    COEF = (3.0, 1.0, 2.0)
DATA A/1.0, -5.0, 1.0, 3.0, 1.0, 5.0, 2.0, 7.0, -4.0/
```

```
    DATA COEF/3.0, 1.0, 2.0/
```

$!$

```
Evaluate B = 3I + A + 2*A**2
```

CALL POLRG (A, COEF, B)
! Print B
CALL WRRRN ('B = $3 I+A+2 * A * * 2$ ', $B)$
END

## Output

```
    B = 3I + A + 2*A**2
```

    \(\begin{array}{rrr}1 & 2 & 3 \\ -20.0 & 35.0 & 32.0\end{array}\)
    \(-11.0 \quad 46.0 \quad-55.0\)
    \(-55.0 \quad-19.0 \quad 105.0\)
    
## MURRV

Multiplies a real rectangular matrix by a vector.

## Required Arguments

A-Real NRA by NCA rectangular matrix. (Input)
$\boldsymbol{X}$ - Real vector of length NX. (Input)
$\boldsymbol{Y}$ - Real vector of length NY containing the product $A * X$ if IPATH is equal to 1 and the product $\operatorname{trans}(\mathrm{A}) * \mathrm{X}$ if IPATH is equal to 2 . (Output)

## Optional Arguments

$N R A$ - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
NX must be equal to NCA if IPATH is equal to 1 . NX must be equal to NRA if IPATH is equal to 2 .
Default: NX $=\operatorname{SIZE}(X, 1)$.
IPATH - Integer flag. (Input) IPATH $=1$ means the product $Y=\mathrm{A} * \mathrm{X}$ is computed. IPATH $=2$ means the product
$Y=\operatorname{trans}(A) * X$ is computed, where $\operatorname{trans}(A)$ is the $\operatorname{transpose}$ of $A$.
Default: IPATH $=1$.
$N Y$ - Length of the vector Y. (Input)
NY must be equal to NRA if IPATH is equal to 1 . NY must be equal to NCA if IPATH is equal to 2 .
Default: NY $=\operatorname{SIZE}(Y, 1)$.

## FORTRAN 90 Interface

Generic: CALL MURRV (A, X, Y [, ...])
Specific: The specific interface names are S_MURRV and D_MURRV.

## FORTRAN 77 Interface

Single: CALL MURRV (NRA, NCA, A, LDA, NX, X, IPATH, NY, Y)
Double: The double precision name is DMURRV.

## Description

If IPATH $=1$, MURRV computes $y=A x$, where $A$ is a real general matrix and $x$ and $y$ are real vectors. If IPATH $=2$, MURRV computes $y=A^{T} x$.

## Example

Multiply a $3 \times 3$ real matrix by a real vector of length 3 . The output vector will be a real vector of length 3.

```
USE MURRV INT
USE WRRRN_INT
```

IMPLICIT NONE

```
!
INTEGER LDA, NCA, NRA, NX, NY
PARAMETER (NCA=3, NRA=3, NX=3, NY=3)
!
    INTEGER IPATH
    REAL A(NRA,NCA), X(NX), Y(NY)
                            Set values for A and X
                            A =( (1.0 0.0 2.0 )
                            ( 0.0 3.0 0.0 )
                            (4.0 1.0 2.0 )
                            X =( 1.0 2.0 1.0 )
DATA A/1.0, 0.0, 4.0, 0.0, 3.0, 1.0, 2.0, 0.0, 2.0/
DATA X/1.0, 2.0, 1.0/
    Compute y = Ax
```

```
    IPATH = 1
    CALL MURRV (A, X, Y)
! Print results
    CALL WRRRN ('y = Ax', Y, 1, NY, 1)
    END
```

Output
$y=A x$
$\begin{array}{rrr}1 & 2 & 3 \\ 3.000 & 6.000^{3} & 8.000^{2}\end{array}$

## MURBV

Multiplies a real band matrix in band storage mode by a real vector.

## Required Arguments

A-Real NLCA + NUCA +1 by n band matrix stored in band mode. (Input)
NLCA - Number of lower codiagonals in A. (Input)
NUCA - Number of upper codiagonals in A. (Input)
$\boldsymbol{X}$ - Real vector of length NX. (Input)
$\boldsymbol{Y}$ - Real vector of length NY containing the product $A * X$ if IPATH is equal to 1 and the product $\operatorname{trans(A)*} \mathrm{X}$ if IPATH is equal to 2 . (Output)

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
NX must be equal to N .
Default: NX $=\operatorname{SIZE}(X, 1)$.
IPATH - Integer flag. (Input)
IPATH $=1$ means the product $\mathrm{Y}=\mathrm{A} * \mathrm{X}$ is computed. IPATH $=2$ means the product $Y=\operatorname{trans}(A) * X$ is computed, where $\operatorname{trans}(A)$ is the transpose of $A$.
Default: IPATH $=1$.
$N Y$ - Length of vector Y. (Input)
NY must be equal to $N$.
Default: NY $=\operatorname{SIZE}(\mathrm{Y}, 1)$.

## FORTRAN 90 Interface

Generic: CALL MURBV (A, NLCA, NUCA, X, Y [, ...])

Specific: The specific interface names are $S_{-} M U R B V$ and $D_{-} M U R B V$.

## FORTRAN 77 Interface

Single: CALL MURBV (N, A, LDA, NLCA, NUCA, NX, X, IPATH, NY, Y)

Double: The double precision name is DMURBV.

## Description

If IPATH $=1$, MURBV computes $y=A x$, where $A$ is a real band matrix and $x$ and $y$ are real vectors. If IPATH $=2$, MURBV computes $y=A^{T} x$.

## Example

Multiply a real band matrix of order 6 , with two upper codiagonals and two lower codiagonals stored in band mode, by a real vector of length 6 . The output vector will be a real vector of length 6.

```
USE MURBV INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, N, NLCA, NUCA, NX, NY
PARAMETER (LDA=5, N=6, NLCA=2, NUCA=2,NX=6,NY=6)
INTEGER IPATH
REAL A(LDA,N), X(NX), Y(NY)
    Set values for A (in band mode)
    A =( ( 0.0 0.0 1.0 2.0 3.0 4.0 )
                                ( 0.0 1.0 2.0 3.0 4.0 5.0 )
                                ( 1.0 2.0 3.0 4.0 5.0 6.0 )
        (-1.0 -2.0 -3.0 -4.0 -5.0 0.0 )
        (-5.0 -6.0 -7.0 -8.0 0.0 0.0 )
    Set values for X
    X = (-1.0 2.0 -3.0 4.0 -5.0 6.0 )
DATA A/0.0, 0.0, 1.0, -1.0, -5.0, 0.0, 1.0, 2.0, -2.0, -6.0, &
    1.0, 2.0, 3.0, -3.0, -7.0, 2.0, 3.0, 4.0, -4.0, -8.0, 3.0, &
    4.0, 5.0, -5.0, 0.0, 4.0, 5.0, 6.0, 0.0, 0.0/
DATA X/-1.0, 2.0, -3.0, 4.0, -5.0, 6.0/
```

$!$
Compute $y=A x$
IPATH = 1
CALL MURBV (A, NLCA, NUCA, X, Y)
!
Print results
CALL WRRRN ('y = Ax', Y, 1, NY, 1)
END

## Output

$\begin{array}{rrrrrrr}1 & 2 & 3^{y}=A x \\ -2.00 & 7.00 & -11.00 & 17.00 & 10.00 & 29.00\end{array}$

## MUCRV

Multiplies a complex rectangular matrix by a complex vector.

## Required Arguments

$\boldsymbol{A}$ - Complex NRA by NCA rectangular matrix. (Input)
$\boldsymbol{X}$ - Complex vector of length NX. (Input)
$\boldsymbol{Y}$ - Complex vector of length NY containing the product $A * X$ if IPATH is equal to 1 and the product $\operatorname{trans}(\mathrm{A}) * \mathrm{X}$ if IPATH is equal to 2 . (Output)

## Optional Arguments

NRA - Number of rows of A. (Input) Default: NRA $=\operatorname{SIZE}(A, 1)$.
$N C A$ - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
NX must be equal to NCA if IPATH is equal to 1 . NX must be equal to NRA if IPATH is equal to 2 .
Default: NX $=\operatorname{SIZE}(X, 1)$.
IPATH - Integer flag. (Input)
IPATH $=1$ means the product $\mathrm{Y}=\mathrm{A} * \mathrm{X}$ is computed. IPATH $=2$ means the product $Y=\operatorname{trans}(A) * X$ is computed, where $\operatorname{trans}(A)$ is the transpose of $A$.
Default: IPATH $=1$.
$N Y$ - Length of the vector Y. (Input)
NY must be equal to NRA if IPATH is equal to 1 . NY must be equal to NCA if IPATH is equal to 2 .
Default: $\mathrm{NY}=\operatorname{SIZE}(\mathrm{Y}, 1)$.

## FORTRAN 90 Interface

Generic: CALL MUCRV (A, X, Y [,...])

Specific: The specific interface names are S_MUCRV and D_MUCRV.

## FORTRAN 77 Interface

Single: CALL mUCRV (NRA, NCA, A, LDA, NX, X, IPATH, NY, Y)
Double: The double precision name is DMUCRV.

## Description

If IPATH $=1$, MUCRV computes $y=A x$, where $A$ is a complex general matrix and $x$ and $y$ are complex vectors. If IPATH $=2$, MUCRV computes $y=A^{T} x$.

## Example

Multiply a $3 \times 3$ complex matrix by a complex vector of length 3 . The output vector will be a complex vector of length 3 .

```
USE MUCRV INT
USE WRCRN_INT
IMPLICIT NONE
INTEGER NCA, NRA, NX, NY
PARAMETER (NCA=3, NRA=3, NX=3, NY=3)
    INTEGER IPATH
    COMPLEX A(NRA,NCA), X(NX), Y(NY)
    A = ( 1.0 + 2.0i 3.0 + 4.0i 1.0 + 0.0i )
        ( 2.0 + 1.0i 3.0 + 2.0i 0.0 - 1.0i )
        (2.0-1.0i 1.0 + 0.0i 0.0 + 1.0i)
        X = ( 1.0 - 1.0i 2.0 - 2.0i 0.0 - 1.0i )
    DATA A/ (1.0,2.0), (2.0,1.0), (2.0,-1.0), (3.0,4.0), (3.0,2.0), &
        (1.0,0.0), (1.0,0.0), (0.0,-1.0), (0.0,1.0)/
    DATA X/(1.0,-1.0), (2.0,-2.0), (0.0,-1.0)/
        Compute y = Ax
    IPATH = 1
CALL MUCRV (A, X, Y)
```

!
! Print results
CALL WRCRN ('y = Ax', Y, 1, NY, 1) END

## Output

$y=A x$
$(17.00,2.00)^{1} \quad(12.00,-3.00)^{2} \quad(4.00,-5.00)^{3}$

## MUCBV

Multiplies a complex band matrix in band storage mode by a complex vector.

## Required Arguments

A-Complex NLCA + NUCA +1 by n band matrix stored in band mode. (Input)
NLCA - Number of lower codiagonals in A. (Input)
NUCA - Number of upper codiagonals in A. (Input)
$\boldsymbol{X}$ - Complex vector of length NX. (Input)
$\boldsymbol{Y}$ - Complex vector of length NY containing the product $A * X$ if IPATH is equal to 1 and the product $\operatorname{trans}(A) * X$ if IPATH is equal to 2 . (Output)

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{N X}$ - Length of the vector X . (Input)
$N X$ must be equal to $N$.
Default: NX $=\operatorname{SIZE}(X, 1)$.
IPATH - Integer flag. (Input)
IPATH $=1$ means the product $\mathrm{Y}=\mathrm{A} * \mathrm{X}$ is computed. IPATH $=2$ means the product
$Y=\operatorname{trans}(A) * X$ is computed, where $\operatorname{trans}(A)$ is the transpose of $A$.
Default: IPATH $=1$.
$N Y$ - Length of vector Y. (Input)
$N Y$ must be equal to $N$.
Default: NY = SIZE (Y,1).

## FORTRAN 90 Interface

Generic: CALL MUCBV (A, NLCA, NUCA, X, Y [, ...])

Specific: The specific interface names are S_MUCBV and D_MUCBV.

## FORTRAN 77 Interface

Single: CALL mUCBV (N, A, LDA, NLCA, NUCA, NX, X, IPATH, NY, Y)
Double: The double precision name is DMUCBV.

## Description

If IPATH $=1$, MUCBV computes $y=A x$, where $A$ is a complex band matrix and $x$ and $y$ are complex vectors. If IPATH $=2$, MUCBV computes $y=A^{T} x$.

## Example

Multiply the transpose of a complex band matrix of order 4, with one upper codiagonal and two lower codiagonals stored in band mode, by a complex vector of length 3 . The output vector will be a complex vector of length 3 .

```
USE MUCBV INT
USE WRCRN__INT
IMPLICIT NONE
Declare variables
INTEGER LDA, N, NLCA, NUCA, NX, NY
PARAMETER (LDA=4, N=4, NLCA=2, NUCA=1, NX=4, NY=4)
INTEGER IPATH
COMPLEX A(LDA,N), X(NX), Y(NY)
                                    Set values for A (in band mode)
        A = ( 0.0+ 0.0i 1.0+ 2.0i 3.0+ 4.0i 5.0+ 6.0i )
            ( -1.0-1.0i -1.0-1.0i -1.0-1.0i -1.0-1.0i )
            ( -1.0+ 2.0i -1.0+ 3.0i -2.0+ 1.0i 0.0+ 0.0i )
            ( 2.0+ 0.0i 0.0+ 2.0i 0.0+ 0.0i 0.0+ 0.0i )
                                    Set values for X
        X = ( 3.0 + 4.0i 0.0 + 0.0i 1.0 + 2.0i -2.0 - 1.0i )
DATA A/ (0.0,0.0), (-1.0,-1.0), (-1.0,2.0), (2.0,0.0), (1.0,2.0), &
    (-1.0,-1.0), (-1.0,3.0), (0.0,2.0), (3.0,4.0), (-1.0,-1.0), &
    (-2.0,1.0), (0.0,0.0), (5.0,6.0), (-1.0,-1.0), (0.0,0.0), &
    (0.0,0.0)/
DATA X/(3.0,4.0), (0.0,0.0),(1.0,2.0), (-2.0,-1.0)/
                            Compute y = Ax
IPATH = 2
CALL MUCBV (A, NLCA, NUCA, X, Y, IPATH=IPATH)
                                    Print results
CALL WRCRN ('Y = Ax', Y, 1, NY, 1)
```


## Output



## ARBRB

Adds two band matrices, both in band storage mode.

## Required Arguments

$\boldsymbol{A}-\mathrm{n}$ by n band matrix with NLCA lower codiagonals and NUCA upper codiagonals stored in band mode with dimension (NLCA + NUCA +1 ) by N . (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
$B-\mathrm{N}$ by n band matrix with NLCB lower codiagonals and NUCB upper codiagonals stored in band mode with dimension (NLCB $+\mathrm{NUCB}+1$ ) by N . (Input)

NLCB - Number of lower codiagonals of B. (Input)
NUCB - Number of upper codiagonals of B. (Input)
$\boldsymbol{C}-\mathrm{N}$ by N band matrix with NLCC lower codiagonals and NUCC upper codiagonals containing the sum $A+B$ in band mode with dimension $(N L C C+N U C C+1)$ by $N$. (Output)

NLCC - Number of lower codiagonals of C. (Input)
NLCC must be at least as large as max(NLCA, NLCB).
NUCC - Number of upper codiagonals of C. (Input)
NUCC must be at least as large as max(NUCA, NUCB).

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A, B and C. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.
$\boldsymbol{L D B}$ - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=\operatorname{SIZE}(B, 1)$.

LDC - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC = SIZE $(C, 1)$.

## FORTRAN 90 Interface

Generic: CALL ARBRB (A, NLCA, NUCA, B, NLCB, NUCB, C, NLCC, NUCC [ , ...])
Specific: The specific interface names are S_ARBRB and D_ARBRB.

## FORTRAN 77 Interface

```
Single: CALL ARBRB (N, A, LDA, NLCA, NUCA, B, LDB, NLCB, NUCB, C,
    LDC, NLCC, NUCC)
```

Double: The double precision name is DARBRB.

## Description

The routine ARBRB adds two real matrices stored in band mode, returning a real matrix stored in band mode.

## Example

Add two real matrices of order 4 stored in band mode. Matrix $A$ has one upper codiagonal and one lower codiagonal. Matrix $B$ has no upper codiagonals and two lower codiagonals. The output matrix $C$, has one upper codiagonal and two lower codiagonals.

```
USE ARBRB_INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER LDA, LDB, LDC, N, NLCA, NLCB, NLCC, NUCA, NUCB, NUCC
PARAMETER (LDA=3, LDB=3, LDC=4, N=4, NLCA=1, NLCB=2, NLCC=2, &
    NUCA=1, NUCB=0, NUCC=1)
REAL A (LDA,N), B(LDB,N), C(LDC,N)
    Set values for A (in band mode)
```



```
        (1.0 1.0 1.0 1.0)
        ( 0.0 3.0 4.0 0.0)
    Set values for B (in band mode)
    B = (\begin{array}{llll}{3.0 3.0 3.0 3.0)}\end{array}]
        ( 1.0 -2.0 1.0 0.0)
        (\begin{array}{llll}{-1.0}&{2.0}&{0.0}&{0.0.0}\end{array})
```

$!$

```
DATA A/0.0, 1.0, 0.0, 2.0, 1.0, 3.0, 3.0, 1.0, 4.0, -1.0, 1.0, &
    0.0/
DATA B/3.0, 1.0, -1.0, 3.0, -2.0, 2.0, 3.0, 1.0, 0.0, 3.0, 0.0, &
        0.0/
            Add A and B to obtain C (in band
                                    mode)
CALL ARBRB (A, NLCA, NUCA, B, NLCB, NUCB, C, NLCC, NUCC)
                    Print results
CALL WRRRN ('C = A+B', C)
END
```


## Output

|  | $C=A+B$ |  |  |  |
| ---: | ---: | ---: | ---: | ---: |
|  | 1 | 2 | 3 | 4 |
| 1 | 0.000 | 2.000 | 3.000 | -1.000 |
| 2 | 4.000 | 4.000 | 4.000 | 4.000 |
| 3 | 1.000 | 1.000 | 5.000 | 0.000 |
| 4 | -1.000 | 2.000 | 0.000 | 0.000 |

## ACBCB

Adds two complex band matrices, both in band storage mode.

## Required Arguments

A - N by N complex band matrix with NLCA lower codiagonals and NUCA upper codiagonals stored in band mode with dimension (NLCA $+\mathrm{NUCA}+1$ ) by N . (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
$B-\mathrm{N}$ by N complex band matrix with NLCB lower codiagonals and NUCB upper codiagonals stored in band mode with dimension (NLCB $+\mathrm{NUCB}+1$ ) by N . (Input)

NLCB - Number of lower codiagonals of B. (Input)
NUCB - Number of upper codiagonals of B. (Input)
C - n by n complex band matrix with NLCC lower codiagonals and NUCC upper codiagonals containing the sum $A+B$ in band mode with dimension ( $\mathrm{NLCC}+\mathrm{NUCC}+1$ ) by N . (Output)

NLCC - Number of lower codiagonals of C. (Input) NLCC must be at least as large as max(NLCA, NLCB).

NUCC - Number of upper codiagonals of C. (Input) NUCC must be at least as large as max(NUCA, NUCB).

## Optional Arguments

$\boldsymbol{N}$ - Order of the matrices A, B and C. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(\mathrm{A}, 1)$.
LDB - Leading dimension of $B$ exactly as specified in the dimension statement of the calling program. (Input)
Default: LDB $=$ SIZE $(B, 1)$.
LDC - Leading dimension of C exactly as specified in the dimension statement of the calling program. (Input)
Default: LDC $=\operatorname{SIZE}(C, 1)$.

## FORTRAN 90 Interface

Generic: CALL ACBCB (A, NLCA, NUCA, B, NLCB, NUCB, C, NLCC, NUCC [, ...])
Specific: $\quad$ The specific interface names are S_ACBCB and D_ACBCB.

## FORTRAN 77 Interface

Single: CALL ACBCB (N, A, LDA, NLCA, NUCA, B, LDB, NLCB, NUCB, C, LDC, NLCC, NUCC)

Double: The double precision name is DACBCB.

## Description

The routine $A C B C B$ adds two complex matrices stored in band mode, returning a complex matrix stored in band mode.

## Example

Add two complex matrices of order 4 stored in band mode. Matrix $A$ has two upper codiagonals and no lower codiagonals. Matrix $B$ has no upper codiagonals and two lower codiagonals. The output matrix $C$ has two upper codiagonals and two lower codiagonals.

```
USE ACBCB INT
USE WRCRN__INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, LDB, LDC, N, NLCA, NLCB, NLCC, NUCA, NUCB, NUCC
PARAMETER (LDA=3, LDB=3, LDC=5, N=3, NLCA=0, NLCB=2, NLCC=2, &
    NUCA=2, NUCB=0, NUCC=2)
```

```
    COMPLEX
        A(LDA,N), B(LDB,N), C(LDC,N)
                            Set values for A (in band mode)
        A = ( 0.0 + 0.0i 0.0 + 0.0i 3.0 - 2.0i )
        ( 0.0 + 0.0i -1.0+ 3.0i 6.0 + 0.0i)
    ( 1.0 + 4.0i 5.0 - 2.0i 3.0 + 1.0i )
        Set values for B (in band mode)
        B = ( 3.0 + 1.0i 4.0 + 1.0i 7.0 - 1.0i )
            ( -1.0-4.0i 9.0 + 3.0i 0.0 + 0.0i )
    (2.0-1.0i 0.0 + 0.0i 0.0 + 0.0i )
    DATA A/ (0.0,0.0), (0.0,0.0), (1.0,4.0), (0.0,0.0), (-1.0,3.0), &
    (5.0,-2.0), (3.0,-2.0), (6.0,0.0), (3.0,1.0)/
    DATA B/ (3.0,1.0), (-1.0,-4.0), (2.0,-1.0), (4.0,1.0), (9.0,3.0), &
        (0.0,0.0), (7.0,-1.0), (0.0,0.0), (0.0,0.0)/
        Compute C = A+B
    CALL ACBCB (A, NLCA, NUCA, B, NLCB, NUCB, C, NLCC, NUCC)
        Print results
    CALL WRCRN ('C = A+B', C)
    END
```


## Output



## NRIRR

Computes the infinity norm of a real matrix.

## Required Arguments

A - Real NRA by NCA matrix whose infinity norm is to be computed. (Input)

ANORM - Real scalar containing the infinity norm of A. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.

NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(\mathrm{A}, 2)$.
$\boldsymbol{L D A}$ - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL NRIRR (A, ANORM [, ...])
Specific: The specific interface names are S_NRIRR and D_NRIRR.

## FORTRAN 77 Interface

Single: CALL NRIRR (NRA, NCA, A, LDA, ANORM)
Double: The double precision name is DNRIRR.

## Description

The routine NRIRR computes the infinity norm of a real rectangular matrix $A$. If $m=$ NRA and $n=$ NCA, then the $\infty$-norm of $A$ is

$$
\|A\|_{\infty}=\max _{1 \leq i \leq m} \sum_{j=1}^{n}\left|A_{i j}\right|
$$

This is the maximum of the sums of the absolute values of the row elements.

## Example

Compute the infinity norm of a $3 \times 4$ real rectangular matrix.

```
USE NRIRR_INT
USE UMACH_INT
IMPLICIT NONE
    INTEGER NCA, NRA
    PARAMETER (NCA=4, NRA=3)
    INTEGER NOUT
    REAL A(NRA,NCA), ANORM
    Set values for A
        A =( ( 1.0 0.0 2.0 0.0 )
                                ( 3.0 4.0 -1.0 0.0)
                            (2.0 1.0 2.0 1.0 )
    DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
        1.0/
    CALL NRIRR (A, ANORM)
    Print results
```

$!$
!

```
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The infinity norm of A is ', ANORM
END
```


## Output

The infinity norm of $A$ is 8.00000

## NR1RR

Computes the 1-norm of a real matrix.

## Required Arguments

$\boldsymbol{A}$ - Real NRA by NCA matrix whose 1-norm is to be computed. (Input)

ANORM - Real scalar containing the 1-norm of A. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.

NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL NR1RR (A, ANORM [, ...])

Specific: The specific interface names are S_NR1RR and D_NR1RR.

## FORTRAN 77 Interface

Single: CALL NR1RR (NRA, NCA, A, LDA, ANORM)

Double: The double precision name is DNR1RR.

## Description

The routine NR1RR computes the 1 -norm of a real rectangular matrix $A$. If $m=$ NRA and $n=\mathrm{NCA}$, then the 1 -norm of $A$ is

$$
\|A\|_{1}=\max _{1 \leq j \leq n} \sum_{i=1}^{m}\left|A_{i j}\right|
$$

This is the maximum of the sums of the absolute values of the column elements.

## Example

Compute the 1 -norm of a $3 \times 4$ real rectangular matrix.

```
USE NR1RR_INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER NCA, NRA
PARAMETER (NCA=4, NRA=3)
INTEGER NOUT
REAL A (NRA,NCA), ANORM
            Set values for A
                A =( (1.0 0.0 2.0 0.0 )
                    ( 3.0 4.0 -1.0 0.0)
                    (2.0 1.0 2.0 1.0 )
DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
    1.0/
CALL NR1RR (A, ANORM) Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The 1-norm of A is ', ANORM
END
```


## Output

The 1 -norm of A is 6.00000

## NR2RR

Computes the Frobenius norm of a real rectangular matrix.

## Required Arguments

$\boldsymbol{A}-$ Real NRA by NCA rectangular matrix. (Input)
ANORM - Frobenius norm of A. (Output)

## Optional Arguments

NRA - Number of rows of A. (Input)
Default: NRA $=\operatorname{SIZE}(A, 1)$.
NCA - Number of columns of A. (Input)
Default: NCA $=\operatorname{SIZE}(A, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL NR2RR (A, ANORM [, ...])

Specific: The specific interface names are S_NR2RR and D_NR2RR.

## FORTRAN 77 Interface

Single: CALL NR2RR (NRA, NCA, A, LDA, ANORM)
Double: The double precision name is DNR2RR.

## Description

The routine NR2RR computes the Frobenius norm of a real rectangular matrix $A$. If $m=$ NRA and $n=$ NCA, then the Frobenius norm of $A$ is

$$
\|A\|_{2}=\left[\sum_{i=1}^{m} \sum_{j=1}^{n} A_{i j}^{2}\right]^{1 / 2}
$$

## Example

Compute the Frobenius norm of a $3 \times 4$ real rectangular matrix.

```
USE NR2RR INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, NCA, NRA
PARAMETER (LDA=3, NCA=4, NRA=3)
INTEGER NOUT
REAL A(LDA,NCA), ANORM
Set values for A
A =( (1.0 0.0 2.0 0.0
    ( 3.0 4.0 -1.0 0.0 )
```

!

```
! ( 2.0 1.0 2.0 1.0 )
DATA A/1.0, 3.0, 2.0, 0.0, 4.0, 1.0, 2.0, -1.0, 2.0, 0.0, 0.0, &
        1.0/
CALL NR2RR (A, ANORM)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The Frobenius norm of A is ', ANORM
END
```


## Output

The Frobenius norm of A is 6.40312

## NR1RB

Computes the 1-norm of a real band matrix in band storage mode.

## Required Arguments

$\boldsymbol{A}-$ Real (NUCA $+\mathrm{NLCA}+1$ ) by N array containing the N by N band matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)
NUCA - Number of upper codiagonals of A. (Input)
ANORM - Real scalar containing the 1-norm of A. (Output)

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.
LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL NR1RB (A, NLCA, NUCA, ANORM [, ...])
Specific: The specific interface names are S_NR1RB and D_NR1RB.

## FORTRAN 77 Interface

Single: CALL NR1RB (N, A, LDA, NLCA, NUCA, ANORM)

Double: The double precision name is DNR1RB.

## Description

The routine NR1RB computes the 1-norm of a real band matrix $A$. The 1-norm of a matrix $A$ is

$$
\|A\|_{1}=\max _{1 \leq j \leq N} \sum_{i=1}^{N}\left|A_{i j}\right|
$$

This is the maximum of the sums of the absolute values of the column elements.

## Example

Compute the 1 -norm of a $4 \times 4$ real band matrix stored in band mode.

```
USE NR1RB_INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER LDA, N, NLCA, NUCA
PARAMETER (LDA=4, N=4, NLCA=2, NUCA=1)
    INTEGER NOUT
    REAL A(LDA,N), ANORM
    Set values for A (in band mode)
            A =( ( 0.0 2.0 2.0 3.0 )
            ( -2.0 -3.0 -4.0 -1.0)
            ( 2.0 1.0 0.0 0.0 )
            (0.0 1.0 0.0 0.0)
DATA A/0.0, -2.0, 2.0, 0.0, 2.0, -3.0, 1.0, 1.0, 2.0, -4.0, 0.0, &
        0.0, 3.0, -1.0, 2*0.0/
            Compute the L1 norm of A
CALL NR1RB (A, NLCA, NUCA, ANORM)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The 1-norm of A is ', ANORM
END
```

!

## Output

The 1 -norm of A is $\quad 7.00000$

## NR1CB

Computes the 1-norm of a complex band matrix in band storage mode.

## Required Arguments

$\boldsymbol{A}$ - Complex (NUCA $+\mathrm{NLCA}+1$ ) by N array containing the N by N band matrix in band storage mode. (Input)

NLCA - Number of lower codiagonals of A. (Input)
$N U C A$ - Number of upper codiagonals of A. (Input)

ANORM - Real scalar containing the 1-norm of A. (Output)

## Optional Arguments

$N$ - Order of the matrix. (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{A}, 2)$.

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA $=\operatorname{SIZE}(A, 1)$.

## FORTRAN 90 Interface

Generic: CALL NR1CB (A, NLCA, NUCA, ANORM [, ...])

Specific: The specific interface names are S_NR1CB and D_NR1CB.

## FORTRAN 77 Interface

Single: CALL NR1CB (N, A, LDA, NLCA, NUCA, ANORM)
Double: The double precision name is DNR1CB.

## Description

The routine NR1CB computes the 1-norm of a complex band matrix $A$. The 1-norm of a complex matrix $A$ is

$$
\|A\|_{1}=\max _{1 \leq j \leq N} \sum_{i=1}^{N}\left[\left|\Re A_{i j}\right|+\left|\mathfrak{\Im} A_{i j}\right|\right]
$$

## Example

Compute the 1-norm of a complex matrix of order 4 in band storage mode.

```
USE NR1CB INT
USE UMACH-
IMPLICIT NONE
! Declare variables
```

```
    INTEGER LDA, N, NLCA, NUCA
    PARAMETER (LDA=4, N=4, NLCA=2, NUCA=1)
!
    INTEGER NOUT
    REAL ANORM
    COMPLEX A(LDA,N)
                A = ( 0.0+0.0i 2.0+3.0i -1.0+1.0i -2.0-1.0i )
                    ( -2.0+3.0i 1.0+0.0i -4.0-1.0i 0.0-4.0i )
                    ( 2.0+2.0i 4.0+6.0i 3.0+2.0i 0.0+0.0i )
                    ( 0.0-1.0i 2.0+1.0i 0.0+0.0i 0.0+0.0i )
DATA A/ (0.0,0.0), (-2.0,3.0), (2.0,2.0), (0.0,-1.0), (2.0,3.0), &
    (1.0,0.0), (4.0,6.0), (2.0,1.0), (-1.0,1.0), (-4.0,-1.0), &
    (3.0,2.0), (0.0,0.0), (-2.0,-1.0), (0.0,-4.0), (0.0,0.0), &
    (0.0,0.0)/
            Compute the L1 norm of A
CALL NR1CB (A, NLCA, NUCA, ANORM)
    Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The 1-norm of A is ', ANORM
END
```


## Output

The 1 -norm of $A$ is 19.0000

## DISL2

This function computes the Euclidean (2-norm) distance between two points.

## Function Return Value

DISL2 - Euclidean (2-norm) distance between the points x and Y . (Output)

## Required Arguments

$\boldsymbol{X}$ — Vector of length max(N * |incx|, 1). (Input)
$\boldsymbol{Y}-$ Vector of length max(N * |incy|, 1). (Input)

## Optional Arguments

$N$ - Length of the vectors x and y . (Input)
Default: $\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)$.
INCX - Displacement between elements of x . (Input)
The $I$-th element of $X$ is $X(1+(I-1)$ * INCX) if INCX is greater than or equal to zero
or $\mathrm{X}(1+(\mathrm{I}-\mathrm{N})$ * INCX) if INCX is less than zero.
Default: $I N C X=1$.

INCY - Displacement between elements of Y. (Input)
The I-th element of $Y$ is $Y(1+(I-1)$ * INCY) if INCY is greater than or equal to zero or $Y(1+(I-N)$ * INCY) if INCY is less than zero.
Default: $\operatorname{INCY}=1$.

## FORTRAN 90 Interface

Generic: DISL2 (X, Y [, ...])
Specific: The specific interface names are S_DISL2 and D_DISL2.

## FORTRAN 77 Interface

Single: DISL2 (N, X, INCX, Y, INCY)
Double: The double precision function name is DDISL2.

## Description

The function DISL2 computes the Euclidean (2-norm) distance between two points $x$ and $y$. The Euclidean distance is defined to be

$$
\left[\sum_{i=1}^{N}\left(x_{i}-y_{i}\right)^{2}\right]^{1 / 2}
$$

## Example

Compute the Euclidean (2-norm) distance between two vectors of length 4.

```
USE DISL2 INT
USE UMACH_INT
IMPLICIT NONE
! INTEGER INCX, INCY, N Declare variables
INTEGER INCX, INCY, N
PARAMETER (N=4)
INTEGER NOUT
REAL VAL, X(N), Y(N)
```

```
Set values for X and Y
```

Set values for X and Y
X = ( 1.0 -1.0 0.0 2.0 )
X = ( 1.0 -1.0 0.0 2.0 )
Y =( 4.0 2.0 1.0 -3.0 )
Y =( 4.0 2.0 1.0 -3.0 )
DATA X/1.0, -1.0, 0.0, 2.0/
DATA Y/4.0, 2.0, 1.0, -3.0/

```
!
\(!\)
\(\mathrm{VAL}=\mathrm{DISL} 2(\mathrm{X}, \mathrm{Y})\)
!
```

Print results

```
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The 2-norm distance is ', VAL
END

\section*{Output}

The 2-norm distance is 6.63325

\section*{DISL1}

This function computes the 1-norm distance between two points.

\section*{Function Return Value}

DISL1 - 1-norm distance between the points X and Y . (Output)

\section*{Required Arguments}
\(\boldsymbol{X}-\) Vector of length \(\max (\mathrm{N} * \mid\) IncX \(\mid, 1)\). (Input)
\(\boldsymbol{Y}-\) Vector of length max( \(\mathrm{N} *|\operatorname{INCY}|, 1\) ). (Input)

\section*{Optional Arguments}
\(N\) - Length of the vectors X and Y . (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)\).
\(\boldsymbol{I N C X}\) - Displacement between elements of X . (Input)
The I-th element of \(X\) is \(X(1+(I-1) *\) INCX \()\) if INCX is greater than or equal to zero or \(\mathrm{X}(1+(\mathrm{I}-\mathrm{N})\) * INCX \()\) if INCX is less than zero.
Default: \(\operatorname{INCX}=1\).
\(\boldsymbol{I N C Y}\) - Displacement between elements of Y . (Input)
The I-th element of \(Y\) is \(Y(1+(I-1)\) * INCY) if INCY is greater than or equal to zero or \(Y(1+(I-N)\) * INCY) if INCY is less than zero.
Default: \(\operatorname{INCY}=1\).

\section*{FORTRAN 90 Interface}

Generic: DISL1 (X, Y [, ...])

Specific: The specific interface names are S_DISL1 and D_DISL1.

\section*{FORTRAN 77 Interface}

Single: DISL1 (N, X, INCX, Y, INCY)
Double: The double precision function name is DDISL1.

\section*{Description}

The function DISL1 computes the 1-norm distance between two points \(x\) and \(y\). The 1-norm distance is defined to be
\[
\sum_{i=1}^{N}\left|x_{i}-y_{i}\right|
\]

\section*{Example}

Compute the 1-norm distance between two vectors of length 4.
```

USE DISL1 INT
USE UMACH_
IMPLICIT NONE
! Declare variables
INTEGER INCX, INCY, N
PARAMETER (N=4)
INTEGER NOUT
REAL VAL, X(N), Y(N)
Set values for X and Y
X =( 1.0 -1.0 0.0 2.0 )
Y =( 4.0 2.0 1.0 -3.0 )
DATA X/1.0, -1.0, 0.0, 2.0/
DATA Y/4.0, 2.0, 1.0, -3.0/
VAL = DISL1(X,Y)
! Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The 1-norm distance is ', VAL
END

```
!

\section*{Output}

The 1 -norm distance is 12.0000

\section*{DISLI}

This function computes the infinity norm distance between two points.

\section*{Function Return Value}

DISLI - Infinity norm distance between the points X and Y . (Output)

\section*{Required Arguments}
\(\boldsymbol{X}-\mathrm{Vector}\) of length \(\max (\mathrm{N} *|\operatorname{INCX}|, 1)\). (Input)
\(\boldsymbol{Y}-\) Vector of length \(\max (\mathrm{N} *|\operatorname{INCY}|, 1)\). (Input)

\section*{Optional Arguments}
\(N\) - Length of the vectors X and Y . (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{X}, 1)\).

INCX - Displacement between elements of X . (Input)
The \(I\)-th element of \(X\) is \(X(1+(I-1) * I N C X)\) if INCX is greater than or equal to zero or \(X(1+(I-N)\) * INCX) if INCX is less than zero.
Default: \(\operatorname{INCX}=1\).

INCY - Displacement between elements of Y. (Input)
The I-th element of \(Y\) is \(Y(1+(I-1)\) * INCY) if INCY is greater than or equal to zero or \(\mathrm{Y}(1+(\mathrm{I}-\mathrm{N})\) * INCY) if INCY is less than zero. Default: \(\operatorname{INCY}=1\).

\section*{FORTRAN 90 Interface}

Generic: DISLI (X, Y [, ...])

Specific: The specific interface names are S_DISLI and D_DISLI.

\section*{FORTRAN 77 Interface}

Single: DISLI (N, X, INCX, Y, INCY)
Double: The double precision function function name is DDISLI.

\section*{Description}

The function DISLI computes the \(\infty\)-norm distance between two points \(x\) and \(y\). The \(\infty\)-norm distance is defined to be
\[
\max _{1 \leq i \leq N}\left|x_{i}-y_{i}\right|
\]

\section*{Example}

Compute the \(\infty\)-norm distance between two vectors of length 4 .
```

USE DISLI_INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER INCX, INCY, N
PARAMETER (N=4)
!
INTEGER NOUT
REAL VAL, X(N), Y(N)
Set values for X and Y
X =( 1.0 -1.0 0.0 2.0)
Y =( 4.0 2.0 1.0 -3.0 )
DATA X/1.0, -1.0, 0.0, 2.0/
DATA Y/4.0, 2.0, 1.0, -3.0/
VAL = DISLI(X,Y)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) ' The infinity-norm distance is ', VAL
END

```

\section*{Output}

The infinity-norm distance is 5.00000

\section*{VCONR}

Computes the convolution of two real vectors.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Vector of length NX . (Input)
\(\mathbf{Y}\) - Vector of length NY. (Input)
\(\mathbf{Z}\) - Vector of length NZ containing the convolution \(\mathrm{Z}=\mathrm{x}\) * Y . (Output)

\section*{Optional Arguments}
\(N X\) - Length of the vector x . (Input)
Default: NX = SIZE (X, 1 ).
\(N Y\) - Length of the vector Y. (Input)
Default: NY \(=\operatorname{SIZE}(\mathrm{Y}, 1)\).
\(N Z\) - Length of the vector \(Z\). (Input)
NZ must be at least NX + NY - 1 .
Default: NZ \(=\operatorname{SIZE}(Z, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL VCONR (X, Y, Z [,...])

Specific: The specific interface names are S_VCONR and D_VCONR.

\section*{FORTRAN 77 Interface}

Single: CALL VCONR (NX, X, NY, Y, NZ, Z)
Double: The double precision name is DVCONR.

\section*{Description}

The routine VCONR computes the convolution \(z\) of two real vectors \(x\) and \(y\). Let \(n_{x}=\mathrm{NX}, n_{y}=\mathrm{NY}\) and \(n_{z}=\) NZ. The vector \(z\) is defined to be
\[
z_{j}=\sum_{k=1}^{n_{x}} x_{j-k+1} y_{k} \quad \text { for } j=1,2, \ldots, n_{z}
\]
where \(n_{z}=n_{x}+n_{y}-1\). If the index \(j-k+1\) is outside the range \(1,2, \ldots, n_{x}\), then \(x_{j-k+1}\) is taken to be zero.

The fast Fourier transform is used to compute the convolution. Define the complex vector \(u\) of length \(n_{z}=n_{x}+n_{y}-1\) to be
\[
u=\left(x_{1}, x_{2}, \ldots, x_{n_{x}}, 0, \ldots, 0\right)
\]

The complex vector \(v\), also of length \(n_{z}\), is defined similarly using \(y\). Then, by the Fourier convolution theorem,
\[
\hat{w}_{i}=\hat{u}_{i} \hat{v}_{i} \quad \text { for } i=1,2, \ldots, n_{z}
\]
where the \(\hat{u}\) indicates the Fourier transform of \(u\) computed via IMSL routines FFTCF and FFTCB (see Chapter 6, Transforms) is used to compute the complex vector \(w\) from \(\hat{w}\). The vector \(z\) is then found by taking the real part of the vector \(w\).

\section*{Comments}

Workspace may be explicitly provided, if desired, by use of \(\mathrm{V} 2 \mathrm{ONR} / \mathrm{DV} 2 \mathrm{ONR}\). The reference is
CALL V2ONR (NX, X, NY, Y, NZ, Z, XWK, YWK, ZWK, WK)
The additional arguments are as follows:
\(\boldsymbol{X W K}\) - Complex work array of length NX + NY - 1 .
\(\boldsymbol{Y} W K\) - Complex work array of length \(N X+N Y-1\).
\(\mathbf{Z W K}\) - Complex work array of length NX \(+\mathrm{NY}-1\).
\(W K-\) Real work array of length 6 * \((N X+N Y-1)+15\).

\section*{Example}

In this example, the convolution of a vector \(x\) of length 8 and a vector \(y\) of length 3 is computed. The resulting vector \(z\) is of length \(8+3-1=10\). (The vector \(y\) is sometimes called a filter.)
```

USE VCONR INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER NX, NY, NZ
PARAMETER (NX=8, NY=3, NZ=NX+NY-1)
REAL X(NX), Y(NY), Z(NZ)
X=((1.0 2.0
Set values for Y
Y =(0.0 0.0 1.0)
DATA X/1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0/
DATA Y/0.0, 0.0, 1.0/
CALL VCONR (X,Y,Z)
Compute vector convolution
Z = X * Y
Print results
CALL WRRRN ('Z = X (*) Y', Z, 1, NZ, 1)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 1 & 2 & 3 & 4 & 5 & \[
\begin{array}{rr}
x & (*) \\
6
\end{array}
\] & 7 & 8 & 9 & 10 \\
\hline 0.000 & 0.000 & 1.000 & 2.000 & 3.000 & 4.000 & 5.000 & 6.000 & 7.000 & 8.000 \\
\hline
\end{tabular}

\section*{VCONC}

Computes the convolution of two complex vectors.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Complex vector of length NX. (Input)
\(\boldsymbol{Y}\) - Complex vector of length NY. (Input)
\(\mathbf{Z}\) - Complex vector of length NZ containing the convolution \(\mathrm{Z}=\mathrm{X} * \mathrm{Y}\). (Output)

\section*{Optional Arguments}
\(\boldsymbol{N X}\) - Length of the vector X . (Input)
Default: NX \(=\operatorname{SIZE}(X, 1)\).
\(N Y\) - Length of the vector \(Y\). (Input)
Default: NY = SIZE (Y,1).
\(N Z\) - Length of the vector \(Z\). (Input)
NZ must be at least NX + NY-1.
Default: NZ \(=\operatorname{SIZE}(\mathrm{Z}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL VCONC (X, Y, Z [,...])

Specific: The specific interface names are \(S_{-} V C O N C\) and \(D\) _VCONC.

\section*{FORTRAN 77 Interface}

Single: CALL VCONC (NX, X, NY, Y, NZ, Z)
Double: The double precision name is DVCONC.

\section*{Description}

The routine VCONC computes the convolution \(z\) of two complex vectors \(x\) and \(y\). Let \(n_{x}=\mathrm{NX}\), then \(n_{y}=N Y\) and \(n_{z}=N Z\). The vector \(z\) is defined to be
\[
z_{j}=\sum_{k=1}^{n_{k}} x_{j-k+1} y_{k} \quad \text { for } j=1,2, \ldots, n_{z}
\]
where \(n_{z}=n_{x}+n_{y}-1\). If the index \(j-k+1\) is outside the range \(1,2, \ldots, n_{x}\), then \(x_{j-k+1}\) is taken to be zero.

The fast Fourier transform is used to compute the convolution. Define the complex vector \(u\) of length \(n_{z}=n_{x}+n_{y}-1\) to be
\[
u=\left(x_{1}, x_{2}, \ldots, x_{n_{z}}, 0, \ldots, 0\right)
\]

The complex vector \(v\), also of length \(n_{z}\), is defined similarly using \(y\). Then, by the Fourier convolution theorem,
\[
\hat{z}_{i}=\hat{u}_{i} \hat{v}_{i} \quad \text { for } i=1,2, \ldots, n_{z}
\]
where the \(\hat{u}\) indicates the Fourier transform of \(u\) computed using IMSL routine FFTCF (see Chapter 6, Transforms). The complex vector \(z\) is computed from \(\hat{w}\) via IMSL routine FFTCB (see Chapter 6, Transforms).

\section*{Comments}

Workspace may be explicitly provided, if desired, by use of \(\mathrm{V} 2 \mathrm{ONC} / \mathrm{DV} 20 \mathrm{NC}\). The reference is CALL V2ONC (NX, X, NY, Y, NZ, Z, XWK, YWK, WK)

The additional arguments are as follows:
\(\boldsymbol{X} \mathbf{W K}\) - Complex work array of length NX + NY - 1 .
YWK - Complex work array of length NX \(+\mathrm{NY}-1\).
\(W K-\) Real work arrary of length 6 * \((N X+N Y-1)+15\).

\section*{Example}

In this example, the convolution of a vector \(x\) of length 4 and a vector \(y\) of length 3 is computed. The resulting vector \(z\) is of length \(4+3-y\) is sometimes called a filter.)
```

USE VCONC_INT
USE WRCRN INT
IMPLICIT NONE
INTEGER NX, NY, NZ
PARAMETER (NX=4,NY=3, NZ=NX+NY-1)
COMPLEX X(NX), Y(NY), Z(NZ)
Set values for X
X = ( 1.0+2.0i 3.0+4.0i 5.0+6.0i 7.0+8.0i )
Set values for Y
Y = (0.0+0i 0.0+0i 1.0+0i )
DATA X/(1.0,2.0), (3.0,4.0), (5.0,6.0), (7.0,8.0)/
DATA Y/(0.0,0.0), (0.0,0.0), (1.0,1.0)/
Compute vector convolution
Z = X * Y
CALL VCONC (X,Y,Z)
CALL WRCRN ('Z = X (*) Y', Z, 1, NZ, 1)
END

```

\section*{Output}


\section*{Extended Precision Arithmetic}

This section describes a set of routines for mixed precision arithmetic. The routines are designed to allow the computation and use of the full quadruple precision result from the multiplication of
two double precision numbers. An array called the accumulator stores the result of this multiplication. The result of the multiplication is added to the current contents of the accumulator. It is also possible to add a double precision number to the accumulator or to store a double precision approximation in the accumulator.
The mixed double precision arithmetic routines are described below. The accumulator array, \(Q A C C\), is a double precision array of length 2 . Double precision variables are denoted by DA and DB. Available operations are:
Initialize a real accumulator, \(\mathrm{QACC} \leftarrow \mathrm{DA}\).
CALL DQINI (DA, QACC)
Store a real accumulator, \(D A \leftarrow Q A C C\).
CALL DQSTO (QACC, DA)
Add to a real accumulator, QACC \(\leftarrow Q A C C+D A\).
CALL DQADD (DA, QACC)

CALL DQMUL (DA, DB, QACC)
There are also mixed double complex arithmetic versions of the above routines. The accumulator, ZACC, is a double precision array of length 4 . Double complex variables are denoted by ZA and ZB. Available operations are:

Initialize a complex accumulator, ZACC \(\leftarrow\) ZA.
CALL ZQINI (ZA, ZACC)
Store a complex accumulator, ZA \(\leftarrow\) ZACC.
CALL ZQSTO (ZACC, ZA)
Add to a complex accumulator, ZACC \(\leftarrow\) ZACC + ZA.
CALL ZQADD (ZA, ZACC)
Add a product to a complex accumulator, \(\mathrm{ZACC} \leftarrow \mathrm{ZACC}+\mathrm{ZA} * \mathrm{ZB}\).
CALL ZQMUL (ZA, ZB, ZACC)

\section*{Example}

In this example, the value of \(1.0 \mathrm{D} 0 / 3.0 \mathrm{D} 0\) is computed in quadruple precision using Newton's method. Four iterations of
\[
x_{k+1}=x_{k}+\left(x_{k}-a x_{k}^{2}\right)
\]
with \(a=3\) are taken. The error \(a x-1\) is then computed. The results are accurate to approximately twice the usual double precision accuracy, as given by the IMSL routine \(\operatorname{DMACH}(4)\), in the Reference Material section of this manual. Since DMACH is machine dependent, the actual accuracy obtained is also machine dependent.
```

    USE IMSL_LIBRARIES
    IMPLICIT NONE
    INTEGER I, NOUT
    DOUBLE PRECISION A, DACC(2), DMACH, ERROR, SACC(2), X(2), X1, X2, EPSQ
    !
CALL UMACH (2, NOUT)
A = 3.0D0
CALL DQINI (1.0001D0/A, X)
Compute X(K+1) = X(K) - A* X(K)*X(K)
+ X(K)
DO 10 I=1, 4
X1 = X(1)
X2 = X(2)
CALL DQADD (X1, X)
CALL DQADD (X2, X)
CALL DQINI (O.ODO, DACC)
CALL DQMUL (X1, X1, DACC)
CALL DQMUL (X1, X2, DACC)
CALL DQMUL (X1, X2, DACC)
CALL DQMUL (X2, X2, DACC)
Compute -A* (X*X)
CALL DQINI (O.ODO, SACC)
CALL DQMUL (-A, DACC (1), SACC)
CALL DQMUL (-A, DACC(2), SACC)
Compute -A* (X*X) + (X + X)
CALL DQADD (SACC (1), X)
CALL DQADD (SACC (2), X)
CONTINUE
CALL DQINI (0.0D0, SACC)
CALL DQMUL (A, X(1), SACC)
CALL DQMUL (A, X(2), SACC)
CALL DQADD (-1.0D0, SACC)
CALL DQSTO (SACC, ERROR)
EPSQ = AMACH(4)
EPSQ = EPSQ * EPSQ
WRITE (NOUT,99999) ERROR, ERROR/EPSQ
!
99999 FORMAT (' A*X - 1 = ', D15.7, ' = ', F10.5, '*MACHEPS**2')
END

```

\section*{Output}
```

A*X - 1 = 0.6162976D-32 = 0.12500*MACHEPS**2

```

\section*{Chapter 10: Linear Algebra Operators and Generic Functions}

\section*{Routines}
1.2. Operators
Computes matrix-matrix or matrix-vector product ..... 1537
Computes transpose matrix-matrix product. ..... 1541
Computes matrix- transpose matrix product ..... 1544
Computes conjugate transpose matrix-matrix product ..... 1547
Computes matrix-conjugate transpose matrix product ..... 1550
Computes the transpose of a matrix ..... 1553
Computes conjugate transpose of a matrix ..... 1556
Computes the inverse matrix ..... 1558
Computes inverse matrix-matrix product ..... 1561
Computes matrix-inverse matrix product ..... 1571
10.2 Functions
Computes the Cholesky factorization of a positive-definite, symmetric or self-adjoint matrix ..... 1574
Computes the condition number of a matrix ..... COND ..... 1577
Computes the determinant of a rectangular matrix ..... 1581
Constructs a square diagonal matrix ..... DIAG ..... 1584
Extracts the diagonal terms of a matrix DIAGONALS ..... 1585
Computes the eigenvalue-eigenvector decomposition of an ordinary or generalized eigenvalue problem EIG ..... 1586
Creates the identity matrix ..... EYE ..... 1590
Computes the Discrete Fourier Transform of one complex sequence. ..... FFT ..... 1592
Discrete Fourier Transform of several complex or real sequences FFT BOX ..... 1594
Computes the inverse of the Discrete Fourier Transform of one complex sequence IFFT ..... 1596
Computes the inverse Discrete Fourier Transform of several complex or real sequences ..... IFFT_BOX ..... 1598
Tests for NaN isNaN ..... 1600
Returns the value for NaN ..... NaN ..... 1601
Computes the norm of an array ..... NORM ..... 1602
Orthogonalizes the columns of a matrix ORTH ..... 1605
Generates random numbers ..... RAND ..... 1608
Computes the mathematical rank of a matrix. ..... RANK ..... 1610
Computes the singular value decomposition of a matrix ..... SVD ..... 1612
Normalizes the columns of a matrix. ..... 1614

\section*{Usage Notes}

This chapter describes numerical linear algebra, Fourier transforms, random number generation, and other utility software packaged as defined operations that are executed with a function notation similar to standard mathematics. The resulting interface alters the way libraries are presented to the user. Many computations of numerical linear algebra are documented here as operators and generic functions. A notation is developed reminiscent of matrix algebra. This allows the Fortran user to express mathematical formulas in terms of operators. The operators can be used with both dense and sparse matrices.

A comprehensive Fortran module, linear_operators, defines the operators and functions. Its use provides this simplification. Subroutine calls and the use of type-dependent procedure names are largely avoided. This makes a rapid development cycle possible, at least for the purposes of experiments and proof-of-concept. The goal is to provide the Fortran programmer with an interface, operators, and functions that are useful and succinct. The modules can be used with or added to existing Fortran programs, but the operators provide a more readable program whenever they apply. This approach may require more hidden working storage. The size of the executable program may be larger than alternatives using subroutines. There are applications wherein the operator and function interface does not have the functionality that is available using subroutine libraries. To retain greater flexibility, some users will continue to require the techniques of calling subroutines.

A parallel computation for many of the defined operators and functions has been implemented. The type of problem solved is a simple one: several independent problems of the same data type and size. Most of the detailed communication for parallel computation is hidden from the user. Those functions having this data type computed in parallel are marked in bold type. The section "Dense Matrix Parallelism Using MPI" gives an introduction on how users should write their codes to use machines on a network.

A number of examples, in addition to those shown in this document, are supplied in the product examples directory. The name of the example code is shown in parentheses in the example heading, for those examples that are included with the product.

\section*{Matrix Optional Data Changes}

To reset tolerances for determining singularity and to allow for other data changes, non-allocated "hidden" variables are defined within the modules. These variables can be allocated first, then assigned values which result in the use of different tolerances or greater efficiency in the executable program. The non-allocated variables, whose scope is limited to the module, are hidden from the casual user. Default values or rules are applied if these arrays are not allocated. In more detail, the inverse matrix operator ".i." applied to a square matrix first uses the \(L U\) factorization code LIN_SOL_GEN and row pivoting. The default value for a small diagonal term is defined to be:
```

sqrt(epsilon(A))*sum(abs(A))/(n*n+1)

```

If the system is singular, a generalized matrix inverse is computed with the \(Q R\) factorization code LIN_SOL_LSQ using this same tolerance. Both row and column pivoting are used. If the system is singular, an error message will be printed and a Fortran 90 STOP is executed. Users may want to change this rule. This is illustrated by continuing and not printing the error message. The following is a additional source to accomplish this, for all following invocations of the operator ".i.":
```

allocate(s_inv_options(1))
s_inv_options (1) = skip_error_processing
B = .i. A

```

There are additional self-documenting integer parameters, packaged in the module linear_operators, that allow users other choices, such as changing the value of the tolerance, as noted above. Included is the ability to have the option apply for just the next invocation of the operator. Options are available that allow optional data to be passed to supporting Fortran 90 subroutines. This is illustrated in the following example:

Operator_ex36.f90
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 4 for LIN GEIG GEN (using operators).
integer, parameter :: n=32
real(kind(1d0)), parameter :: one=1d0, zero=0d0
real(kind(ldO)) a(n,n), b(n,n), bta(n), err
complex(kind(1d0)) alpha(n), v(n,n)
! Generate random matrices for both A and B.
A = rand (A); B = rand (B)
! Set the option, a larger tolerance than default for lin_sol_lsq.
allocate(d_eig_options(6))
d_eig_options(1) = options_for_lin_geig_gen
d_eig_options(2) = 4
d_eig_options(3) = d_lin_geig_gen_for_lin_sol_lsq
d_eig_options(4) = 2
d_eig_options(5) = d_options(d_lin_sol_lsq_set_small,\&
sqrt(epsilon(one))*norm(B,1))
d_eig_options(6) = d_lin_sol_lsq_no_sing_mess
! Compute the generalized eigenvalues.
alpha = EIG(A, B=B, D=bta, W=V)
! Check the residuals.
err = norm((A .x. V .x. diag(bta)) - (B .x. V .x. diag(alpha)),1)/\&
(norm(A,1)*norm(bta,1)+norm(B,1)*norm(alpha,1))

```
```

    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 4 for LIN_GEIG_GEN (operators) is correct.'
    end if
    ! Clean up the allocated array. This is good housekeeping.
deallocate(d_eig_options)
end

```

Note that in this example one first allocates the array by which the user will pass the new options for EIG to use. This array is named d_eig_options in accordance with the name of the unallocated option array specified in the documentation for EIG. A size of 6 is specified because a total of six options must be passed to EIG to accomplish the resetting of the singular value tolerance and to turn off the printing of the error message when the matrix is singular. The first entry of d_eig_options specifies which of the options for EIG will be set. The next entry designates the number of entries which follows that apply to "options_for_lin_geig_gen". The third entry specifies the option value of LIN_GEIG_GEN to be set, d_lin_geig_gen_for_lin_sol_lsq. The fourth entry specifies the number of entries that follow which apply to LIN_SOL_LSQ. Finally, the fifth and sixth entries set the two LIN_SOL_LSQ options that we desire.

\section*{Dense Matrix Computations}

\section*{, M PI}

For a detailed description of MPI Capability see "Dense Matrix Parallelism Using MPI."

This section is concerned with methods for computing with dense matrices. Consider a Fortran 90 code fragment that solves a linear system of algebraic equations, \(A y=b\), then computes the residual \(r=b-A y\). A standard mathematical notation is often used to write the solution,
\[
y=A^{-1} b
\]

A user thinks: "matrix and right-hand side yields solution." The code shows the computation of this mathematical solution using a defined Fortran operator ".ix.", and random data obtained with the function, rand. This operator is read "inverse matrix times." The residuals are computed with another defined Fortran operator ". x.", read "matrix times vector." Once a user understands the equivalence of a mathematical formula with the corresponding Fortran operator, it is possible to write this program with little effort. The last line of the example before end is discussed below.
```

USE linear_operators
integer,parameter :: n=3; real A(n,n), y(n), b(n), r(n)
A=rand(A); b=rand(b); y = A .ix. b
r = b - (A .x. y ) ! Parentheses are needed

```
end

The IMSL Fortran Numerical Library provides additional lower-level software that implements the operation ".ix.", the function rand, matrix multiply ".x.", and others not used in this
example. Standard matrix products and inverse operations of matrix algebra are shown in the following table:
\begin{tabular}{|c|c|c|}
\hline Defined Array Operation & Matrix Operation & Alternative in Fortran 90 \\
\hline A . x . B & \(A B\) & matmul (A, B) \\
\hline .i. A & \(A^{-1}\) & \[
\begin{aligned}
& \text { LIN_SOL_GEN } \\
& \text { LIN_SOL_LSQ }
\end{aligned}
\] \\
\hline .t. A, .h. A & \(A^{T}, A^{H}\) & \begin{tabular}{l}
transpose (A) \\
conjg (transpose (A))
\end{tabular} \\
\hline A .ix. B & \(A^{-1} B\) & \[
\begin{aligned}
& \text { LIN_SOL_GEN } \\
& \text { LIN_SOL_LSQ }
\end{aligned}
\] \\
\hline B .xi. A & \(B A^{-1}\) & \[
\begin{aligned}
& \text { LIN_SOL_GEN } \\
& \text { LIN_SOL_LSQ }
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& A \text {.tx. B, or (.t. A) .x. B } \\
& \text { A.hx. B, or (.h. A) .x. B }
\end{aligned}
\] & \(A^{T} B, A^{H} B\) & \begin{tabular}{l}
matmul(transpose (A), B) \\
matmul (conjg(transpose (A)), B)
\end{tabular} \\
\hline \[
\begin{array}{|l}
\hline \text { B .xt. A, or B .x. } \\
\text { B . (.t. A) } \\
\text { Bh. or B .x. } \\
\text { (.h. A) }
\end{array}
\] & \(B A^{T}, B A^{H}\) & \[
\begin{aligned}
& \text { matmul(B, transpose(A)) } \\
& \text { matmul(B, conjg(transpose (A))) }
\end{aligned}
\] \\
\hline
\end{tabular}

The IMSL operators apply generically to all standard precisions and floating-point data types real and complex - and to objects that are broader in scope than arrays with a fixed number of dimensions. For example, the matrix product ". x." applies to matrix times vector and matrix times matrix represented as Fortran 90 arrays. It also applies to "independent matrix products." For this, use the notion: a box of problems to refer to independent linear algebra computations, of the same kind and dimension, but different data. The racks of the box are the distinct problems. In terms of Fortran 90 arrays, a rank-3, assumed-shape array is the data structure used for a box. The first two dimensions are the data for a matrix; the third dimension is the rack number. Each problem is independent of other problems in consecutive racks of the box. We use parallelism of an underlying network of processors, and MPI, when computing these disjoint problems.

In addition to the operators .ix., .xi., .i., and .x., additional operators .t., .h., .tx., .hx., .xt., and . xh. are provided for complex matrices. Since the transpose matrix is defined for complex matrices, this meaning is kept for the defined operations. In order to write one defined operation for both real and complex matrices, use the conjugate-transpose in all cases. This will result in only real operations when the data arrays are real.

For sums and differences of vectors and matrices, the intrinsic array operations " + " and " - " are available. It is not necessary to have separate defined operations. A parsing rule in Fortran 90 states that the result of a defined operation involving two quantities has a lower precedence than any intrinsic operation. This explains the parentheses around the next-to-last line containing the sub-expression "A.x. y" found in the example. Users are advised to always include parentheses around array expressions that are mixed with defined operations, or whenever there is possible confusion without them. The next-to-last line of the example results in computing the residual associated with the solution, namely \(r=b-A y\). Ideally, this residual is zero when the system has a unique solution. It will be computed as a non-zero vector due to rounding errors and conditioning of the problem.

\section*{Dense Matrix Functions}

For a detailed description of MPI Capability see "Dense Matrix Parallelism Using MPI."

Several decompositions and functions required for numerical linear algebra follow. The convention of enclosing optional quantities in brackets, "[ ]" is used. The functions that use MPI for parallel execution of the box data type are marked in bold.
\begin{tabular}{|c|c|}
\hline Defined Array Functions & Matrix Operation \\
\hline \(\mathrm{S}=\mathrm{SVD}(\mathrm{A}[, \mathrm{U}=\mathrm{U}, \mathrm{V}=\mathrm{V}])\) & \(A=U S V^{T}\) \\
\hline \[
\begin{aligned}
& \mathrm{E}=\mathrm{EIG}(\mathrm{~A}[[, \mathrm{~B}=\mathrm{B}, \mathrm{D}=\mathrm{D}], \\
& \mathrm{V}=\mathrm{V}, \mathrm{~W}=\mathrm{W}])
\end{aligned}
\] & \[
\begin{aligned}
& (A V=V E), A V D=B V E \\
& (A W=W E), A W D=B W E
\end{aligned}
\] \\
\hline \(\mathrm{R}=\mathrm{CHOL}(\mathrm{A})\) & \(A=R^{T} R\) \\
\hline \(\mathrm{Q}=0 \mathrm{RTH}(\mathrm{A}[, \mathrm{R}=\mathrm{R}]\) ) & \((A=Q R), Q^{T} Q=I\) \\
\hline U=UNIT ( A ) & \(\left[u_{1}, \ldots\right]=\left[a_{1} /\left\|a_{1}\right\|, \ldots\right]\) \\
\hline \(\mathrm{F}=\mathrm{DET}\) ( A ) & \(\operatorname{Det}(A)=\) determinant \\
\hline K=RANK ( A ) & \(\operatorname{rank}(A)=\operatorname{rank}\) \\
\hline P=NORM(A[, [type=]i]) & \[
\begin{aligned}
& p=\|A\|_{1}=\max _{j}\left(\sum_{i=1}^{m} \mid a_{i j}\right) \\
& p=\|A\|_{2}=s_{1}=\text { largest singular value } \\
& p=\|A\|_{\infty \leftrightarrow h u s e(1)}=\max _{i}\left(\sum_{j=1}^{n}\left|a_{i j}\right|\right)
\end{aligned}
\] \\
\hline C=COND ( A ) & \(s_{1} / s_{\text {rank }(A)}\) \\
\hline \(\mathrm{Z}=\mathrm{EYE}(\mathrm{N})\) & \(Z=I_{N}\) \\
\hline A=DIAG (X) & \(A=\operatorname{diag}\left(x_{1}, \ldots\right)\) \\
\hline X=DIAGONALS (A) & \(x=\left(a_{11}, \ldots\right)\) \\
\hline \[
\begin{aligned}
& \mathrm{Y}=\operatorname{FFT}(\mathrm{X},[\mathrm{WORK}=\mathrm{W}]) ; \\
& \mathrm{X}=\operatorname{IFFT}(\mathrm{Y},[\mathrm{WORK}=\mathrm{W}])
\end{aligned}
\] & Discrete Fourier Transform, Inverse \\
\hline \[
\begin{aligned}
& \mathrm{Y}=\mathrm{FFT}_{1} \mathrm{BOX}(\mathrm{X},[\mathrm{WORK}=\mathrm{W}]) ; \\
& \mathrm{X}=\mathrm{IFFT} \text { _BOX(Y, [WORK=W] })
\end{aligned}
\] & Discrete Fourier Transform for Boxes, Inverse \\
\hline A=RAND (A) & Random numbers, \(0<A<1\) \\
\hline \(\mathrm{L}=\) isNaN (A) & Test for NaN, if (l) then... \\
\hline
\end{tabular}

In certain functions, the optional arguments are inputs while other optional arguments are outputs. To illustrate the example of the box SVD function, a code is given that computes the singular value decomposition and the reconstruction of the random matrix box, \(A\), using the computed factors, \(R=U S V^{T}\). Mathematically \(R=A\), but this will be true, only approximately, due to rounding errors. The value units_of_error \(=\|A-R\| /(| | A \| \varepsilon)\), shows the merit of this approximation.

\section*{Dense Matrix Parallelism Using MPI}

\section*{General Remarks}

The central theme we use for the computing functions of the box data type is that of delivering results to a distinguished node of the machine. One of the design goals was to shield much of the complexity of distributed computing from the user.

The nodes are numbered by their "ranks." Each node has rank value MP_RANK. There are MP_NPROCS nodes, so MP_RANK \(=0,1, \ldots, M P \_N P R O C S-1\). The root node has MP_RANK \(=0\). Most of the elementary MPI material is found in Gropp, Lusk, and Skjellum (1994) and Snir, Otto, Huss-Lederman, Walker, and Dongarra (1996). Although IMSL Fortran Numerical Library users are for the most part shielded from the complexity of MPI, it is desirable for some users to learn this important topic. Users should become familiar with any referenced MPI routines and the documentation of their usage. MPI routines are not discussed here, because that is best found in the above references.
The IMSL Fortran Numerical Library algorithm for allocating the racks of the box to the processors consists of creating a schedule for the processors, followed by communication and execution of this schedule. The efficiency may be improved by using the nodes according to a specific priority order. This order can reflect information such as a powerful machine on the network other than the user's work station, or even transient network behavior. The IMSL Fortran Numerical Library allows users to define this order, but a default order is provided. A setup function establishes an order based on timing matrix products of a size given by the user. See Parallel Example 4 for an illustration of this usage.

\section*{Getting Started with Modules MPI_setup_int and MPI_node_int}

The MPI_setup_int and MPI_node_int modules are part of the IMSL Fortran Numerical Library and not part of MPI itself. Following a call to the function MP_SETUP (), the module MPI_node_int will contain information about the number of processors, the rank of a processor, the communicator for IMSL Fortran Numerical Library, and the usage priority order of the node machines. Since MPI_node_int is used by MPI_setup_int, it is not necessary to explicitly use this module. If neither MP_SETUP () nor MPI_Init () is called, then the box data type will compute entirely on one node. No routine from MPI will be called.
```

MODULE MPI_NODE_INT

```
```

INTEGER, ALLOCATABLE :: MPI NODE PRIORITY(:)
INTEGER, SAVE :: MP_LIBRARY_WORLD = huge(1)
LOGICAL, SAVE :: MPI_ROOT_WORKS = .TRUE.
INTEGER, SAVE :: MP_RANK = 0, MP_NPROCS = 1
END MODULE

```

When the function MP_SETUP () is called with no arguments, the following events occur:
- If MPI has not been initialized, it is first initialized. This step uses the routines MPI_Initialized() and possibly MPI_Init(). Users who choose not to call MP_SETUP () must make the required initialization call before using any IMSL Fortran Numerical Library code that relies on MPI for its execution. If the user's code calls an IMSL Fortran Numerical Library function utilizing the box data type and MPI has not been initialized, then the computations are performed on the root node. The only MPI routine always called in this context is MPI_Initialized(). The name MP_SETUP is pushed onto the subprogram or call stack.
- If MP_LIBRARY_WORLD equals its initial value (=huge (1)) then MPI_COMM_WORLD, the default MPI communicator, is duplicated and becomes its handle. This uses the routine MPI_Comm_dup (). Users can change the handle of MP_LIBRARY_WORLD as required by their application code. Often this issue can be ignored.
- The integers MP_RANK and MP_NPROCS are respectively the node's rank and the number of nodes in the communicator, MP_LIBRARY_WORLD. Their values require the routines MPI_Comm_size() and MPI_Comm_rank(). The default values are important when MPI is not initialized and a box data type is computed. In this case the root node is the only node and it will do all the work. No calls to MPI communication routines are made when MP_NPROCS = 1 when computing the box data type functions. A program can temporarily assign this value to force box data type computation entirely at the root node. This is desirable for problems where using many nodes would be less efficient than using the root node exclusively.
- The array MPI_NODE_PRIORITY (:) is unallocated unless the user allocates it. The IMSL Fortran Numerical Library codes use this array for assigning tasks to processors, if it is allocated. If it is not allocated, the default priority of the nodes is \(\left(0,1, \ldots, M P \_N P R O C S-1\right)\). Use of the function call MP_SETUP (N) allocates the array, as explained below. Once the array is allocated its size is MP_NPROCS. The contents of the array is a permutation of the integers \(0, \ldots\), MP_NPROCS-1. Nodes appearing at the start of the list are used first for parallel computing. A node other than the root can avoid any computing, except receiving the schedule, by setting the value MPI_NODE_PRIORITY (I) < 0 . This means that node |MPI_NODE_PRIORITY (I) | will be sent the task schedule but will not perform any significant work as part of box data type function evaluations.
- The LOGICAL flag MPI ROOT WORKS designates whether or not the root node participates in the major computation of the tasks. The root node communicates with the other nodes to complete the tasks but can be designated to do no other work. Since there may be only one processor, this flag has the default value .TRUE., assuring that one node exists to do work. When more than one processor is available users can consider assigning MPI_ROOT_WORKS=.FALSE. This is desirable when the alternate nodes have equal or greater computational resources compared with the root node. Parallel Example 4 illustrates this
usage. A single problem is given a box data type, with one rack. The computing is done at the node, other than the root, with highest priority. This example requires more than one processor since the root does no work.

When the generic function MP_SETUP ( N ) is called, where N is a positive integer, a call to MP_SETUP () is first made, using no argument. Use just one of these calls to MP_SETUP(). This initializes the MPI system and the other parameters described above. The array
MPI_NODE_PRIORITY (: ) is allocated with size MP_NPROCS. Then DOUBLE PRECISION matrix products \(C=A B\), where \(A\) and \(B\) are \(N\) by \(N\) matrices, are computed at each node and the elapsed time is recorded. These elapsed times are sorted and the contents of MPI_NODE_PRIORITY (: ) are permuted in accordance with the shortest times yielding the highest priority. All the nodes in the communicator MP_LIBRARY_WORLD are timed. The array MPI_NODE_PRIORITY (: ) is then broadcast from the root to the remaining nodes of MP_LIBRARY_WORLD using the routine MPI_Bcast (). Timing matrix products to define the node priority is relevant because the effort to compute \(C\) is comparable to that of many linear algebra computations of similar size. Users are free to define their own node priority and broadcast the array MPI_NODE_PRIORITY (: ) to the alternate nodes in the communicator.

To print any IMSL Fortran Numerical Library error messages that have occurred at any node, and to finalize MPI, use the function call MP_SETUP ('Final'). Case of the string 'Final' is not important. Any error messages pending will be discarded after printing on the root node. This is triggered by popping the name 'MP_SETUP' from the subprogram stack or returning to Level 1 in the stack. Users can obtain error messages by popping the stack to Level 1 and still continuing with MPI calls. This requires executing call e1pop ('MP_SETUP'). To continue on after summarizing errors execute call elpsh ('MP_SETUP'). More details about the error processor are found in Reference Material chapter of this manual.
Messages are printed by nodes from largest rank to smallest, which is the root node. Use of the routine MPI_Finalize () is made within MP_SETUP('Final'), which shuts down MPI. After MPI_Finalize () is called, the value of MP_NPROCS \(=0\). This flags that MPI has been initialized and terminated. It cannot be initialized again in the same program unit execution. No MPI routine is defined when MP_NPROCS has this value.

\section*{Using Processors}

There are certain pitfalls to avoid when using IMSL Fortran Numerical Library and box data types as implemented with MPI. A fundamental requirement is to allow all processors to participate in parts of the program where their presence is needed for correctness. It is incorrect to have a program unit that restricts nodes from executing a block of code required when computing with the box data type. On the other hand it is appropriate to restrict computations with rank-2 arrays to the root node. This is not required, but the results for the alternate nodes are normally discarded. This will avoid gratuitous error messages that may appear at alternate nodes.
Observe that only the root has a correct result for a box data type function. Alternate nodes have the constant value one as the result. The reason for this is that during the computation of the functions, sub-problems are allocated to the alternate nodes by the root, but for only the root to utilize the result. If a user needs a value at the other nodes, then the root must send it to the nodes. See Parallel Example 3 for an illustration of this usage. Convergence information is computed at the root node and broadcast to the others. Without this step some nodes would not terminate the
loop even when corrections at the root become small. This would cause the program to be incorrect.

\section*{Sparse Matrix Computations}

\section*{Introduction}

This section is concerned with methods for computing with sparse matrices. Our primary goal is to give the appearance of simplicity and allow ease-of-use in dealing with these calculations. The underlying principle in our design is to use Fortran 2003 standard support for derived types with initialized and allocatable components. To perform data storage and conversions we use overloaded assignment to hide complexity. The operations currently supported are:
- defining entries of the matrices,
- adding sparse matrices,
- forming products of sparse matrices and dense vectors or matrices,
- solving linear systems of algebraic equations
- condition number computation
- conversion of sparse matrices or dense arrays to the converse
- storage management operations

The definition of the sparse matrices starts with a triplet consisting of the row and column indices and a value at that entry. By setting a flag in the derived type SLU_Options, repeated values may be accumulated to yield a value that is the sum of all triplets for that matrix entry. A diagram for constructing a single precision sparse \(10000 \times 10000\) matrix, H , is illustrated with the pseudocode fragment:
Use linear_operators
Integer I, J; Real(Kind(1.e0)) value, x(10000)
Type(s_sparse) A
Type(s_hbc_sparse) H
1. Define non-zero values of A with repeated overloaded assignments
\[
A=s \_e n t r y(I, ~ J, ~ v a l u e) .
\]
2. Convert to computational Harwell-Boeing form with the overloaded assignment \(\mathrm{H}=\mathrm{A}\).
3. Compute with sparse matrix \(H\), e.g. \(x=H\).ix. \(x\).

A basic feature is that there are four sparse matrix derived types, Types (s_hbc_sparse), (d_hbc_sparse), (c_hbc_sparse), and (z_hbc_sparse). These respectively handle single, double, complex and double-complex data. The defined operators work with a sparse matrix and a corresponding dense array of the same precision and data type. There is no mixing of data types such as a sparse double precision matrix multiplied by a single precision vector. To accommodate that case an intermediate double precision quantity will be created that ascends the single precision
vector to a double precision vector. The table below shows the operations that are valid with sparse matrix types.
\begin{tabular}{|c|c|c|c|}
\hline Mathematical Operation & Operation Notation & Input Terms & Output Terms \\
\hline \(y=H^{-1} x\) & y = H . \(\mathrm{ix} . \mathrm{x}\) & \(H_{n \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(y=x^{T} H^{-1} \equiv H^{-T} x\) & \(y=x\).xi. \(H\) & \(H_{n \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(Y=H^{-1} X_{n \times r}\) & \(\mathbf{Y}=\mathbf{H}\).ix. \(\mathbf{X}\) & \(H_{n \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{n}, 1: \mathrm{r})\) \\
\hline \(Y=X_{r \times n} H^{-1} \equiv\left(H^{-T} X^{T}\right)^{T}\) & \(\mathbf{Y}=\mathrm{X} . \mathrm{xi} . \mathrm{H}\) & \(H_{n \times n}\) sparse, \(\mathrm{X}(1: \mathrm{r}, 1: \mathrm{n}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{n})\) \\
\hline \(y=H x\) & y = H.x. x & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{m})\) \\
\hline \(y=x^{T} H \equiv H^{T} x\) & \(y=x . x . H\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq m\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(Y=H X_{n \times r}\) & \(\mathbf{Y}=\mathbf{H} . \mathrm{x} . \mathrm{X}\) & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{m}, 1: \mathrm{r})\) \\
\hline \(Y=X_{r \times m} H\) & \(\mathbf{Y}=\mathbf{X} . \mathrm{x} . \mathrm{H}\) & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{r}, 1: \mathrm{k}), k \geq m\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{n})\) \\
\hline \(K=H^{T}\) & \(\mathrm{K}=. \mathrm{t} . \mathrm{H}\) & \(H_{m \times n}\) sparse & \(K_{n \times m}\) sparse \\
\hline \(K=H^{H}=\bar{H}^{T}\) & \(\mathrm{K}=. \mathrm{h} . \mathrm{H}\) & \(H_{m \times n}\) sparse, complex & \(K_{n \times m}\) sparse \\
\hline \(y=H^{T} x\) & y = H .tx. x & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq m\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(Y=H^{T} X_{m \times r}\) & Y = H .tx. X & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq m\) & \(\mathrm{Y}(1: \mathrm{n}, 1: \mathrm{r})\) \\
\hline \(y=x^{T} H\) & Y = x .tx. H & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq m\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(Y=X_{r \times m}^{T} H\) & Y = X .tx. H & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq m\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{n})\) \\
\hline \(y=H x^{T}\) & y = H.xt. x & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & y(1:m) \\
\hline \(Y=H X_{n \times r}^{T}\) & Y = H .xt. \(\mathbf{X}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}, 1: \mathrm{r}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{m}, 1: \mathrm{r})\) \\
\hline \(y=x H^{T}\) & \(y=x . x t . H\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{m})\) \\
\hline \(Y=X_{r \times n} H^{T}\) & Y = X .xt. H & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{r}, 1: \mathrm{k}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{m})\) \\
\hline \(y=H^{H} x=\bar{H}^{T} x\) & y = H .hx. x & \(H_{m \times n}\) sparse \(^{3}, \mathrm{x}(1: \mathrm{k}), k \geq m\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline \(Y=H^{H} X_{m \times r}=\bar{H}^{T} X_{m \times r}\) & Y = H.hx. X & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq m\) & \(\mathrm{Y}(1: \mathrm{n}, 1: \mathrm{r})\) \\
\hline \(y=x^{H} H=\bar{x}^{T} H\) & Y = x .hx. \(\mathbf{H}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq m\) & \(\mathrm{y}(1: \mathrm{n})\) \\
\hline
\end{tabular}
\({ }^{3}\) The operators .hx. and .xh. apply to sparse complex matrices only. For real matrices use the .tx. and .xt. operators.
\begin{tabular}{|l|l|l|l|}
\hline Mathematical Operation & Operation Notation & Input Terms & Output Terms \\
\hline\(Y=X_{r \times m}^{H} H=\bar{X}_{r \times m}^{T} H\) & \(\mathbf{Y}=\mathbf{X} . \mathbf{h x .} \mathbf{H}\) & \(H_{m \times n}\) sparse, \(\mathrm{X}(1: \mathrm{k}, 1: \mathrm{r}), k \geq m\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{n})\) \\
\hline\(y=H x^{H}=H \bar{x}^{T}\) & \(\mathbf{y}=\mathbf{H} \cdot \mathbf{x h} \cdot \mathbf{x}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{m})\) \\
\hline\(Y=H X_{n \times r}^{H}=H \bar{X}_{n \times r}^{T}\) & \(\mathbf{Y}=\mathbf{H} \cdot \mathbf{x h} \cdot \mathbf{X}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}, 1: \mathrm{r}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{m}, 1: \mathrm{r})\) \\
\hline\(y=x H^{H}=x \bar{H}^{T}\) & \(\mathbf{y}=\mathbf{x} \cdot \mathbf{x h} \mathbf{H}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{k}), k \geq n\) & \(\mathrm{y}(1: \mathrm{m})\) \\
\hline\(Y=X_{r \times n} H^{H}=X_{r \times n} \bar{H}^{T}\) & \(\mathbf{Y}=\mathbf{X} \cdot \mathbf{x h . ~} \mathbf{H}\) & \(H_{m \times n}\) sparse, \(\mathrm{x}(1: \mathrm{r}, 1: \mathrm{k}), k \geq n\) & \(\mathrm{Y}(1: \mathrm{r}, 1: \mathrm{m})\) \\
\hline
\end{tabular}

\section*{Derived Type Definitions}

A derived type is used for the entries (triplets or coordinate format) of a sparse matrix, which consists of row and column coordinates and a corresponding value:
```

type s_entry
intēger irow
integer jcol
real(kind(1.e0)) value
end type

```

Additionally, type (d_entry), type (c_entry), and type (z_entry) are defined similarly. These support double precision, complex and complex-double precision accuracy and types.
Thus for a sparse matrix \(A\), the entry at the intersection of row irow and column jcol is the scalar value. We define a sparse matrix representation in terms of a collection of triplets. This is a convenient way for a user to define a sparse matrix. This representation is used to define the matrix entries in a user's program using overloaded assignment. There is no implied order on the collection of triplets that define this sparse matrix. Our experience shows that for writing application code the technique of using triplets to define the matrix entries is convenient and provides a workable transition from mathematical definitions of the entries to computer code. Also note that there is generally no need for the programmer to allocate the components of a matrix of type s_sparse when using the overloaded assignment: s_sparse = s_entry. The software handles this detail by reallocating and expanding those components of the s_sparse matrix as required. (For this task we use the Fortran 2003 intrinsic subroutine move_alloc (), when it is available. This routine provides an efficient way to perform a reallocation.) The amount reallocated is controlled by an expansion factor that is a component of the derived type SLU_options.
```

    type s_sparse
        integer :: mrows = 0
        integer :: ncols = 0
        integer :: numnz = 0
        integer, allocatable, dimension(:) :: irow
        integer, allocatable, dimension(:) :: jcol
        real(kind(1.e0)), allocatable, dimension(:) :: value
        type (SLU_options) options
    end type
    ```

When performing matrix computations we use the Harwell-Boeing column-oriented derived type. The row indices, for each column, are unique and increasing. The values in the colptr(1:ncols) component mark the start of the row indices and corresponding matrix entries for that column. The value colptr (ncols+1)-1 will equal the value numnz after the matrix is defined with non-zero entries. The row indices for each column are in array irow (: ). They are unique and sorted into increasing order.
```

type s_hbc_sparse
integer : : mrows = 0
integer : : ncols $=0$
integer : : numnz $=0$
integer, allocatable, dimension(:) : : irow
integer, allocatable, dimension(:) : colptr
real(kind(1.e0)), allocatable, dimension(:) : : value
type(SLU_options) options
end type

```

Additionally we support types (d_hbc_sparse), type (c_hbc_sparse), and type (z_hbc_sparse). These will have analogous support for the operations defined with type (s_hbc_sparse) and others that follow. From now on we only mention type (s_hbc_sparse).

All components of the type (s_sparse) object are self-explanatory except for the one named type (SLU_options). This component contains various parameters for managing the data structure, and for computing matrix products and linear system solutions. Normally these components do not need to be changed from their default values.

The derived type SLU_Options carries extra required information. That data needed for SuperLU \({ }^{4}\) is labeled with a comment. The remaining data is needed by IMSL codes that call on SuperLU. Of particular importance is the Sequence attribute statement. This prevents the Fortran compiler from rearranging the order of the components. Maintaining this order is required since the derived type SLU_Options is passed to a IMSL C code that uses the information as a C structure. The Sequence statement orders the Fortran-defined data so that it matches the C code structure.

\section*{Type SLU_options}
```

Sequence
Integer :: unique = 1 ! Each new entry is unique -IMSL
Integer :: Accumulate = 0
! Accumulate or assemble duplicated entries in
! a ?_sparse matrix. This flag is checked
! when executing an overloaded assignment
! with a Harwell-Boeing = ?_sparse matrix.
! The default is not to accumulate (0)
! Assign the value 1 to accumulate.
Integer :: handle(2) = 0

```
\({ }^{4}\) SuperLU is used to support the defined operations .ix. and .xi., and the condition number function, cond (). SuperLU is well-tested. Distributed and threaded versions are available but these are not used here in our software at present.
```

            ! Each HBC matrix requiring an LU
            ! decomposition will have allocated
            ! arrays whose start is pointed to by
            ! this value. In cases where the OS
            ! uses }64\mathrm{ bit addressing 8 bytes are used.
    Integer :: Info = - 1
                            ! Flag returned after LU factorization (SuperLU)
    Integer :: Fact = 0 !DOFACT - SuperLU
    Integer :: Equil = 1 !YES
    Integer :: ColPerm = 3 !COLAMD
    Integer :: Trans = 0 !NOTRANS
    Integer :: IterRefine = 1 !REFINE
    Integer :: PrintStat = 0 !NO
    Integer :: SymmetricMode = 0 ! NO
    Integer :: PivotGrowth = 0 !NO
    Integer :: ConditionNumber = 0 ! NO
    Integer :: RowPerm = 0 !NO
    Integer :: ReplaceTinyPivot = 0 ! NO
    Integer :: SolveInitialized = 0 !NO
    Integer :: RefineInitialized = 0 !NO
    Real (Kind(1.d0)) :: DiagPivotThresh = 1.d0 ! SuperLU
    Real (Kind(1.d0)) :: expansion_factor = 1.2 ! VNI -
    ! The factor to use when expanding storage. Any value > 1.
! can be used such that the integer part of this factor times
! any integer > 9 provides at least a value of 1 increase.
Integer :: Cond Iteration Max = 30
! Maximum number of Lanczos and inverse iterations with sparse COND().
Integer Alignment_Dummy
End Type

```

\section*{Overloaded Assignments}

A natural way to define a sparse matrix is in terms of its triplets. The basic tool used here to define all the non-zero entries is overloaded assignment. Fortran 90, and further updates to the standard, supports a hidden subroutine call, packaged in a module, when an assignment is executed between differing derived types. Thus if a Fortran program has a declaration type (s_sparse) A, then the overloaded assignment statement
```

A = s_entry(I, J, value)

```
has the effect of calling subroutines that result in joining the matrix entry value at the intersection of row I and column J. The components of A are managed to hold any number of values. The number of rows, columns and non-zero values are updated as new triplets are assigned. Also the arrays that hold the triplets are re-allocated and expanded, as required, to hold newly assigned triplets.

The code snippet for this operation, and others that follow, will require use of the module linear_operators. If new space is required in the assignment, a reallocation of the components of the matrix A will occur. The user does not have to manage the details.
```

Use linear_operators
Type(s_sparse) A

```
```

{For all entries in A, A = s_entry(I, J, Value)}

```

\section*{Sparse = Collection of Triplets}

The Harwell-Boeing sparse matrix data types are used for computations. An assignment, \(\mathrm{H}=\mathrm{A}\), implies deallocating any allocated components of \(H\), allocating new storage, and sorting the collection of triplets provided as input in the sparse matrix A. If the accumulation flag is set in \(\mathrm{H} \% \mathrm{O}\) tions\%accumulate, the duplicate row indices in a column are reduced to a single entry and the corresponding values are added to yield a final value. The assignment \(\mathrm{H}=0\) deallocates the allocated components and returns \(H\) to its initialized state, except for any changes to the component SLU_options. A similar comment holds for the assignment, \(\mathrm{A}=0\).
```

Use linear_operators
Type(s_sparse) A
Type(s_hbc_sparse) H
...
{For all nonzero matrix entries, A = s_entry(I, J, Value)}
H}=
A = 0 ! Clear and deallocate components of A
H=0 ! Clear and deallocate components of H

```

\section*{Sparse = Dense}

The non-zero entries of the dense array are converted to a Harwell-Boeing sparse matrix. As a first step any allocated components are cleared and then allocated as needed to hold the non-zero values of the dense array. The specific dimensions of array \(D\) are arbitrary.
```

Use linear_operators
Type(s_hbc_sparse) H
Integer, parameter :: M=1000, N=1000
Real (kind(1.e0)) D(M,N)
{Define entries of D}
H = D

```

\section*{Dense = Sparse}

For some applications it is convenient to expand a sparse matrix into a dense matrix. The specific dimensions of array \(D\) are arbitrary.
```

Use linear_operators
Type(s_hbc_sparse) H
Integer, parameter :: M=1000, N=1000
Real (kind(1.e0)) D(M,N)
{Define entries of H}

```
\[
D=H
\]

\section*{Scalar = s_hbc_entry(Sparse, I, J)}

This assignment gets the value at the intersection of row I and column \(J\) of the Harwell-Boeing sparse matrix. There must be type agreement with the function and sparse matrix type. Use a prefix of \(d_{-}, c_{-}\), or \(z_{-}\)for double, complex, or double complex values.

Use inear_operators
Type(s_hbc_sparse) H
Real (kind(1.e0)) value
\{Define entries of \(H, ~ I ~ a n d ~ J\}\)
value = s_hbc_entry (H, I, J)

\section*{.X.}

Computes matrix-matrix or matrix-vector product.

\section*{Operator Return Value}

Matrix containing the product of A and B. (Output)

\section*{Required Operands}

A- Left operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
\(\boldsymbol{B}\) - Right operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
If A has rank one, B must have rank two.
If B has rank one, A must have rank two.
If A has rank three, B must have rank three. If B has rank three, A must have rank three.

\section*{FORTRAN 90 Interface}
A.x. B

\section*{Description}

Computes the product of matrix or vector \(A\) and matrix or vector \(B\). The results are in a precision and data type that ascends to the most accurate or complex operand.

Rank three operation is defined as follows:
```

do i = 1, min(size(A,3), size(B,3))
X(:,:,i) = A(:,:,i) .x. B(:,:,i)
end do

```
.x. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex03.f90)}
```

use linear_operators
implicit none
! This is the equivalent of Example 3 for LIN_SOL_GEN using operators.
integer, parameter :: n=32
real(kind(le0)) :: one=1e0, zero=0e0, A(n,n), b(n), x(n)
real(kind(1e0)) change_new, change_old
real(kind(1d0)) :: d_zero=0d0, c(n), d(n,n), y(n)
! Generate a random matrix and right-hand side.
A = rand(A); b= rand(b)
! Save double precision copies of the matrix and right-hand side.
D = A
C = b
! Compute single precision inverse to compute the iterative refinement.
A = .i. A
! Start solution at zero. Update it to an accurate solution
! with each iteration.
y = d_zero
change_old = huge (one)
iterative refinement: do
! Compute the rēsidual with higher accuracy than the data.
b = c - (D .x. y)
! Compute the update in single precision.
x = A .x. b
y = x + y
change_new = norm(x)
! Exit when changes are no longer decreasing.
if (change_new >= change_old) exit iterative_refinement
change_old = change_new
end do itērative refinēment
write (*,*) 'Example 3 for LIN_SOL_GEN (operators) is correct.'
end

```

\section*{Sparse Matrix Example}

Consider the one-dimensional Dirichlet problem
\[
\frac{d^{2} u}{d x^{2}}=f(x), \quad a<x<b, \quad u(a)=u_{a}=u_{1}, u(b)=u_{b}=u_{N}
\]

Using a standard approach to solving this involves approximating the second derivative operator with central divided differences
\[
\frac{d^{2} u}{d x^{2}} \doteq \frac{u_{i-1}-2 u_{i}+u_{i+1}}{h^{2}}, \quad h=(b-a) /(N-1), \quad i=2, \ldots, N-1, \quad N>2
\]

This leads to the sparse linear algebraic system \(M u=w\). The definitions for these terms are implied in the following Fortran program.
```

Subroutine document exl
! Illustrate a 1D Pöisson equation with Dirichlet boundary conditions.
! This module defines the structures and overloaded assignment code.
Use linear_operators
Implicit None
!
Integer :: I
Integer, Parameter :: N = 1000
Real (Kind(1.d0)) :: f, h, r, w (N), a = 0.d0, b = 1.d0, \&
u_a = 0.d0, u_b = 1.d0, u (N)
Type (d_sparse) M
Type (d`hbc sparse) K
External f
! Define the difference used.
h = (b-a) / (N-1)
r = 1.d0 / h ** 2
! Fill in the matrix entries.
! Isolated equation for the left boundary condition.
M = d_entry (1, 1, r)
Do I = 2, N - 1
M = d_entry (I, I-1, r)
M = d_entry (I, I,-2*r)
M = d_entry (I, I+1, r)
End Do
! Isolated equation for the right boundary condition.
M = d_entry (N, N, r)
! Fill in the right-hand side (a dense vector).
Do I = 2, N - 1
W (I) = f (a+(I-1)*h)
End Do
! Insert the known end conditions. These should be satisfied
! almost exactly, up to rounding errors.
w (1) = u_a * r
w (N) = u_b * r
! Ready to solve}..
! Conversion to Harwell-Boeing format using overloaded assignment
K=M
! Solve the system using an IMSL defined operator.
u = K .ix. w
! The parentheses are needed because of precedence rules.
! Compute residuals and overwrite w(:) with these values.
w = w - (K .x. u)
End Subroutine
!

```
```

Function f (x)
Real (Kind(1.d0)) :: f, x
! Define a hat function, peaked at x=0.5.
If (x <= 0.5d0) Then
f = x
Else
f = 1.d0 - x
End If
End Function

```

\section*{Parallel Example (parallel_ex03.f90)}

This example shows the box data type used while obtaining an accurate solution of several systems. Important in this example is the fact that only the root will achieve convergence, which controls program flow out of the loop. Therefore the nodes must share the root's view of convergence, and that is the reason for the broadcast of the update from root to the nodes. Note that when writing an explicit call to an MPI routine there must be the line INCLUDE 'mpif.h', placed just after the IMPLICIT NONE statement. Any number of nodes can be used.
use linear_operators
use mpi_setup_int
implicit none
INCLUDE 'mpif.h'
! This is the equivalent of Parallel Example 3 for .i. and iterative
! refinement with box date types, operators and functions.
integer, parameter : : \(n=32\), \(n r=4\)
integer IERROR
real(kind(1e0)) : : one=1e0, zero=0e0
real (kind (le0)) : : A \((n, n, n r), b(n, 1, n r), x(n, 1, n r)\)
real(kind(1e0)) change_old(nr), change_new(nr)
real(kind(1d0)) : : d_zero=0d0, \(c(n, 1, n \bar{r}), D(n, n, n r), y(n, 1, n r)\)
! Setup for MPI.
MP_NPROCS=MP_SETUP ()
! Generate a random matrix and right-hand side.
\(A=\operatorname{rand}(A) ; b=\operatorname{rand}(b)\)
! Save double precision copies of the matrix and right-hand side.
\(D=A\)
\(\mathrm{c}=\mathrm{b}\)
```

! Get single precision inverse to compute the iterative refinement.

```
\(\mathrm{A}=\).i. A
```

! Start solution at zero. Update it to a more accurate solution

```
! with each iteration.
y = d_zero
change_old = huge (one)
ITERATIVE_REFINEMENT: DO
```

! Compute the residual with higher accuracy than the data.
b = c - (D .x. y)
! Compute the update in single precision.
x = A. .x. b
y = x + y
change_new = norm(x)
! All processors must share the root's test of convergence.
CALL MPI_BCAST(change_new, nr, MPI_REAL, 0, \&
MP_LIBRARY_WORLD, IERROR)
! Exit when changes are no longer decreasing.
if (ALL(change_new >= change_old)) exit iterative_refinement
change old = change_new
end DO ITERRATIVE_REFINEMENT
IF(MP_RANK == 0) write (*,*) 'Parallel Example 3 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
en\overline{d}

```
.tx.

Computes transpose matrix-matrix or transpose matrix-vector product.

\section*{Operator Return Value}

Matrix containing the product of \(\mathrm{A}^{\mathrm{T}}\) and B . (Output)

\section*{Required Operands}
\(\boldsymbol{A}\) - Left operand matrix. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types,
?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
\(\boldsymbol{B}\) - Right operand matrix or vector. This is an array of rank 1,2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
If A has rank three, B must have rank three. If B has rank three, A must have rank three.

\section*{FORTRAN 90 Interface}

A .tx. B

\section*{Description}

Computes the product of the transpose of matrix A and matrix or vector B. The results are in a precision and data type that ascends to the most accurate or complex operand.

Rank three operation is defined as follows:
```

do i = 1, min(size(A,3), size(B,3))
X(:,:,i) = A(:,:,i) .tx. B(:,:,i)
end do

```
.tx. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex05.f90)}
```

use linear operators
implicit none
! This is the equivalent of Example 1 for LIN_SOL_SELF using operators
! and functions.
integer, parameter :: m=64, n=32
real(kind(le0)) :: one=1.0e0, err
real(kind(le0)) A (n,n), b (n,n), C(m,n), d(m,n), x(n,n)
! Generate two rectangular random matrices.
C = rand(C); d=rand(d)
! Form the normal equations for the rectangular system.
A = C .tx. C; b = C .tx. d
! Compute the solution for Ax = b, A is symmetric.
x = A .ix. b
! Check the results.
err = norm(b - (A .x. x))/(norm(A) +norm(b))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 1 for LIN_SOL_SELF (operators) is correct.'
end if
end

```

\section*{Sparse Matrix Example}
```

use wrrrn_int
use linear_operators
type (s_sparse) S
type (s_hbc_sparse) H

```
```

integer, parameter :: N=3
real (kind(1.e0)) x(N,N), y(N,N), B(N,N)
real (kind(1.e0)) err
S = s_entry (1, 1, 2.0)
S = s_entry (1, 3, 1.0)
S = s_entry (2, 2, 4.0)
S = s_entry (3, 3, 6.0)
H=S}\mp@subsup{S}{}{-}\mathrm{ ! sparse
X = H ! dense equivalent of H
B = rand(B)
Y = H .tx. B
call wrren ( 'H', X)
call wrren ( 'B', b)
call wrrrn ( 'H .tx. B ', y)
! Check the results.
err = norm(y - (X .tx. B))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .tx. operator is correct.'
end if
end

```

\section*{Output}

H
\begin{tabular}{rrr}
1 & 2 & 3 \\
2.000 & 0.000 & 1.000 \\
0.000 & 4.000 & 0.000 \\
0.000 & 0.000 & 6.000
\end{tabular}

B
\begin{tabular}{lrrr} 
& 1 & 2 & 3 \\
1 & 0.8711 & 0.4467 & 0.4743 \\
2 & 0.8315 & 0.7257 & 0.4518 \\
3 & 0.6839 & 0.0561 & 0.6972
\end{tabular}


\section*{Parallel Example (parallel_ex05.f90)}
```

    use linear operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 5 using box data types,
! operators and functions.

```
```

    integer, parameter :: m=64, n=32, nr=4
    real(kind(1e0)) :: one=1e0, err(nr)
    real(kind(le0)), dimension(n,n,nr) :: A, b, x
    real(kind(le0)), dimension(m,n,nr) :: C, d
    ! Setup for MPI.
mp_nprocs = mp_setup()
! Generate two rectangular random matrices, only
! at the root node.
if (mp_rank == 0) then
C = rand (C); d=rand (d)
endif
! Form the normal equations for the rectangular system.
A = C .tx. C; b = C .tx. d
! Compute the solution for Ax = b.
x = A .ix. b
! Check the results.
err = norm(b - (A .x. x)) /(norm(A) +norm(b))
if (ALL(err <= sqrt(epsilon(one))) .AND. MP_RANK == 0) \&
write (*,*) 'Parallel Example 5 is correct.'
! See to any error messages and quit MPI.
mp_nprocs = mp_setup('Final')
end
.xt.
4MPI
CAPABLE

```

Computes matrix- transpose matrix product.

\section*{Operator Return Value}

Matrix containing the product of A and \(\mathrm{B}^{\mathrm{T}}\). (Output)

\section*{Required Operands}
\(\boldsymbol{A}\) - Left operand matrix or vector. This is an array of rank 1,2 , or 3 . It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ? hbc_sparse.
\(\boldsymbol{B}\) - Right operand matrix. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types,
?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
If A has rank three, B must have rank three.
If B has rank three, A must have rank three.

\section*{FORTRAN 90 Interface}

A . xt. B

\section*{Description}

Computes the product of matrix or vector A and the transpose of matrix B. The results are in a precision and data type that ascends to the most accurate or complex operand.

Rank three operation is defined as follows:
```

    do i = 1, min(size(A, 3), size(B, 3))
            X(:,:,i) = A(:,:,i) .xt. B(:,:,i)
        end do
    ```
. xt. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex14.f90)}
```

    use linear_operators
    implicit none
    !
integer, parameter :: n=32
real(kind(1d0)) :: one=1d0, zero=0d0
real(kind(1d0)) A(n,n), P(n,n), Q (n,n), \&
S_D(n), U_D(n,n), V_D (n,n)
! Generate a random matrix.
A = rand(A)
! Compute the singular value decomposition.
S_D = SVD (A, U=U_D, V=V_D)
! Compute the (left) orthogonal factor.
P = U_D . xt. V_D
! Compute the (right) self-adjoint factor.
Q = V_D .x. diag(S_D) .xt. V_D
! Check the results.
if (norm( EYE(n) - (P .xt. P)) \&
<= sqrt(epsilon(one))) then
if (norm(A - (P .x. Q))/norm(A) \&
<= sqrt(epsilon(one))) then

```
```

                    write (*,*) 'Example 2 for LIN_SOL_SVD (operators) is correct.
    end if
    end if
end

```

\section*{Sparse Matrix Example}
```

use wrrrn_int
use linear
type (s sparse) S
type (s_hbc_sparse) H
integer, parameter :: N=3
real (kind(1.e0)) x(N,N), y(N,N), a(N,N)
real (kind(1.e0)) err
S = s_entry (1, 1, 2.0)
S = s_entry (1, 3, 1.0)
S = s_entry (2, 2, 4.0)
S = s entry (3, 3, 6.0)
H = S ! sparse
X = H ! dense equivalent of H
A = rand(A)
Y = A .xt. H
call wrrrn ( 'A', A)
call wrrrn ( 'H', X)
call wrrrn ( 'A.xt. H', y)
! Check the results.
err = norm(y - (A .xt. X))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .xt. operator is correct.
end if
end

```

\section*{Output}

A
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 0.5423 & 0.2380 & 0.9250 \\
2 & 0.0844 & 0.1323 & 0.1937 \\
3 & 0.4146 & 0.3135 & 0.7757
\end{tabular}

\section*{H}
\begin{tabular}{|c|c|c|c|}
\hline & 1 & 2 & 3 \\
\hline 1 & 2.000 & 0.000 & 1.000 \\
\hline 2 & 0.000 & 4.000 & 0.000 \\
\hline 3 & 0.000 & 0.000 & 6.000 \\
\hline & A & \[
\text { .xt. } \mathrm{H}_{2}
\] & 3 \\
\hline 1 & 2.010 & 0.952 & 5.550 \\
\hline 2 & 0.363 & 0.529 & 1.162 \\
\hline 3 & 1.605 & 1.254 & 4.654 \\
\hline \multicolumn{4}{|l|}{Sparse example for .xt. operator is correct.} \\
\hline
\end{tabular}

\section*{Parallel Example (parallel_ex15.f90)}

A "Polar Decomposition" of several matrices are computed. The box data type and the SVD () function are used. Orthogonality and small residuals are checked to verify that the results are correct.
```

    use linear_operators
    use mpi setup int
    implicit none
    ! This is the equivalent of Parallel Example 15 using operators and,
! functions for a polar decomposition.
integer, parameter :: n=33, nr=3
real(kind(1d0)) :: one=1d0, zero=0d0
real(kind(1d0)),dimension(n,n,nr) :: A, P, Q, \&
S D (n,nr), U D, V D
real(kin\overline{d}(1d0)) TEM\overline{P}1(nr\overline{)},\operatorname{TEMP2(nr)}
! Setup for MPI:
mp_nprocs = mp_setup()
! Generate a random matrix.
if(mp_rank == 0) A = rand(A)
! Compute the singular value decomposition.
S_D = SVD(A, U=U_D, V=V_D)
! Compute the (left) orthogonal factor.
P = U_D . xt. V_D
! Compute the (right) self-adjoint factor.
Q = V_D .x. diag(S_D) .xt. V_D
! Check the results for orthogonali}ty an
! small residuals.
TEMP1 = NORM(spread(EYE (n),3,nr) - (p .xt. p))
TEMP2 = NORM(A - (P .X. Q)) / NORM(A)
if (ALL(TEMP1 <= sqrt(epsilon(one))) .and. \&
ALL(TEMP2 <= sqrt(epsilon(one)))) then
if(mp_rank == 0)\&
write (*,*) 'Parallel Example 15 is correct.'
end if
! See to any error messages and exit MPI.
mp_nprocs = mp_setup('Final')
end

```
.hx.


Computes conjugate transpose matrix-matrix product.

\section*{Operator Return Value}

Matrix containing the product of \(\mathrm{A}^{\mathrm{H}}\) and B . (Output)

\section*{Required Operands}
\(\boldsymbol{A}\) — Left operand matrix. This is an array of rank 2 or 3 . It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
\(\boldsymbol{B}\) - Right operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
If A has rank three, B must have rank three.
If B has rank three, A must have rank three.

\section*{FORTRAN 90 Interface}

A .hx. B

\section*{Description}

Computes the product of the conjugate transpose of matrix A and matrix or vector B. The results are in a precision and data type that ascends to the most accurate or complex operand.

Rank three operation is defined as follows:
```

do i = 1, min(size(A,3), size(B,3))
X(:,:,i) = A(:,:,i) .hx. B(:,:,i)
end do

```
.hx. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex32.f90)}
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 4 (using operators) for LIN_EIG_GEN.
integer, parameter :: n=17
real(kind(1d0)), parameter :: one=1d0
real(kind(1d0)), dimension(n,n) :: A, C
real(kind(1d0)) variation(n), eta
complex(kind(1d0)), dimension(n,n) :: U, V, e(n), d(n)
! Generate a random matrix.

```
```

    A = rand (A)
    ! Compute the eigenvalues, left- and right- eigenvectors.
D = EIG(A, W=V); E = EIG(.t.A, W=U)
! Compute condition numbers and variations of eigenvalues.
variation = norm(A)/abs(diagonals( U .hx. V))
! Now perturb the data in the matrix by the relative factors
! eta=sqrt(epsilon) and solve for values again. Check the
! differences compared to the estimates. They should not exceed
! the bounds.
eta = sqrt(epsilon(one))
C = A + eta* (2*rand(A)-1)*A
D = EIG(C)
! Looking at the differences of absolute values accounts for
! switching signs on the imaginary parts.
if (count(abs(d)-abs(e) > eta*variation) == 0) then
write (*,*) 'Example 4 for LIN EIG GEN (operators) is correct.'
end if
end

```

\section*{Sparse Matrix Example}
```

use wrcrn_int
use linear_operators
type (c_sparse) S
type (c hbc sparse) H
integer, parameter : : N=3
complex (kind(1.e0)) x(N,N), y(N,N), A(N,N)
real (kind(1.e0)) err
S = c_entry (1, 1, (2.0, 1.0) )
S = c_entry (1, 3, (1.0, 3.0))
S = c_entry (2, 2, (4.0, -1.0))
S = c_entry (3, 3, (6.0, 2.0))
H = S ! sparse
X = H ! dense equivalent of H
A= rand(A)
Y = H .hx. A
call wrcrn ( 'H', X)
call wrcrn ( 'A', a)
call wrcrn ( 'H .hx. A ', y)
! Check the results.
err = norm(y - (X .hx. A))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .hx. operator is correct.
end if
end

```

\section*{Output}

H

```

A

| 1 | $(0.6278$, | $0.8475)$ | $(0.8007,0.4179)$ | $(0.4512,0.2601)$ |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 2 | $(0.1249$, | $0.4675)$ | $(0.7957,0.1609)$ | $(0.4228,0.0507)$ |
| 3 | $(0.4608$, | $0.0891)$ | $(0.3181,0.9180)$ | $(0.9961,0.1939)$ |

H .hx. A
1223
$(2.103,1.067)(2.019,0.035)(1.163,0.069)$
$(0.032,1.995)(3.022,1.439) \quad(1.640,0.626)$
$(6.113,-1.423)(5.799,2.888)(7.596,-1.922)$
Sparse example for .hx. operator is correct.

```

\section*{Parallel Example}
```

use linear_operators
use mpi_setup_int
integer, parameter :: N=32, nr=4
complex (kind(1.e0)) A(N,N,nr), B(N,N,nr), Y(N,N,nr)
! Setup for MPI
mp_nprocs = mp_setup()
if (mp rank == 0) then
A = 票利(A)
B = rand (B)
end if
Y = A .hx. B
mp_nprocs = mp_setup ('Final')
end

```
.xh.


Computes a matrix-conjugate transpose matrix product.

\section*{Operator Return Value}

Matrix containing the product of A and \(\mathrm{B}^{\mathrm{H}}\). (Output)

\section*{Required Operands}

A- Left operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ? hbc sparse. (Input)
Note that A and B cannot both be ? hbc sparse.
\(\boldsymbol{B}\) - Right operand matrix. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types,
?_hbc_sparse. (Input)
Note that A and B cannot both be ?_hbc_sparse.
If A has rank three, \(B\) must have rank three.
If \(B\) has rank three, A must have rank three.

\section*{FORTRAN 90 Interface}
A.xh. B

\section*{Description}

Computes the product of matrix or vector A and the conjugate transpose of matrix B . The results are in a precision and data type that ascends to the most accurate or complex operand.
Rank three operation is defined as follows:
```

do i = 1, min(size(A,3), size(B,3))
X(:,:,i) = A(:,:,i) .xh. B(:,:,i)
end do

```
.xh. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example}
```

use wrcrn int
use linear__operators
integer, parameter :: N=3
complex (kind(1.e0)) A(N,N), B(N,N), Y(N,N)
A = rand(A)
B = rand(B)
Y = A . xh. B
call wrcrn ( 'A', a)
call wrcrn ( 'H', b)
call wrcrn ( 'A .xh. B ', y)
end

```

\section*{Output}

A
\begin{tabular}{llllll}
1 & \((0.8071,0.0054)\) & \((0.5617,0.2508)\) & \((0.0223,0.5555)\) \\
2 & \((0.9380,0.5181)\) & \((0.8895,0.9512)\) & \((0.7951,0.6010)\) \\
3 & \((0.8349,0.7291)\) & \((0.4162,0.5255)\) & \((0.7388,0.0309)\)
\end{tabular}

B
\begin{tabular}{llllll} 
& & 1 \\
1 & \((0.5342\), & \(0.2246)\) & \((0.9045,0.0550)\) & \((0.4576,0.3173)\) \\
2 & \((0.5531,0.3362)\) & \((0.0757,0.3970)\) & \((0.6807,0.8625)\) \\
3 & \((0.3553\), & \(0.9157)\) & \((0.0951,0.7807)\) & \((0.4853,0.0617)\)
\end{tabular}
\begin{tabular}{lllll}
1 & \((1.141\), & \(0.265)\) & \((1.085,-0.113)\) & \((0.586,-0.884)\) \\
2 & \((2.029\), & \(0.900)\) & \((2.198,-0.587)\) & \((2.058,-1.036)\) \\
3 & \((1.363\), & \(0.434)\) & \((1.477,-0.619)\) & \((1.775,-0.811)\)
\end{tabular}

\section*{Sparse Matrix Example}
```

use wrcrn_int
use linear_operators
type (c_sparse) S
type (c_hbc_sparse) H
integer, parameter :: N=3
complex (kind(1.e0)) x (N,N), y(N,N), A(N,N)
real (kind(1.e0)) err
S = c_entry (1, 1, (2.0, 1.0) )
S = c_entry (1, 3, (1.0, 3.0))
S = c_entry (2, 2, (4.0, -1.0))
S = c_entry (3, 3, (6.0, 2.0))
H = S ! sparse
X = H ! dense equivalent of H
A= rand(A)
Y = A . xh. H
call wrcrn ( 'A', a)
call wrcrn ( 'H', X)
call wrcrn ( 'A .xh. H ', y)
! Check the results.
err = norm(y - (A .xh. X))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .xh. operator is correct.'
end if
end

```

\section*{Output}

A
\(1 \quad 2 \quad 3\)
```

1 ( 0.8526, 0.3532) ( 0.1822, 0.3938) ( 0.8008, 0.1308)
2 (0.5599, 0.8914) (0.7541, 0.5163) (0.8713, 0.9580)
3 ( 0.9947, 0.2735) ( 0.6237, 0.2137) ( 0.3802, 0.8903)

```

H
```

1 ( 2.000, 1.000) ( 0.000, 0.000) ( 1.000, 3.000)
2 ( 0.000, 0.000) ( 4.000,-1.000) ( 0.000, 0.000)
3(0.000,0.000) ( 0.000, 0.000) ( 6.000, 2.000)

```

A .xh. H
\begin{tabular}{lrlrr} 
& \multicolumn{1}{r}{} \\
1 & \((3.252,-2.418)\) & \((0.335,1.757)\) & \((5.066,-0.817)\) \\
2 & \((5.757,-0.433)\) & \((2.500,2.819)\) & \((7.144,4.005)\) \\
3 & \((5.314,-0.698)\) & \((2.281,1.478)\) & \((4.062,4.581)\) \\
Sparse example for.\(x h\). operator is correct.
\end{tabular}

\section*{Parallel Example}
```

use linear_operators
use mpi_setup_int
integer, parameter :: N=32, nr=4
complex (kind(1.e0)) A(N,N,nr), B(N,N,nr), Y(N,N,nr)
! Setup for MPI
mp_nprocs = mp_setup()
if (mp_rank == 0) then
A = rand (A)
B = rand (B)
end if
Y = A . xh. B
mp_nprocs = mp_setup ('Final')
end

```

\section*{.t.}

Computes the transpose of a matrix.

\section*{Operator Return Value}

Matrix containing the transpose of A. (Output)

\section*{Required Operand}

A-Matrix for which the transpose is to be computed. This is a real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input).

\section*{FORTRAN 90 Interface}
.t. A

\section*{Description}

Computes the transpose of matrix A. The operation may be read transpose, and the results are the mathematical objects in a precision and data type that matches the operand. Since this is a unary operation, it has higher Fortran 90 precedence than any other intrinsic unary array operation.
.t. can be used with either dense or sparse matrices.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex07.f90)}
```

use linear_operators
implicit none
! This is the equivalent of Example 3 (using operators) for LIN_SOL_SELF.
integer tries
integer, parameter :: m=8, n=4, k=2
integer ipivots(n+1)
real(kind(1d0)) :: one=1.0d0, err
real(kind(1dO)) a(n,n), b(n,1), c(m,n), x(n,1), \&
e(n), ATEMP (n,n)
type(d_options) :: iopti(4)
! Generate a random rectangular matrix.
C = rand(C)
! Generate a random right hand side for use in the inverse
! iteration.
b = rand(b)
! Compute the positive definite matrix.
A = C .tx. C; A = (A+.t.A)/2
! Obtain just the eigenvalues.
E = EIG(A)
! Use packaged option to reset the value of a small diagonal.
iopti(4) = 0
iopti(1) = d_options(d_lin_sol_self_set_small,\&
epsilon(one)*abs(E(1)))
! Use packaged option to save the factorization.
iopti(2) = d_lin_sol_self_save_factors
! Suppress error messages and stopping due to singularity
! of the matrix, which is expected.
iopti(3) = d_lin_sol_self_no_sing_mess

```
```

    ATEMP = A
    ! Compute A-eigenvalue*I as the coefficient matrix.
! Use eigenvalue number k.
A = A - e(k)*EYE(n)
do tries=1,2
call lin_sol_self(A, b, x, \&
pivots=ipivots, iopt=iopti)
! When code is re-entered, the already computed factorization
! is used.
iopti(4) = d_lin_sol_self_solve_A
! Reset right-hand side in the direction of the eigenvector.
B = UNIT(x)
end do
! Normalize the eigenvector.
x = UNIT(x)
! Check the results.
b=ATEMP .x. x
err = dot_product(x(1:n,1), b(1:n,1)) - e(k)
! If any result is not accurate, quit with no printing.
if (abs(err) <= sqrt(epsilon(one))*E(1)) then
write (*,*) 'Example 3 for LIN_SOL_SELF (operators) is correct.'
end if
end

```

\section*{Sparse Matrix Example}
```

use wrrrn_int
use linear__operators
type (s_sparse) S
type (s_hbc sparse) H, HT
integer, parameter :: N=3
real (kind(1.e0)) X(3,3), XT(3,3)
real (kind(1.e0)) err
S = s_entry (1, 1, 2.0)
S = s_entry (1, 3, 1.0)
S = s_entry (2, 2, 4.0)
S = s_entry (3, 3, 6.0)
H = S ! sparse
X = H ! dense equivalent of H
HT = .t. H
XT = HT ! dense equivalent of HT
call wrrrn ( 'H', X)
call wrrrn ( 'H Transpose', XT)
! Check the results.
err = norm(XT - (.t. X))

```
```

        if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Sparse example for .t. operator is correct.'
        end if
    end

```

\section*{Output}

.h.
Computes the conjugate transpose of a matrix.

\section*{Operator Return Value}

Matrix containing the conjugate transpose of A . (Output)

\section*{Required Operand}
\(\boldsymbol{A}\) - Matrix for which the conjugate transpose is to be computed. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)

\section*{FORTRAN 90 Interface}
.h. A

\section*{Description}

Computes the conjugate transpose of matrix A. The operation may be read adjoint, and the results are the mathematical objects in a precision and data type that matches the operand. Since this is a unary operation, it has higher Fortran 90 precedence than any other intrinsic unary array operation.
.h. can be used with either dense or sparse matrices.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex34.f90)}
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 2 (using operators) for LIN_GEIG_GEN.
integer, parameter :: n=32
real(kind(1d0)), parameter :: one=1d0, zero=0d0
real(kind(1d0)) err, alpha(n)
complex(kind(1d0)), dimension(n,n) :: A, B, C, D, V
! Generate random matrices for both A and B.
C = rand(C); D = rand(D)
A = C +..h.C; B = D .hx. D; B = (B + .h. B) / 2
ALPHA = EIG (A, B=B, W=V)
! Check that residuals are small. Use a real array for alpha
! since the eigenvalues are known to be real.
err= norm((A .x. V) - (B .x. V .x. diag(alpha)),1)/\&
(norm(A,1) +norm(B,1)*norm(alpha,1))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 2 for LIN_GEIG_GEN (operators) is correct.'
end if
end

```
Sparse Matrix Example
```

use wrcrn_int
use linear_operators
type (c_sparse) S
type (c_hbc_sparse) H, HT
integer, parameter :: N=3
complex (kind(1.e0)) X (3,3), XT (3,3)
real (kind(1.e0)) err
S = c_entry (1, 1, (2.0, 1.0) )
S = c_entry (1, 3, (1.0, 3.0))
S = c_entry (2, 2, (4.0, -1.0))
S = c_entry (3, 3, (6.0, 2.0))
H=S
X = H ! dense equivalent of H
HT = .h. H
XT = HT ! dense equivalent of HT
call wrcrn ( 'H', X)
call wrcrn ( 'H Conjugate Transpose', XT)
! Check the results.

```
```

        err = norm(XT - (.h. X))
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Sparse example for .h. operator is correct.'
    end if
    end

```

\section*{Output}

\section*{H}


\section*{.i.}

\section*{勻MPI \\ CAPABLE}

Computes the inverse matrix.

\section*{Operator Return Value}

Matrix containing the inverse of A. (Output)

\section*{Required Operand}

A-Matrix for which the inverse is to be computed. This is an array of rank 2 or 3. It may be real, double, complex, double complex. (Input)

\section*{Optional Variables, Reserved Names}

This operator uses the routines LIN_SOL_GEN or LIN_SOL_LSQ (See Chapter 1, "Linear Systems").

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for .i. & Option Value \\
\hline Use_lin_sol_gen_only & 1 \\
Use_lin_sol_lsq_only & 2 \\
I_options_for_lin_sol_gen & 3 \\
I_options_for_lin_sol_lsq & 4 \\
Skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_inv_options (:) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_inv_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_GEN and LIN_SOL_LSQ located in Chapter 1, "Linear Systems" for the specific options for these routines.

\section*{FORTRAN 90 Interface}
.i. A

\section*{Description}

Computes the inverse matrix for square non-singular matrices using LIN_SOL_GEN, or the Moore-Penrose generalized inverse matrix for singular square matrices or rectangular matrices using LIN_SOL_LSQ. The operation may be read inverse or generalized inverse, and the results are in a precision and data type that matches the operand.

This operator requires a single operand. Since this is a unary operation, it has higher Fortran 90 precedence than any other intrinsic array operation.

\section*{Examples}

Dense Matrix Example (operator_ex02.f90)
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 2 for LIN_SOL_GEN using operators
! and functions.
integer, parameter :: n=32
real(kind(le0)) :: one=1e0, err, det_A, det_i
real(kind(le0)), dimension(n,n) :: A, inv
! Generate a random matrix.

```
```

        A = rand(A)
    ! Compute the matrix inverse and its determinant.
inv = .i.A; det_A = det(A)
! Compute the determiñant for the inverse matrix.
det_i = det(inv)
! Check the quality of both left and right inverses.
err = (norm(EYE (n)-(A .x. inv)) +norm(EYE(n)-(inv.x.A)))/cond(A)
if (err <= sqrt(epsilon(one)) .and. abs(det_A*det_i - one) <= \&
sqrt(epsilon(one))) \&
write (*,*) 'Example 2 for LIN_SOL_GEN (operators) is correct.'
end

```

\section*{Parallel Example (parallel_ex02.f90)}
```

    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 2 for .i. and det() with box
! data types, operators and functions.
integer, parameter :: n=32, nr=4
integer J
real(kind(le0)) :: one=1e0
real(kind(le0)), dimension(nr) :: err, det_A, det_i
real(kind(le0)), dimension(n,n,nr) :: A, inv, R, S
! Setup for MPI.
MP NPROCS=MP_SETUP()
! Generāte a rando\overline{m} matrix.
A = rand(A)
! Compute the matrix inverse and its determinant.
inv = .i.A; det_A = det(A)
! Compute the determinant for the inverse matrix.
det_i = det(inv)
! Check the quality of both left and right inverses.
DO J=1,nr; R(:,:,J)=EYE(N); END DO
S=R; R=R-(A .x. inv); S=S-(inv .x. A)
err = (norm(R) +norm(S))/cond(A)
if (ALL(err <= sqrt(epsilon(one)) . and. \&
abs(det_A*det_i - one) <= sqrt(epsilon(one)))\&
. and. MP_RANK == 0) \&
write (*,*) 'Parallel Example 2 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{.ix.}

Computes the product of the inverse of a matrix and a vector or matrix.

\section*{Operator Return Value}

Matrix containing the product of \(\mathrm{A}^{-1}\) and B. (Output)

\section*{Required Operands}

A - Left operand matrix. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. (Input)
\(\boldsymbol{B}\) - Right operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, or double complex. (Input)

\section*{Optional Variables, Reserved Names}

This operator uses the routines LIN_SOL_GEN or LIN_SOL_LSQ (See Chapter 1, "Linear Systems").

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for .ix. & Option Value \\
\hline use_lin_sol_gen_only & 1 \\
use_lin_sol_lsq_only & 2 \\
ix_options_for_lin_sol_gen & 3 \\
ix_options_for_lin_sol_lsq & 4 \\
Skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_invx_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_invx_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_GEN and LIN_SOL_LSQ located in Chapter 1, "Linear Systems" for the specific options for these routines.

\section*{FORTRAN 90 Interface}

A .ix. B

\section*{Description}

Computes the product of the inverse of matrix \(A\) and vector or matrix \(B\), for square non-singular matrices or the corresponding Moore-Penrose generalized inverse matrix for singular square matrices or rectangular matrices. The operation may be read generalized inverse times. The results are in a precision and data type that matches the most accurate or complex operand.
.ix. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex01.f90)}
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 1 for LIN_SOL_GEN, with operators
! and functions.
integer, parameter :: n=32
real(kind(1e0)) :: one=1.0e0, err
real(kind(le0)), dimension(n,n) :: A, b, x
! Generate random matrices for A and b:
A = rand(A); b=rand(b)
! Compute the solution matrix of Ax = b.
x = A .ix. b
! Check the results.
err = norm(b - (A .x. x))/(norm(A)*norm(x) +norm(b))
if (err <= sqrt(epsilon(one))) \&
write (*,*) 'Example 1 for LIN_SOL_GEN (operators) is correct.'
end

```

\section*{Sparse Matrix Example 1}
```

use wrrrn_int
use linear_operators
type (s_sparse) S
type (s_hbc_sparse) H
integer, parameter :: N=3
real (kind(1.e0)) x (N,N), y(N,N), B(N,N)
real (kind(1.e0)) err
S = s_entry (1, 1, 2.0)
S = s_entry (1, 3, 1.0)
S = s_entry (2, 2, 4.0)
S = s_entry (3, 3, 6.0)
H = S- ! sparse

```
```

X = H ! dense equivalent of H
B= rand (B)
Y = H .ix. B
call wrrrn ('H', X)
call wrren ( 'B', b)
call wrrrn ( 'H .ix. B ', y)
! Check the results.
err = norm(y - (X .ix. B))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .ix. operator is correct.'
end if
end

```

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{2}{c}{\begin{tabular}{l} 
H
\end{tabular}} \\
& 1 & 2 & 3 \\
1 & 2.000 & 0.000 & 1.000 \\
2 & 0.000 & 4.000 & 0.000 \\
3 & 0.000 & 0.000 & 6.000
\end{tabular}
\begin{tabular}{lrrr} 
& \multicolumn{2}{c}{ B } & \\
& 1 & 2 & 3 \\
1 & 0.8292 & 0.5697 & 0.1687 \\
2 & 0.9670 & 0.7296 & 0.0603 \\
3 & 0.1458 & 0.2726 & 0.8809
\end{tabular}
\begin{tabular}{rrrr} 
& \multicolumn{2}{c}{ H .ix. B } & 3 \\
& 1 & 2 & 3 \\
1 & 0.4025 & 0.2621 & 0.0109 \\
2 & 0.2417 & 0.1824 & 0.0151 \\
3 & 0.0243 & 0.0454 & 0.1468
\end{tabular}

\section*{Sparse Matrix Example 2: Plane Poisson Problem with Dirichlet Boundary Conditions}

We want to calculate a numerical solution, which approximates the true solution of the Poisson (boundary value) problem in the solution domain \(\Omega\), a rectangle in \(\mathbb{R}^{2}\) The equation is
\[
\Delta u=\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}}=f \text { in } \Omega
\]

There are Dirichlet boundary conditions \(u=g\) on \(\partial_{1} \Omega\)
There are further Neumann boundary conditions \(\frac{\partial u}{\partial n}=h\) on \(\partial_{2} \Omega\)
The boundary arcs comprising \(\partial_{1} \Omega \cup \partial_{2} \Omega=\partial \Omega\) are mutually exclusive of each other. The functions \(f, g, h\) are defined on their respective domains.

We will solve an instance of this problem by using finite differences to approximate the derivatives. This will lead to a sparse system of linear algebraic equations. Note that particular cases of this problem can be solved with methods that are likely to be more efficient or more appropriate than the one illustrated here. We use this method to illustrate our matrix data handling routines and defined operators.

The area of the rectangle \(\Omega\) is \(a \times b\) with the origin fixed at the lower left or SW corner. The dimension along the \(x\) axis is \(a\) and along the \(y\) axis is \(b\). A rectangular \(n \times m\) uniform grid is defined on \(\Omega\) where each sub-rectangle in the grid has sides \(\Delta x=a /(n-1)\) and \(\Delta y=a /(m-1)\). What is perhaps novel in our development is that the boundary values are written into the \((m \times n)^{2}\) linear system as trivial equations. This leads to more unknowns than standard approaches to this problem but the complexity of describing the equations into computer code is reduced. The boundary conditions are naturally in place when the solution is obtained. No reshaping is required.

We number the approximate values of \(u\) at the grid points and collapse them into a single vector. Given a coordinate of the \(\operatorname{grid}(i, j), \quad((i=1, \ldots, n), j=1, \ldots, m)\), we use the mapping \(J=i+(j-1) n\) to define coordinate \(J\) of this vector. This mapping enables us to define the matrix that is used to solve for the values of \(u\) at the grid points.

For the Neumann boundary conditions we take \(\partial_{2} \Omega\) to be the East face of the rectangle. Along that edge we have \(\frac{\partial u}{\partial n}=\frac{\partial u}{\partial x}\), and we impose the smooth interface \(h=0\).

Our use of finite differences is standard. For the differential equation we approximate
\[
\frac{\partial^{2} u}{\partial x^{2}}+\frac{\partial^{2} u}{\partial y^{2}} \doteq\left(\frac{u_{i-1, j}-2 u_{i, j}+u_{i+1, j}}{\Delta x^{2}}\right)+\left(\frac{u_{i, j-1}-2 u_{i, j}+u_{i, j+1}}{\Delta y^{2}}\right)=f\left(x_{i}, y_{j}\right)
\]
at the inner grid points \((i, j), \quad((i=2, \ldots, n-1), j=2, \ldots, m-1)\). For the Neumann condition we approximate
\[
\frac{\partial u}{\partial x} \doteq\left(\frac{u_{n, j}-u_{n-1, j}}{\Delta x}\right)=0, \quad j=1, \ldots, m
\]

The remaining equations come from the Dirichlet conditions given on \(\partial_{1} \Omega\).
To illustrate three examples of solutions to this problem we consider
1. A Laplace Equation with the boundary conditions
\[
\begin{aligned}
& u=0, \text { on the South Edge } \\
& u=0.7, \text { on the East Edge } \\
& u=1, \text { on the North Edge } \\
& u=0.3, \text { on the West Edge }
\end{aligned}
\]

The function \(f=0\) for all \((x, y)\). Graphical results are shown below with the title "Problem Case 1"
2. A Poisson equation with the boundary conditions \(u=0\) on all of the edges and \(f(x, y)=-\sin (\pi x) \sin (\pi y)\). This problem has the solution \(u(x, y)=-f(x, y) /\left(2 \pi^{2}\right)\), and this identity provides a way of verifying that the accuracy is within the truncation error implied by the difference equations. Graphical results are shown with the title "Problem Case 2" The residual function verifies the expected accuracy.
3. The Laplace Equation with the boundary conditions of Problem Case 1 except that the boundary condition on the East Edge is replaced by the Neumann condition \(\frac{\partial u}{\partial x}=0\). Graphical results are shown as "Problem Case 3."
```

Subroutine document ex2
! Illustrate a 2D Poisson equation with Dirichlet and
! Neumann boundary conditions.
! These modules defines the structures and overloaded assignment code.
Use linear_operators
Implicit None
Integer :: I, J, JJ, MY_CASE, IFILE
Integer, Parameter : : N N
Real (Kind(1.d0)) :: a = 1.d0, b = 1.d0
Real (Kind(1.d0)) :: delx, dely, r, s, pi, scale
Real (Kind(1.d0)) :: u(N*M), w(N*M), P(N, M)
Real (Kind(1.e0)) :: TS, TE
CHARACTER(LEN=12) : : PR_LABEL (3)=\&
(/'Laplace','Poisson','Neumann'/)
! Mapping function (in-line) for grid coordinate to
! matrix-vector indexing.
JJ (I, J) = I + (J-1) * N
! Define sparse matrices to hold problem data.
Type (d_sparse) C
Type (d_hbc_sparse) D
! Define differences and related parameters.
delx = a / (N-1)
dely = b / (M-1)
r = 1.d0 / delx ** 2
s = 1.d0 / dely ** 2
Do MY_CASE = 1, 3
For MY CA\overline{SE =}

1. Solve boundary value problem with f=0 and Dirichlet
boundary conditions.
2. Solve Poisson equation with f such that a solution is known.
Use zero boundary condtions.
3. Solve boundary value problem with Dirichlet condtions as in 1.
except on the East edge. There the partial WRT x is zero.
Set timer for building the matrix.
Call cpu_time (TS)
Do I = 2, N - 1
```
```

    Do J = 2, M - 1
    ! Write entries for second partials WRT x and y.
C = d_entry (JJ(I, J), JJ(I-1, J), r)
C = d-entry (JJ(I, J), JJ(I+1, J), r)
C = d_entry (JJ(I, J), JJ(I, J),-2*(r+s))
C = d_entry (JJ(I, J), JJ(I, J-1), s)
C = d_entry (JJ(I, J), JJ(I, J+1), s)
!
! Define components of the right-hand side.
w (JJ(I, J)) = f((I-1)*delx, (J-1)*dely, MY_CASE)
End Do
End Do
! Write entries for Dirichlet boundary conditions.
! First do the South edge, then the West, then the North.
Select Case (MY_CASE)
Case (1:2)
Do I = 1, N
C = d_entry (JJ(I, 1), JJ(I, 1), r+s)
w (JJ(I, 1)) = g ((I-1)*delx, 0.d0, MY_CASE) * (r+s)
End Do
Do J = 2, M - 1
C = d_entry (JJ(1, J), JJ(1, J), r+s)
w (JJ(1, J)) = g (0.d0, (J-1)*dely, MY CASE) * (r+s)
End Do
Do I = 1, N
C = d entry (JJ(I, M), JJ(I, M), r+s)
w (JJ(I, M)) = g ((I-1)*delx, b, MY_CASE) * (r+s)
End Do
Do J = 2, M - 1
C = d_entry (JJ(N, J), JJ(N, J), (r+s))
w (JJ(N, J)) = g (a, (J-1)*dely, MY CASE) * (r+s)
End Do
Case (3)
! Write entries for the boundary values but avoid the East edge.
Do I = 1, N - 1
C = d_entry (JJ(I, 1), JJ(I, 1), r+s)
w (JJ-(I, 1)) = g ((I-1)*delx, 0.d0, MY_CASE) * (r+s)
End Do
Do J = 2, M - 1
C = d_entry (JJ(1, J), JJ(1, J), r+s)
w (JJ(1, J)) = g (0.d0, (J-1)*dely, MY_CASE) * (r+s)
End Do
Do I = 1, N - 1
C = d_entry (JJ(I, M), JJ(I, M), r+s)
w (JJ\overline{(I, M) ) = g ((I-1)*delx, b, MY_CASE) * (r+s)}
End Do
! Write entries for the Neumann condition on the East edge.
Do J = 1, M
C = d_entry (JJ(N, J), JJ(N, J), 1.d0/delx)
C = d entry (JJ(N, J), JJ(N-2, J),-1.d0/delx)
w (JJ(- N, J)) = 0.d0
End Do
End Select
!
! Convert to Harwell-Boeing format for solving.

```
```

    D = C
    !
Call cpu_time (TE)
Write (*','(A,F6.2," S. - ",A)') "Time to build matrix = ", \&
TE - TS, PR_LABEL(MY_CASE)
! Clear sparse triplets.
C=0
!
! Turn off iterative refinement for maximal performance.
! This is generally not recommended unless
! the problem is known not to require it.
If (MY CASE == 2) D%options%iterRefine = 0
! This is the solve step.
Call cpu time (TS)
u = D.商x. w
Call cpu_time (TE)
Write (*','(A,I6," is",F6.2," S")') \&
"Time to solve system of size = ", N * M, TE - TS
! This is a second solve step using the factorization
! from the first step.
Call cpu_time (TS)
u = D .ix. w
Call cpu_time (TE)
!
If(MY_CASE == 1) then
Write (*,'(A,I6," is",F6.2," S")') \&
"Time for a 2nd system of size (iterative refinement) =", \&
N * M, TE - TS
Else
Write (*,'(A,I6," is",F6.2," S")') \&
"Time for a 2nd system of size (without refinement) =", \&
N * M, TE - TS
End if
! Convert solution vector to a 2D array of values.
P = reshape (u , (/ N, M /))
If (MY CASE == 2) Then
pi = dconst ('pi')
!
scale = - 0.5 / pi ** 2
Do I = 1, N
Do J = 1, M
! This uses the known form of the solution to compute residuals.
P (I, J) = P (I, J) - scale * f ((I-1)*delx, \&
(J-1)*dely, MY_CASE)
End Do
End Do
!
write (*,*) minval (P), " = min solution error "
write (*,*) maxval (P), " = max solution error "
End If
Write (*,'(A,1pE12.4/)') "Condition number of matrix", cond (D)
! Clear all matrix data for next problem case.
D = 0
!
End Do ! MY_CASE

```
```

Contains
Function f (x, y, MY_CASE)
implicit none
! Define the right-hand side function associated with the
! "del" operator.
Real (Kind(1.dO)) x, y, f, pi
Integer MY CASE
if(MY_CASE == 2) THEN
pi}\mp@subsup{}{}{-}=\mathrm{ dconst ('pi')
f = - Sin (pi*x) * Sin (pi*y)
Else
f}=0.d
End If
End Function
!
Function g (x, y, MY_CASE)
implicit none
! Define the edge values, except along East edge, x = a.
Real (Kind(1.d0)) x, y, g
Integer MY_CASE
! Fill in a constant value along each edge.
If (MY_CASE == 1 .Or. MY_CASE == 3) Then
If }\overline{(y = = .d0) Then
g = 0.d0
Return
End If
If (y == b) Then
g = 1.d0
Return
End If
If (x == 0.d0) Then
g = 0.3d0
Return
End If
If (x == a) Then
g = 0.7d0
End If
Else
g = 0.d0
!
End If
!
End Function
End Subroutine

```


Problem Case 1


Problem Case 2


Problem Case 3

\section*{Parallel Example (parallel_ex01.f90)}
```

    use linear operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 1 for .ix., with box data types
! and functions.
integer, parameter :: n=32, nr=4
real(kind(1e0)) :: one=1e0
real(kind(le0)), dimension(n,n,nr) :: A, b, x, err(nr)
! Setup for MPI.
MP_NPROCS=MP_SETUP()
! Generate random matrices for A and b:
A = rand(A); b=rand(b)
! Compute the box solution matrix of Ax = b.
x = A .ix. b
! Check the results.
err = norm(b - (A .x. x))/(norm(A)*norm(x) +norm(b))
if (ALL(err <= sqrt(epsilon(one))) .and. MP_RANK == 0) \&
write (*,*) 'Parallel Example 1 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')

```
.xi.

Computes the product of a matrix or vector and the inverse of a matrix.

\section*{Operator Return Value}

Matrix containing the product of A and \(\mathrm{B}^{-1}\). (Output)

\section*{Required Operands}

A-Right operand matrix or vector. This is an array of rank 1, 2, or 3. It may be real, double, complex, or double complex. (Input)
\(\boldsymbol{B}\) - Left operand matrix. This is an array of rank 2, or 3. It may be real, double, complex, double complex, or one of the computational sparse matrix derived types,
?_hbc_sparse. (Input)

\section*{Optional Variables, Reserved Names}

This operator uses the routines LIN_SOL_GEN or LIN_SOL_LSQ (See Chapter 1, "Linear Systems").

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for .xi. & Option Value \\
\hline use_lin_sol_gen_only & 1 \\
use_lin_sol_lsq_only & 2 \\
xi_options_for_lin_sol_gen & 3 \\
xi_options_for_lin_sol_lsq & 4 \\
Skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_xinv_options (:) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_xinv_options_once (:) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_GEN and LIN_SOL_LSQ located in Chapter 1, "Linear Systems" for the specific options for these routines.

\section*{FORTRAN 90 Interface}
A .xi. B

\section*{Description}

Computes the product of matrix A and the inverse of matrix B , for square non-singular matrices or the corresponding Moore-Penrose generalized inverse matrix for singular square matrices or rectangular matrices. The operation may be read times generalized inverse. The results are in a precision and data type that matches the most accurate or complex operand.
.xi. can be used with either dense or sparse matrices. It is MPI capable for dense matrices only.

\section*{Examples}

\section*{Dense Matrix Example}
```

    use linear_operators
    implicit none
    integer, parameter :: n=32
    real(kind(1e0)) : : one=1.0e0, err
    real(kind(le0)), dimension(n,n) :: A, b, x
    ! Generate random matrices for A and b:
A = rand(A); b=rand (b)
! Compute the solution matrix of xA = b.
x = b .xi. A
! Check the results.
err = norm(b - (x .x. A))/(norm(A)*norm(x) +norm(b))
if (err <= sqrt(epsilon(one))) \&
write (*,*) 'Example for .xi. operator is correct.'
end

```

\section*{Sparse Matrix Example}
```

use wrrrn_int
use linear_operators

```
```

type (s sparse) S
type (s_hbc_sparse) H
integer, parameter :: N=3
real (kind(1.e0)) x(N,N), y(N,N), a(N,N)
real (kind(1.e0)) err
S = s_entry (1, 1, 2.0)
S = s_entry (1, 3, 1.0)
S = s_entry (2, 2, 4.0)
S = s_entry (3, 3, 6.0)
H = S ! sparse
X = H ! dense equivalent of H
A = rand(A)
Y = A .xi. H
call wrrrn ( 'A', A)
call wrrrn ( 'H', X)
call wrrrn ( 'A .xi. H', y)
! Check the results.
err = norm(y - (A .xi. X))
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Sparse example for .xi. operator is correct.'
end if
end

```

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{2}{c}{ A } & \\
& 1 & 2 & 3 \\
1 & 0.5926 & 0.5015 & 0.5368 \\
2 & 0.4001 & 0.9529 & 0.6988 \\
3 & 0.0412 & 0.0633 & 0.3821
\end{tabular}

H
\begin{tabular}{rrrr} 
& 1 & 2 & 3 \\
1 & 2.000 & 0.000 & 1.000 \\
2 & 0.000 & 4.000 & 0.000 \\
3 & 0.000 & 0.000 & 6.000
\end{tabular}


\section*{Parallel Example}
```

use linear_operators
use mpi_setup_int
implicit none
! This is the equivalent of Parallel Example l for .xi., with box data types
! and functions.

```
```

    integer, parameter :: n=32, nr=4
    real(kind(1e0)) :: one=1e0
    real(kind(le0)), dimension(n,n,nr) :: A, b, x, err(nr)
    ! Setup for MPI.
MP_NPROCS=MP_SETUP()
! Generate random matrices for A and b:
A = rand(A); b=rand (b)
! Compute the box solution matrix of xA = b.
x = b .xi. A
! Check the results.
err = norm(b - (x .x. A))/(norm(A)*norm(x) +norm(b))
if (ALL(err <= sqrt(epsilon(one))) .and. MP_RANK == 0) \&
write (*,*) 'Parallel Example 1 is correc}\overline{t}.
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{CHOL}

\section*{『- MPI}

Computes the Cholesky factorization of a positive-definite, symmetric or self-adjoint matrix.

\section*{Function Return Value}

Matrix containing the Cholesky factorization of A. The factor is upper triangular, \(R^{T} R=A\). (Output)

\section*{Required Argument}

A - Matrix to be factored. This argument must be a rank-2 or rank-3 array that contains a positive-definite, symmetric or self-adjoint matrix. It may be real, double, complex, double complex. (Input)
For rank-3 arrays each rank-2 array, (for fixed third subscript), is a positive-definite, symmetric or self-adjoint matrix. In this case, the output is a rank-3 array of Cholesky factors for the individual problems.

\section*{Optional Arguments, Packaged Options}

This function uses LIN_SOL_SELF (See Chapter 1, "Linear Systems"), using the appropriate options to obtain the Cholesky factorization.

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for CHOL & Option Value \\
\hline Use_lin_sol_gen_only & 4 \\
Use_lin_sol_lsq_only & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_chol_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_chol_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_SELF located in Chapter 1, "Linear Systems" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

CHOL (A)

\section*{Description}

Computes the Cholesky factorization of a positive-definite, symmetric or self-adjoint matrix, \(A\). The factor is upper triangular, \(R^{T} R=A\).

\section*{Examples}

Dense Matrix Example (operator_ex06.f90)
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 2 for LIN_SOL_SELF using operators
! and functions.
integer, parameter : : m=64, n=32
real(kind(le0)) :: one=1e0, zero=0e0, err
real(kind(le0)) A(n,n), b(n), C(m,n), d(m), cov(n,n), x(n)
! Generate a random rectangular matrix and right-hand side.
C = rand(C); d=rand(d)
! Form the normal equations for the rectangular system.
A = C .tx. C; b = C .tx. d
COV = .i. CHOL(A); COV = COV .xt. COV
! Compute the least-squares solution.
x = C .ix. d

```
```

! Compare with solution obtained using the inverse matrix.
err = norm(x - (COV .x. b))/norm(cov)
! Scale the inverse to obtain the sample covariance matrix.
COV = sum((d - (C.x. x))**2)/(m-n) * COV
! Check the results.
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 2 for LIN_SOL_SELF (operators) is correct.'
end if
end

```

\section*{Parallel Example (parallel_ex06.f90)}
```

    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 6 for box data types, operators
! and functions.
integer, parameter :: m=64, n=32, nr=4
real(kind(1e0)) :: one=1e0, zero=0e0, err(nr)
real(kind(le0)), dimension(m,n,nr) :: C, d(m,1,nr)
real(kind(le0)), dimension(n,n,nr) :: A, cov
real(kind(le0)), dimension(n,1,nr) :: b, x
! Setup for MPI:
mp_nprocs=mp_setup ()
! Generate a random rectangular matrix and right-hand side.
if(mp rank == 0) then
C = rand (C); d=rand (d)
endif
! Form the normal equations for the rectangular system.
A = C .tx. C; b = C .tx. d
COV = .i. CHOL(A); COV = COV .xt. COV
! Compute the least-squares solution.
x = C .ix. d
! Compare with solution obtained using the inverse matrix.
err = norm(x - (COV .x. b))/norm(cov)
! Check the results.
if (ALL(err <= sqrt(epsilon(one))) .and. mp_rank == 0) \&
write (*,*) 'Parallel Example 6 is correct.'
! See to any eror messages and quit MPI
mp_nprocs=mp_setup('Final')
end

```

\section*{COND}

Computes the condition number of a matrix.

\section*{Function Return Value}

Computes condition number of matrix A. This is a scalar for the case where A is rank-2 or a sparse matrix. It is a rank-1 array when \(A\) is a dense rank- 3 array. (Output)

\section*{Required Argument}

A - Matrix for which the condition number is to be computed. The matrix may be real, double, complex, double-complex, or one of the computational sparse matrix derived types, ?_hbc_sparse. For an array of type real, double, complex, or double-complex the array may be of rank-2 or rank-3.

For a dense rank-3 array, each rank-2 array section, (for fixed third subscript), is a separate problem. In this case, the output is a rank-1 array of condition numbers for each problem. (Input)

\section*{Optional Arguments, Packaged Options}

NORM_CHOICE - Integer indicating the type of norm to be used in computing the condition number.
\begin{tabular}{lllllll} 
NORM_CHOICE & \multicolumn{2}{l}{ CONDITION } & \multicolumn{2}{l}{ Square Matrix } & \multicolumn{2}{l}{ Rectangular Matrix } \\
& Number & Dense & Sparse & Dense & Sparse \\
1 & \(l_{1}\) & Yes & Yes & No & No \\
2 (Default) & \(l_{2}\) & Yes & Yes & Yes & No \\
huge (1) & \(l_{\infty}\) & Yes & Yes & No & No
\end{tabular}

This function uses LIN_SOL_SVD (see Chapter 1, "Linear Systems").
The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for COND & Option Value \\
\hline ?_cond_set_small & 1 \\
?_cond_for_lin_sol_svd & 2 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_cond_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_cond_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_SVD located in Chapter 1, "Linear Systems" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

COND (A \([, \ldots]\) )

\section*{Description}

The mathematical definitions of the condition numbers which this routine estimates are:
\[
\begin{aligned}
& l_{1} \text { condition number } \kappa_{1}(A)=\|A\|_{1} \cdot\left\|A^{-1}\right\|_{1} \\
& l_{2} \text { condition number } \kappa_{2}(A)=\|A\|_{2} \cdot\left\|A^{-1}\right\|_{2} \\
& l_{\infty} \text { condition number } \kappa_{\infty}(A)=\|A\|_{\infty} \cdot\left\|A^{-1}\right\|_{\infty}
\end{aligned}
\]

COND can be used with either dense or sparse matrices as follows:
\begin{tabular}{|l|l|l|l|l|}
\hline \multicolumn{3}{|c|}{ Square Matrix } & \multicolumn{2}{l|}{ Rectangular Matrix } \\
\hline & Dense & Sparse & Dense & Sparse \\
\hline\(l_{1}\) & Yes & Yes & No & No \\
\(l_{2}\) & Yes & Yes & Yes & No \\
\(l_{\infty}\) & Yes & Yes & No & No \\
\hline
\end{tabular}

The generic function COND can be used with either dense or sparse square matrices. This function uses LIN_SOL_SVD for dense square and rectangular matrices in computing \(\kappa_{2}(A)=s_{1} / s_{n}\). The function uses LIN_SOL_GEN for dense square matrices in computing \(\kappa_{1}(A)\) and \(\kappa_{\infty}(A)\). For sparse square matrices, the values returned for \(\kappa_{1}(A)\) and \(\kappa_{\infty}(A)\) are provided by the SuperLU linear equation solver. The condition number \(\kappa_{2}(A)=s_{1} / s_{n}\) is computed by an algorithm that first approximates \(S_{1}\) by computing the singular values of the \(k \times k\) bidiagonal matrix obtained using the Lanczos method found in Golub and Van Loan, Ed. 3, p. 495. Here \(k\) is set using the value A\%Options\%Cond_Iteration_Max, which has the default value of 30 . The value \(S_{n}^{-2}\) is obtained using the power method, Golub and Van Loan, p. 330, iterating with the inverse \(\operatorname{matrix}\left(A^{T} A\right)^{-1}=A^{-1} A^{-T}\). For complex matrices \(A^{T}\) is replaced by \(A^{H}=\bar{A}^{T}\). The dominant
eigenvalue of this inverse matrix is \(s_{n}^{-2}\). The number of iterations is limited by the parameter value \(k\) or relative accuracy equal to the cube root of machine epsilon. Some timing tests indicate that computing \(\kappa_{2}(A)\) for sparse matrices by this algorithm typically requires about twice the time as for a single linear solve using the defined operator A.ix. b.

For computation of \(\kappa_{2}(A)\) with rectangular sparse matrices one can use a dense matrix representation for the matrix. This is not recommended except for small problem sizes. For overdetermined systems of sparse least-squares equations \(A x \cong b\) a related square system is given by
\[
C\left[\begin{array}{l}
x \\
r
\end{array}\right] \equiv\left[\begin{array}{cc}
A & I_{m \times m} \\
0_{n \times n} & A^{T}
\end{array}\right]\left[\begin{array}{l}
x \\
r
\end{array}\right]=\left[\begin{array}{l}
b \\
0
\end{array}\right]
\]

One can form \(C\), which has more than twice the number of non-zeros as \(A\). But \(C\) is still sparse. One can use the condition number of \(C\) as an estimate of the accuracy for the solution vector \(x\) and the residual vector \(r\). Note that this version of the condition number is not the same as the \(l_{2}\) condition number of \(A\) but is relevant to determining the accuracy of the least-squares system.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex02.f90)}
```

use wrrrn_int
use linear_operators
integer, parameter :: N=3
real (kind(1.e0)) A(N,N)
real (kind(1.e0)) C1, C2, CINF
DATA A/2.0, 2.0, -4.0, 0.0, -1.0, 2.0, 0.0, 0.0, 5.0/
CINF = COND (A, norm_choice=huge(1))
C1 = COND (A, norm_choice=1)
C2 = COND (A)
call wrrrn ( 'A', A)
write (*,*) 'L1 condition number= ', C1
write (*,*) 'L2 condition number= ', C2
write (*,*) 'L infinity condition number= ', CINF
end

```

\section*{Output}
\begin{tabular}{rrrr} 
& \multicolumn{3}{c}{ A } \\
& 1 & 2 & \multicolumn{1}{c}{3} \\
1 & 2.000 & 0.000 & 0.000 \\
2 & 2.000 & -1.000 & 0.000 \\
3 & -4.000 & 2.000 & 5.000 \\
& & \\
L1 & condition number \(=\) & 12.0 \\
L2 & condition number \(=\) & 10.405088
\end{tabular}
```

L infinity condition number= 22.0

```

\section*{Sparse Matrix Example}
```

use wrrrn int
use linear_operators
type (s sparse) S
type (s_hbc_sparse) H
integer, parameter :: N=3
real (kind(1.e0)) X(N,N)
real (kind(1.e0)) C1, C2, CINF
S = s_entry (1, 1, 2.0)
S = s_entry (2, 1, 2.0)
S = s_entry (3, 1, -4.0)
S = s_entry (3, 2, 2.0)
S = s_entry (2, 2, -1.0)
S = s_entry (3, 3, 5.0)
H = S ! sparse
X = H ! dense equivalent of H
CINF = COND (H, norm_choice=huge(1))
C1 = COND (H, norm_choice=1)
C2 = COND (H)
call wrrrn ( 'H', X)
write (*,*) 'L1 condition number= ', C1
write (*,*) 'L2 condition number= ', C2
write (*,*) 'L infinity condition number= ', CINF
end

```

\section*{Output}

\section*{H}
\begin{tabular}{lrrr} 
& 1 & 2 & 3 \\
1 & 2.000 & 0.000 & 0.000 \\
2 & 2.000 & -1.000 & 0.000 \\
3 & -4.000 & 2.000 & 5.000 \\
& & \\
L1 condition number \(=\) & 12.0 \\
L2 condition number \(=\) & 10.405088 \\
L infinity condition & number \(=\) & 22.0
\end{tabular}

\section*{Parallel Example (parallel_ex02.f90)}
```

    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 2 for .i. and det() with box
! data types, operators and functions.
integer, parameter :: n=32, nr=4

```
```

    integer J
    real(kind(1e0)) :: one=1e0
    real(kind(1e0)), dimension(nr) :: err, det_A, det_i
    real(kind(1e0)), dimension(n,n,nr) :: A, iñv, R, \overline{S}
    ! Setup for MPI.
MP_NPROCS=MP_SETUP()
! Genera\overline{te a random matrix.}
A = rand (A)
! Compute the matrix inverse and its determinant.
inv = .i.A; det_A = det(A)
! Compute the determiñant for the inverse matrix.
det_i = det(inv)
! Check the quality of both left and right inverses.
DO J=1,nr; R(:,:,J)=EYE(N); END DO
S=R; R=R-(A .x. inv); S=S-(inv .x. A)
err = (norm(R)+norm(S))/cond(A)
if (ALL(err <= sqrt(epsilon(one)) .and. \&
abs(det_A*det_i - one) <= sqrt(epsilon(one)))\&
.and. MP_RANK == 0) \&
write (`,*) 'Parallel Example 2 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{DET}


Computes the determinant of a rectangular matrix.

\section*{Function Return Value}

Determinant of matrix A. This is a scalar for the case where A is rank 2. It is a rank-1 array of determinant values for the case where \(A\) is rank 3. (Output)

\section*{Required Argument}

A - Matrix for which the determinant is to be computed. This argument must be a rank-2 or rank-3 array that contains a rectangular matrix. It may be real, double, complex, double complex. (Input)

For rank-3 arrays, each rank-2 array (for fixed third subscript), is a separate matrix. In this case, the output is a rank-1 array of determinant values for each problem.

\section*{Optional Arguments, Packaged Options}

This function uses LIN_SOL_LSQ (see Chapter 1, "Linear Systems") to compute the \(Q R\) decomposition of \(A\), and the logarithmic value of \(\operatorname{det}(A)\), which is exponentiated for the result.

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for DET & Option Value \\
\hline ?_det_for_lin_sol_lsq & 1 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_det_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_det_options_once (:) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_LSQ located in Chapter 1, "Linear Systems" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

DET (A)

\section*{Description}

Computes the determinant of a rectangular matrix, \(A\). The evaluation is based on the \(Q R\) decomposition,
\[
Q A P=\left[\begin{array}{cc}
R_{k \times k} & 0 \\
0 & 0
\end{array}\right]
\]
and \(k=\operatorname{rank}(A)\). Thus \(\operatorname{det}(A)=s \times \operatorname{det}(R)\) where \(s=\operatorname{det}(Q) \times \operatorname{det}(P)= \pm 1\).
Even well-conditioned matrices can have determinants with values that have very large or very tiny magnitudes. The values may overflow or underflow. For this class of problems, the use of the logarithmic representation of the determinant found in LIN_SOL_GEN or LIN_SOL_LSQ is required.

\section*{Examples}

Dense Matrix Example (operator_ex02.f90)
use linear_operators
implicit nōne
! This is Example 2 for LIN_SOL_GEN using operators and functions.
```

    integer, parameter :: n=32
    real(kind(1e0)) :: one=1e0, err, det_A, det_i
    real(kind(1e0)), dimension(n,n) :: A, inv
    ! Generate a random matrix.
A = rand(A)
! Compute the matrix inverse and its determinant.
inv = .i.A; det_A = det(A)
! Compute the determiñant for the inverse matrix.
det_i = det(inv)
! Check the quality of both left and right inverses.
err = (norm(EYE(n)-(A.x. inv)) +norm(EYE (n)-(inv.x.A)))/cond(A)
if (err <= sqrt(epsilon(one)) .and. abs(det_A*det_i - one) <= \&
sqrt(epsilon(one))) \&
write (*,*) 'Example 2 for LIN_SOL_GEN (operators) is correct.'
end

```

\section*{Parallel Example (parallel_ex02.f90)}
```

    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 2 for .i. and det() with box
! data types, operators and functions.
integer, parameter : : n=32, nr=4
integer J
real(kind(le0)) :: one=1e0
real(kind(le0)), dimension(nr) :: err, det_A, det_i
real(kind(1e0)), dimension(n,n,nr) :: A, inv, R, \overline{S}
! Setup for MPI.
MP NPROCS=MP SETUP()
! Generate a random matrix.
A = rand (A)
! Compute the matrix inverse and its determinant.
inv = .i.A; det_A = det(A)
! Compute the determinant for the inverse matrix.
det_i = det(inv)
! Check the quality of both left and right inverses.
DO J=1,nr; R(:,:,J)=EYE(N); END DO
S=R; R=R-(A .x. inv); S=S-(inv.x. A)
err = (norm(R) +norm(S))/cond(A)
if (ALL(err <= sqrt(epsilon(one)) .and. \&
abs(det_A*det_i - one) <= sqrt(epsilon(one)))\&
.and. MP RANK == 0) \&

```

```

! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{DIAG}

Constructs a square diagonal matrix.

\section*{Function Return Value}

Square diagonal matrix of rank-2 if A is rank-1 or rank-3 array if A is rank-2. (Output)

\section*{Required Argument}
\(\boldsymbol{A}\) - This is a rank-1 or rank-2 array of type real, double, complex, or double complex, containing the diagonal elements. The output is a rank-2 or rank-3 array, respectively. (Input)

\section*{FORTRAN 90 Interface}
```

DIAG (A)

```

\section*{Description}

Constructs a square diagonal matrix from a rank-1 array or several diagonal matrices from a rank2 array. The dimension of the matrix is the value of the size of the rank- 1 array.

The use of DIAG may be obviated by observing that the defined operations \(C=\operatorname{diag}(x) \cdot x\). A or \(D=B . x . \operatorname{diag}(x)\) are respectively the array operations \(C=\operatorname{spread}(x\), DIM=1, NCOPIES=size (A, 1)) *A, and D = B*spread (x, DIM=2, NCOPIES=size (B, 2)). These array products are not as easy to read as the defined operations using DIAG and matrix multiply, but their use results in a more efficient code.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex13.f90)}
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 1 for LIN_SOL_SVD using operators
! and functions.
integer, parameter :: m=128, n=32
real(kind(1d0)) :: one=1d0, err
real(kind(ld0)) A(m,n), b(m), x(n), U(m,m),V(n,n), S(n),g(m)
! Generate a random matrix and right-hand side.
A = rand(A); b = rand (b)
! Compute the least-squares solution matrix of }Ax=b\mathrm{ .
S = SVD(A, U = U, V = V)
g = U .tx. b; x = V .x. diag(one/S) .x. g(1:n)
! Check the results.
err = norm(A .tx. (b - (A .x. x))) /(norm(A) +norm(x))

```
```

if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 1 for LIN_SOL_SVD (operators) is correct.'
end if
end

```

\section*{DIAGONALS}

Extracts the diagonal terms of a matrix.

\section*{Function Return Value}

Array containing the diagonal terms of matrix A. It is rank-1 or rank-2 depending on the rank of A. When A is a rank-3 array, the result is a rank-2 array consisting of each separate set of diagonals. (Output)

\section*{Required Argument}

A - Matrix from which to extract the diagonal. This is a rank-2 or rank-3 array of type real, double, complex, or double complex. The output is a rank-1 or rank-2 array, respectively. (Input)

\section*{FORTRAN 90 Interface}

DIAGONALS (A)

\section*{Description}

Extracts a rank-1 array whose values are the diagonal terms of the rank-2 array A. The size of the array is the smaller of the two dimensions of the rank-2 array.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex32.f90)}
```

    use linear operators
    implicit none
    ! This is the equivalent of Example 4 (using operators) for LIN_EIG_GEN.
integer, parameter :: n=17
real(kind(1d0)), parameter :: one=1d0
real(kind(1d0)), dimension(n,n) :: A, C
real(kind(1d0)) variation(n), eta
complex(kind(1d0)), dimension(n,n) :: U, V, e(n), d(n)
! Generate a random matrix.
A = rand(A)
! Compute the eigenvalues, left- and right- eigenvectors.
D = EIG (A, W=V); E = EIG(.t.A, W=U)

```
```

! Compute condition numbers and variations of eigenvalues.
variation = norm(A)/abs(diagonals( U .hx. V))
! Now perturb the data in the matrix by the relative factors
! eta=sqrt(epsilon) and solve for values again. Check the
! differences compared to the estimates. They should not exceed
! the bounds.
eta = sqrt(epsilon(one))
C = A + eta* (2*rand (A) -1)*A
D = EIG(C)
! Looking at the differences of absolute values accounts for
! switching signs on the imaginary parts.
if (count(abs(d)-abs(e) > eta*variation) == 0) then
write (*,*) 'Example 4 for LIN_EIG_GEN (operators) is correct.'
end if
end

```

\section*{EIG}

\section*{勻MPI \\ CAPABLE}

Computes the eigenvalue-eigenvector decomposition of an ordinary or generalized eigenvalue problem.

\section*{Function Return Value}

Rank-1 or rank-2 complex array of eigenvalues. (Output)

\section*{Required Argument}

A-Matrix for which the eigenexpansion is to be computed. This is a square rank-2 array or a rank- 3 array with square first rank- 2 sections of type single, double, complex, or double complex. (Input)

\section*{Optional Arguments, Packaged Options}

B-Matrix B for the generalized problem, \(A x=e B x\). B must be the same type as A. (Input)
\(\boldsymbol{D}\) - Array containing the real diagonal matrix factors of the generalized eigenvalues.
(Output)
\(\boldsymbol{V}\) - Array of real eigenvectors for both the ordinary and generalized problem. Used only for the generalized problem when matrix \(B\) is singular. (Output)
\(W\) - Array of complex eigenvectors for both the ordinary and generalized problem. Do not use optional argument V when w is present. (Output)

This function uses LIN_EIG_SELF, LIN_EIG_GEN, and LIN_GEIG_GEN to compute the decompositions. See Chapter 2, "Eigensystem Analysis".

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for EIG & Option Value \\
\hline Options_for_lin_eig_self & 1 \\
options_for_lin_eig_gen & 2 \\
Options_for_lin_geig_gen & 3 \\
Skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_eig_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_eig_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_EIG_SELF, LIN_EIG_GEN, and LIN_GEIG_GEN located in Chapter 2, "Eigensystems Analysis" for the specific \(\overline{\text { options for these restines. }}\)

\section*{FORTRAN 90 Interface}
\[
\operatorname{EIG}(\mathrm{A}[, \ldots])
\]

\section*{Description}

Computes the eigenvalue-eigenvector decomposition of an ordinary or generalized eigenvalue problem.
For the ordinary eigenvalue problem, \(A x=e x\), the optional input " \(B=\) " is not used. With the generalized problem, \(A x=e B x\), the matrix \(B\) is passed as the array in the right-side of " \(\mathrm{B}=\) ". The optional output " \(D=\) " is an array required only for the generalized problem and then only when the matrix \(B\) is singular.
The array of real eigenvectors is an optional output for both the ordinary and the generalized problem. It is used as " \(\mathrm{V}=\) " where the right-side array will contain the eigenvectors. If any eigenvectors are complex, the optional output " \(\mathrm{W}=\) " must be present. In that case " \(\mathrm{V}=\) " should not be used.

\section*{Examples}

\section*{Dense Matrix Example 1 (operator_ex26.f90)}
```

use linear_operators
implicit none

```
```

! This is the equivalent of Example 2 (using operators) for LIN_EIG_SELF.
integer, parameter :: n=8
real(kind(1e0)), parameter :: one=1e0
real(kind(1e0)), dimension(n,n) :: A, d(n), v_s
! Generate a random self-adjoint matrix.
A = rand(A); A = A + .t.A
! Compute the eigenvalues and eigenvectors.
D = EIG (A,V=v_s)
! Check the results for small residuals.
if (norm((A .x. v_s) - (v_s .x. diag(D)))/abs(d(1)) <= \&
sqrt(epsilon(one))) then
write (*,*) 'Example 2 for LIN_EIG_SELF (operators) is correct.'
end if
end

```
        Dense Matrix Example 2 (operator_ex33.f90)
    use linear_operators
    implicit none
! This is the equivalent of Example 1 (using operators) for LIN_GEIG_GEN.
    integer, parameter : : \(n=32\)
    real(kind(1d0)), parameter : : one=1d0
    real(kind(1d0)) \(A(n, n), B(n, n)\), bta(n), beta_t(n), err
    complex(kind(1d0)) alpha(n), alpha_t(n), V(n, n)
! Generate random matrices for both \(A\) and \(B\).
    \(A=\operatorname{rand}(A) ; B=\operatorname{rand}(B)\)
! Compute the generalized eigenvalues.
    alpha \(=\operatorname{EIG}(A, B=B, D=b t a)\)
! Compute the full decomposition once again, \(A * V=B * V *\) values,
! and check for any error messages.
    alpha_t \(=\operatorname{EIG}\left(A, B=B, D=b e t a \_t, W=V\right)\)
! Use values from the first decomposition, vectors from the
! second decomposition, and check for small residuals.
    \(\operatorname{err}=\operatorname{norm}((\mathrm{A} . \mathrm{x} \cdot \mathrm{V} . \mathrm{x} \cdot \operatorname{diag}(\mathrm{bta}))-(\mathrm{B} . \mathrm{x} \cdot \mathrm{V} . \mathrm{x} \cdot \operatorname{diag}(\mathrm{alpha})), 1) / \&\)
        \((\operatorname{norm}(A, 1) * \operatorname{norm}(b t a, 1)+\operatorname{norm}(B, 1) * \operatorname{norm}(a l p h a, 1))\)
    if (err <= sqrt(epsilon(one))) then
        write (*,*) 'Example 1 for LIN_GEIG_GEN (operators) is correct.'
    end if
    end

\section*{Parallel Example (parallel_ex04.f90)}

Here an alternate node is used to compute the majority of a single application, and the user does not need to make any explicit calls to MPI routines. The time-consuming parts are the evaluation of the eigenvalue-eigenvector expansion, the solving step, and the residuals. To do this, the rank2 arrays are changed to a box data type with a unit third dimension. This uses parallel computing. The node priority order is established by the initial function call, MP_SETUP ( n ) . The root is restricted from working on the box data type by assigning MPI_ROOT_WORKS=.false. This example anticipates that the most efficient node, other than the root, will perform the heavy computing. Two nodes are required to execute.
```

    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 4 for matrix exponential.
! The box dimension has a single rack.
integer, parameter : : n=32, k=128, nr=1
integer i
real(kind(le0)), parameter :: one=1e0, t_max=one, delta_t=t_max/(k-1)
real(kind(1e0)) err(nr), sizes(nr), A(n, n,nr)
real(kind(le0)) t(k), y(n,k,nr), y_prime(n,k,nr)
complex(kind(le0)), dimension(n,nr) : : x (n,n,nr), z_0, \&
Z_1(n,nr,nr), y_0, d
! Setup for MPI. Establish a node priority order.
! Restrict the root from significant computing.
! Illustrates using the 'best' performing node that
! is not the root for a single task.
MP_NPROCS=MP_SETUP(n)
MPI_ROOT_WORKS=.false.
! Generate a random coefficient matrix.
A = rand (A)
! Compute the eigenvalue-eigenvector decomposition
! of the system coefficient matrix on an alternate node.
D = EIG(A, W=X)
! Generate a random initial value for the ODE system.
y_0 = rand (y_0)
! Solve complex data system that transforms the initial
! values, X z_0=y_0.
z_1= X .ix. y_0 ; z_0(:,nr) = z_1 (:,nr,nr)
! The grid of points where a solution is computed:
t = (/(i*delta_t,i=0,k-1)/)
! Compute Y and y' at the values t(1:k).

```
```

! With the eigenvalue-eigenvector decomposition AX = XD, this
! is an evaluation of EXP(A t)y_0 = y(t).
y = X . x.exp (spread (d (:, nr ),2,k)*spread (t,1,n))*spread (z_0 (:,nr), 2,k)
! This is y', derived by differentiating y(t).
y_prime = X .x. \&
spread(\overline{d}(:,nr),2,k)*exp (spread(d(:,nr),2,k)*spread(t, 1,n))* \&
spread(z_0(:,nr),2,k)
! Check results. Is y' - Ay = 0?
err = norm(y_prime-(A .x. y))
sizes=norm(y_prime) +norm(A)*norm(y)
if (ALL(err <= sqrt(epsilon(one))*sizes) .and. MP_RANK == 0) \&
write (*,*) 'Parallel Example 4 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{EYE}

Creates the identity matrix.

\section*{Function Return Value}

Identity matrix of size \(N \mathrm{x} N\) and type real . (Output)

\section*{Required Argument}
\(N\) - Size of output identity matrix. (Input)

\section*{FORTRAN 90 Interface}

EYE (N)

\section*{Description}

Creates a rank-2 square array whose diagonals are all the value one. The off-diagonals all have value zero.

\section*{Examples}

\section*{Dense Matrix Example (operator_ex07.f90)}
```

    use linear_operators
    implicit none
    ! This is the equivalent of Example 3 (using operators) for LIN_SOL_SELF.

```
```

        integer tries
        integer, parameter :: m=8, n=4, k=2
        integer ipivots(n+1)
        real(kind(1d0)) :: one=1.0d0, err
        real(kind(1d0)) a(n,n), b(n,1), c(m,n), x(n,1), &
        e(n), ATEMP(n,n)
    type(d_options) :: iopti(4)
    ! Generate a random rectangular matrix.
C = rand (C)
! Generate a random right hand side for use in the inverse
! iteration.
b = rand(b)
! Compute the positive definite matrix.
A = C .tx. C; A = (A+.t.A)/2
! Obtain just the eigenvalues.
E = EIG(A)
! Use packaged option to reset the value of a small diagonal.
iopti(4) = 0
iopti(1) = d_options(d_lin_sol_self_set_small,\&
epsilon(one)*abs(E(\overline{1})))
! Use packaged option to save the factorization.
iopti(2) = d_lin_sol_self_save_factors
! Suppress error messages and stopping due to singularity
! of the matrix, which is expected.
iopti(3) = d_lin_sol_self_no_sing_mess
ATEMP = A
! Compute A-eigenvalue*I as the coefficient matrix.
! Use eigenvalue number k.
A = A - e(k)*EYE(n)
do tries=1,2
call lin_sol_self(A, b, x, \&
pivots=ipivots, iopt=iopti)
! When code is re-entered, the already computed factorization
! is used.
iopti(4) = d_lin_sol_self_solve_A
! Reset right-hand side in the direction of the eigenvector.
B = UNIT(x)
end do
! Normalize the eigenvector.
x = UNIT(x)
! Check the results.
b=ATEMP .x. x

```
```

    err = dot_product(x(1:n,1), b(1:n,1)) - e(k)
    ! If any result is not accurate, quit with no printing.
if (abs(err) <= sqrt(epsilon(one))*E(1)) then
write (*,*) 'Example 3 for LIN_SOL_SELF (operators) is correct.'
end if
end

```

\section*{FFT}

Computes the Discrete Fourier Transform of one complex sequence.

\section*{Function Return Value}

Complex array containing the Discrete Fourier Transform of x . The result is the complex array of the same shape and rank as X . (Output)

\section*{Required Argument}
\(\boldsymbol{X}\) - Array containing the sequence for which the transform is to be computed. x is an assumed shape complex array of rank 1,2 or 3 . If \(x\) is real or double, it is converted to complex internally prior to the computation. (Input)

\section*{Optional Arguments, Packaged Options}

WORK - A COMPLEX array of the same precision as the data. For rank-1 transforms the size of WORK is \(n+15\). To define this array for each problem, set \(\operatorname{WORK}(1)=0\). Each additional rank adds the dimension of the transform plus 15 . Using the optional argument WORK increases the efficiency of the transform.

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for FFT & Option Value \\
\hline Options_for_fast_dft & 1 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_fft_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_fft_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See FAST DFT located in Chapter 6, "Transforms" for the specific options for this routine.

\section*{FORTRAN 90 Interface}
\(\operatorname{FFT}(X[, \ldots])\)

\section*{Description}

Computes the Discrete Fourier Transform of a complex sequence. This function uses FAST DFT, FAST_2DFT, and FAST_3DFT from Chapter 6.

\section*{Examples (operator_ex37.f90)}
```

    use rand_gen_int
    use fft i
    use ifft_int
    use linear_operators
    implicit none
    ! This is Example 4 for FAST_DFT (using operators).
integer j
integer, parameter :: n=40
real(kind(1e0)) :: err, one=1e0
real(kind(le0)), dimension(n) :: a, b, c, yy(n,n)
complex(kind(le0)), dimension(n) :: f, fa, fb
! Generate two random periodic sequences 'a' and 'b'.
a=rand(a); b=rand(b)
! Compute the convolution 'c' of 'a' and 'b'.
Yy(1:,1)=b
do j=2,n
YY (2:,j)=Yy (1:n-1,j-1)
YY(1,j)=YY(n,j-1)
end do
C=YY .x. a
! Compute f=inverse(transform(a)*transform(b)).
fa= fft(a)
fb}=\textrm{fft}(\textrm{b}
f=ifft(fa*fb)
! Check the Convolution Theorem:
! inverse(transform(a)*transform(b)) = convolution(a,b).
err = norm(c-f)/norm(c)
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 4 for FAST_DFT (operators) is correct.'
end if
end

```

\section*{FFT_BOX}


Computes the Discrete Fourier Transform of several complex or real sequences.

\section*{Function Return Value}

Complex array containing the Discrete Fourier Transform of the sequences in X . If x is an assumed shape complex array of rank 2,3 or 4 , the result is a complex array of the same shape and rank consisting of the DFT for each of the last rank's indices. (Output)

\section*{Required Argument}
\(\boldsymbol{X}\) - Box containing the sequences for which the transform is to be computed. x is an assumed shape complex array of rank 2,3 or 4 . If \(x\) is real or double, it is converted to complex internally prior to the computation. (Input)

\section*{Optional Arguments, Packaged Options}

WORK - A COMPLEX array of the same precision as the data. For rank-1 transforms the size of WORK is \(n+15\). To define this array for each problem, set \(\operatorname{WORK}(1)=0\). Each additional rank adds the dimension of the transform plus 15 . Using the optional argument WORK increases the efficiency of the transform

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for FFT & Option Value \\
\hline Options_for_fast_dft & 1 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_fft_box_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_fft_box_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See FAST_DFT located in Chapter 6, "Transforms" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

FFT_BOX (X [, ...])

\section*{Description}

Computes the Discrete Fourier Transform of a box of complex sequences. This function uses FAST_DFT, FAST_2DFT, and FAST_3DFT from Chapter 6.

\section*{Examples}

\section*{Parallel Example}
```

    use rand_gen_int
    use fft box int
    use iff\overline{t}bo\overline{x}_int
    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is FFT_BOX example.
integer i,j
integer, parameter :: n=40, nr=4
real(kind(le0)) :: err(nr), one=1e0
real(kind(le0)) : : a (n,1,nr), b (n,nr), c(n,1,nr), yy(n,n,nr)
complex(kind(le0)), dimension(n,nr) :: f, fa, fb, cc, aa
real(kind(1e0)), parameter::zero_par=0.e0
real(kind(1e0)) : :dummy_par(0)
integer iseed_par
type(s_options)::iopti_par(2)
! setup for MPI
MP NPROCS = MP SETUP()
! Set Random Number generator seed
iseed_par = 53976279
iopti_par(1)=s_options(s_rand_gen_generator_seed,zero_par)
iopti_par(2)=s_options(iseed_par,\overline{zero_par)}
call rand_gen(dummy_par,iopt=iopti_par)
! Generate two random periodic sequences 'a' and 'b'.
a=rand(a); b=rand(b)
! Compute the convolution 'c' of 'a' and 'b'.
do i=1,nr
aa(1:,i) = a(1:,1,i)
yy(1:,1,i)=b(1:,i)
do j=2,n
YY(2:,j,i)=Yy(1:n-1,j-1,i)
yy(1,j,i)=yy(n,j-1,i)
end do
end do

```
```

    c=yy .x. a
    ! Compute f=inverse(transform(a)*transform(b)).
fa = fft_box(aa)
fb = fft_box(b)
f=ifft_bōx(fa*fb)
! Check the Convolution Theorem:
! inverse(transform(a)*transform(b)) = convolution(a,b).
do i=1,nr
CC(1:,i) = C(1:,1,i)
end do
err = norm(cc-f)/norm(cc)
if (ALL(err <= sqrt(epsilon(one))).AND. MP_RANK == 0) then
write (*,*) 'FFT_BOX is correct.'
end if
MP_NPROCS = MP_SETUP('Final')
end

```

\section*{IFFT}

Computes the inverse of the Discrete Fourier Transform of one complex sequence.

\section*{Function Return Value}

Complex array containing the inverse of the Discrete Fourier Transform of x. The result is the complex array of the same shape and rank as X . (Output)

\section*{Required Argument}
\(\boldsymbol{X}\) - Array containing the sequence for which the inverse transform is to be computed. X is an assumed shape complex array of rank 1,2 or 3 . If \(x\) is real or double, it is converted to complex internally prior to the computation. (Input)

\section*{Optional Arguments, Packaged Options}

WORK - a COMPLEX array of the same precision as the data. For rank-1 transforms the size of WORK is \(n+15\). To define this array for each problem, set \(\operatorname{WORK}(1)=0\). Each additional rank adds the dimension of the transform plus 15 . Using the optional argument WORK increases the efficiency of the transform.

The option and derived type names are given in the following tables:
\begin{tabular}{|l|l|}
\hline Option Name for IFFT & Option Value \\
\hline options_for_fast_dft & 1 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_ifft_options (:) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_ifft_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See FAST_DFT located in Chapter 6, "Transforms" for the specific options for this routine.

\section*{FORTRAN 90 Interface}
\(\operatorname{IFFT}(X[, \ldots])\)

\section*{Description}

Computes the inverse of the Discrete Fourier Transform of a complex sequence. This function uses FAST_DFT, FAST_2DFT, and FAST_3DFT from Chapter 6.

\section*{Example (operator_ex37.f90)}
```

        use rand_gen_int
        use fft int
        use ifft_int
        use linear_operators
        implicit none
    ! This is the equivalent of Example 4 for FAST_DFT (using operators).
integer j
integer, parameter :: n=40
real(kind(1e0)) :: err, one=1e0
real(kind(1e0)), dimension(n) :: a, b, c, yy(n,n)
complex(kind(le0)), dimension(n) :: f, fa, fb
! Generate two random periodic sequences 'a' and 'b'.
a=rand(a); b=rand(b)
! Compute the convolution 'c' of 'a' and 'b'.
yy(1:,1)=b
do j=2,n
yy(2:,j)=yy(1:n-1,j-1)
yy(1,j)=yy(n,j-1)
end do
c=yy .x. a
! Compute f=inverse(transform(a)*transform(b)).
fa = fft(a)
fb = fft(b)
f=ifft(fa*fb)

```
```

! Check the Convolution Theorem:
! inverse(transform(a)*transform(b)) = convolution(a,b).
err = norm(c-f)/norm(c)
if (err <= sqrt(epsilon(one))) then
write (*,*) 'Example 4 for FAST_DFT (operators) is correct.'
end if
end

```

\section*{IFFT_BOX}

Computes the inverse Discrete Fourier Transform of several complex or real sequences.

\section*{Function Return Value}

Complex array containing the inverse of the Discrete Fourier Transform of the sequences in X . If \(x\) is an assumed shape complex array of rank 2,3 or 4 , the result is a complex array of the same shape and rank consisting of the inverse DFT for each of the last rank's indices.
(Output)

\section*{Required Argument}
\(\boldsymbol{X}\) - Box containing the sequences for which the inverse transform is to be computed. x is an assumed shape complex array of rank 2,3 or 4 . If \(x\) is real or double, it is converted to complex internally prior to the computation. (Input)

\section*{Optional Arguments, Packaged Options}

WORK - A COMPLEX array of the same precision as the data. For rank-1 transforms the size of WORK is \(n+15\). To define this array for each problem, set WORK \((1)=0\). Each additional rank adds the dimension of the transform plus 15 . Using the optional argument WORK increases the efficiency of the transform.

The option and derived type names are given in the following tables:
\begin{tabular}{|l|l|}
\hline Option Names for IFFT & Option Value \\
\hline Options_for_fast_dft & 1 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_ifft_box_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_ifft_box_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See FAST_DFT located in Chapter 6, "Transforms" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

IFFT_BOX (X [, ...])

\section*{Description}

Computes the inverse of the Discrete Fourier Transform of a box of complex sequences. This function uses FAST_DFT, FAST_2DFT, and FAST_3DFT from Chapter 6.

\section*{Parallel Example}
```

    use rand_gen_int
    use fft_boo_int
    use ifft_box_int
    use linear_operators
    use mpi_setup_int
    implicit none
    ! This is FFT_BOX example.
integer i,j
integer, parameter :: n=40, nr=4
real(kind(le0)) :: err(nr), one=1e0
real(kind(le0)) :: a (n,1,nr), b (n, nr), c(n,1,nr), yy(n,n,nr)
complex(kind(le0)), dimension(n,nr) :: f, fa, fb, cc, aa
real(kind(1e0)), parameter::zero_par=0.e0
real(kind(1e0))::dummy_par(0)
integer iseed_par
type(s_options)::iopti_par(2)
! setup for MPI
MP_NPROCS = MP_SETUP()
! Set Random Number generator seed
iseed_par = 53976279
iopti_par(1)=s_options(s_rand_gen_generator_seed,zero_par)
iopti_par(2)=s_options(iseed_\overline{par,}\mathrm{ zero_par)}

```
```

    call rand gen(dummy par,iopt=iopti_par)
    ! Generate two random periodic sequences 'a' and 'b'.
a=rand(a); b=rand(b)
! Compute the convolution 'c' of 'a' and 'b'.
do i=1,nr
aa(1:,i) = a(1:,1,i)
YY(1:,1,i)=b(1:,i)
do j=2,n
yy(2:,j,i)=yy(1:n-1,j-1,i)
YY(1,j,i)=yY(n,j-1,i)
end do
end do
c=yY .x. a
! Compute f=inverse(transform(a)*transform(b)).
fa = fft_box(aa)
fb = fft box(b)
f=ifft_box(fa*fb)
! Check the Convolution Theorem:
! inverse(transform(a)*transform(b)) = convolution(a,b).
do i=1,nr
cc(1:,i) = c(1:,1,i)
end do
err = norm(cc-f)/norm(cc)
if (ALL(err <= sqrt(epsilon(one))).AND. MP_RANK == 0) then
write (*,*) 'FFT_BOX is correct.'
end if
MP_NPROCS = MP_SETUP('Final')
end

```

\section*{isNaN}

Tests for NaN .

\section*{Function Return Value}

Logical indicating whether or not A contains NaN. The output value tests . true. only if there is at least one NaN in the scalar or array. (Output)

\section*{Required Argument}

A - The argument can be a scalar or array of rank-1, rank-2 or rank-3. The values can be any of the four intrinsic floating-point types. (Input)

\section*{FORTRAN 90 Interface}
isNaN( A)

\section*{Description}

This is a generic logical function used to test scalars or arrays for occurrence of an IEEE 754 Standard format of floating point (ANSI/IEEE 1985) NaN, or not-a-number. Either quiet or signaling NaNs are detected without an exception occurring in the test itself. The individual array entries are each examined, with bit manipulation, until the first NaN is located. For non-IEEE formats, the bit pattern tested for single precision is transfer \((\operatorname{not}(0), 1)\). For double precision numbers x , the bit pattern tested is equivalent to assigning the integer array \(i(1: 2)=\operatorname{not}(0)\), then testing this array with the bit pattern of the integer array transfer \((x, i)\). This function is likely to be required whenever there is the possibility that a subroutine blocked the output with NaNs in the presence of an error condition.

\section*{Example}
```

    use isnan_int
    implicit none
    ! This is the equivalent of Example 1 for NaN.
integer, parameter :: n=3
real(kind(le0)) A(n,n); real(kind(1d0)) B (n,n)
real(kind(1e0)), external :: s_NaN
real(kind(ld0)), external :: d_NaN
! Assign NaNs to both A and B:
A = s_Nan(1e0); B = d_Nan(1d0)
! Check that NaNs are noted in both A and B:
if (isNan(A) .and. isNan(B)) then
write (*,*) 'Example 1 for NaN is correct.'
end if
end

```

\section*{NaN}

Returns the value for NaN .

\section*{Function Return Value}

Returns, as a scalar, a value corresponding to the IEEE 754 Standard format of floating point (ANSI/IEEE 1985) for NaN. For other floating point formats a special pattern is returned that tests .true. using the function isNaN. (Output)

\section*{Required Argument}
\(\boldsymbol{X}\) - Scalar value of the same type and precision as the desired result, NaN . This input value is used only to match the type of output. (Input)

\section*{FORTRAN 90 Interface}

\author{
NaN (A)
}

\section*{Description}

NaN returns, as a scalar, a value corresponding to the IEEE 754 Standard format of floating point (ANSI/IEEE 1985) for NaN .

The bit pattern used for single precision is transfer ( \(\operatorname{not}(0), 1\) ). For double precision, the bit pattern for single precision is replicated by assigning the temporary integer array \(i(1: 2)=\operatorname{not}(0)\), and then using the double-precision bit pattern transfer(i,x) for the output value.

\section*{Example}

Arrays are assigned all NaN values, using single and double-precision formats. These are tested using the logical function routine, isNaN.
```

    use isnan_int
    implicit none
    ! This is the equivalent of Example 1 for NaN.
integer, parameter :: n=3
real(kind(le0)) A(n,n); real(kind(1d0)) B(n,n)
real(kind(le0)), external :: s_NaN
real(kind(1d0)), external :: d_NaN
! Assign NaNs to both A and B:
A = s_Nan(1e0); B = d_Nan(1d0)
! Check that NaNs are noted in both A and B:
if (isNan(A) .and. isNan(B)) then
write (*,*) 'Example 1 for NaN is correct.'
end if
end

```

\section*{NORM}

Computes the norm of an array.

\section*{Function Return Value}

Norm of A. This is a scalar for the case where A is rank 1 or rank 2. For rank-3 arrays, the norms of each rank-2 array, in dimension 3, are computed. (Output)

\section*{Required Argument}

A - An array of rank-1, rank-2, or rank-3, containing the values for which the norm is to be computed. It may be real, double, complex, or double complex. (Input)

\section*{Optional Arguments, Packaged Options}

TYPE - Integer indicating the type of norm to be computed.
\(1=l_{1}\)
\(2=l_{2} \quad\) (default)
huge (1) \(=l_{\infty}\)

Use of the option number ?_reset_default_norm will switch the default from the \(l_{2}\) to the \(l_{1}\) or \(l_{\infty}\) norms. (Input)

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for NORM & Option Value \\
\hline ?_norm_for_lin_sol_svd & 1 \\
?_reset_default_norm & 2 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{l} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_norm_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_norm_options_once ( : ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_SVD located in Chapter 1, "Linear Systems" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

NORM (A [,...])

\section*{Description}

Computes the \(l_{2}, l_{1}\) or \(l_{\infty}\) norm. The \(l_{1}\) and \(l_{\infty}\) norms are likely to be less expensive to compute than the \(l_{2}\) norm.
\[
\begin{aligned}
& \|A\|_{1}=\max _{j}\left(\sum_{i=1}^{m}\left|a_{i j}\right|\right) \\
& \|A\|_{2}=s_{1}=\text { largest singular value } \\
& \|A\|_{\infty \leftrightarrow \text { huge }(1)}=\max _{i}\left(\sum_{j=1}^{n}\left|a_{i j}\right|\right)
\end{aligned}
\]

If the \(l_{2}\) norm is required, this function uses LIN_SOL_SVD (see Chapter 1, "Linear Systems"), to compute the largest singular value of \(A\). For the other norms, Fortran 90 intrinsics are used.

\section*{Examples}
```

Compute three norms of an array
use norm_int
real (kīnd(1e0)) A(5), n_1, n_2, n_inf
A = rand (A)
! I1
n_1 = norm(A, TYPE=1)
write (*,*) n_1
! I2
n_2 = norm(A)
write (*,*) n_2
! I infinity
n_inf = norm(A, TYPE=huge(1))
write (*,*) n_inf
end

```

\section*{Parallel Example (parallel_ex14.f90)}

A "Polar Decomposition" of several matrices are computed. The box data type and the SVD () function are used. Orthogonality and small residuals are checked to verify that the results are correct.
```

use linear operators
use mpi_setup_int
implicit none
! This is Parallel Example 15 using operators and
! functions for a polar decomposition.
integer, parameter :: n=33, nr=3
real(kind(1d0)) :: one=1d0, zero=0d0
real(kind(1d0)),dimension(n,n,nr) :: A, P, Q, \&
S_D(n,nr), U_D, V_D
real(kind(1d0)) TEMP1(nr), TEMP2(nr)
! Setup for MPI:
mp_nprocs = mp_setup()
! Generate a random matrix.
if(mp_rank == 0) A = rand(A)
! Compute the singular value decomposition.

```
```

    S_D = SVD(A, U=U D, V=V D)
    ! Compute the (left) orthogonal factor.
P = U_D . xt. V_D
! Compute the (right) self-adjoint factor.
Q = V_D .x. diag(S_D) .xt. V_D
! Check the results for orthogonality and
! small residuals.
TEMP1 = NORM(spread(EYE (n),3,nr) - (p .xt. p))
TEMP2 = NORM(A - (P . X. Q)) / NORM(A)
if (ALL(TEMP1 <= sqrt(epsilon(one))) .and. \&
ALL(TEMP2 <= sqrt(epsilon(one)))) then
if(mp_rank == 0)\&
write- (*,*) 'Parallel Example 15 is correct.'
end if
! See to any error messages and exit MPI.
mp_nprocs = mp_setup('Final')
end

```

\section*{ORTH}


Orthogonalizes the columns of a matrix.

\section*{Function Return Value}

Orthogonal matrix \(Q\) from the decomposition \(A=Q R\). If \(A\) is rank \(3, Q\) is rank-3. (Output)

\section*{Required Argument}

A- Matrix A to be decomposed. Must be an array of rank-2 or rank-3 (box data) of type real, double, complex, or double complex. (Input)

\section*{Optional Arguments, Packaged Options}
\(\boldsymbol{R}\) - Upper-triangular or upper trapezoidal matrix \(R\), from the QR decomposition. If this optional argument is present, the decomposition is complete. If A is rank-3, \(R\) is rank-3. (Output)

The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Name for ORTH & Option Value \\
\hline Skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_orth_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_orth_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes".

\section*{FORTRAN 90 Interface}

ORTH (A \([, \ldots]\) )

\section*{Description}

Orthogonalizes the columns of a matrix. The decomposition \(A=Q R\) is computed using a forward and backward sweep of the Modified Gram-Schmidt algorithm.

\section*{Examples}

\section*{(Operator_ex19.f90)}
```

use linear_operators
use \overline{lin_sol_tri_int}
use ran\overline{d_int}
use Numerical_Libraries
implicit none
! This is the equivalent of Example 3 (using operators) for LIN_SOL_TRI.
integer i, nopt
integer, parameter :: n=128, k=n/4, ncoda=1, lda=2
real(kind(1e0)), parameter :: s_one=1e0, s_zero=0e0
real(kind(1e0)) A(lda,n), EVAL(k)
type(s_options) :: iopt(2)
real(kind(le0)) d(n), b(n), d_t(2*n,k), c_t(2*n,k), perf_ratio, \&
b_t(2*n,k), y_t (2*n,k), \overline{eval_t (k), rēs (n,k)}
logical small
! This flag is used to get the k largest eigenvalues.
small = .false.
! Generate the main diagonal and the co-diagonal of the
! tridiagonal matrix.
b=rand(b); d=rand(d)
A(1,1:) =b; A (2,1:)=d
! Use Numerical Libraries routine for the calculation of k
! largest eigenvalues.
CALL EVASB (N, K, A, LDA, NCODA, SMALL, EVAL)
EVAL_T = EVAL

```
```

! Use IMSL Fortran Numerical Librarytridiagonal solver for inverse iteration
! calculation of eigenvectors.
factorization_choice: do nopt=0,1
! Create k tridiagonal problems, one for each inverse
! iteration system.
b_t(1:n,1:k) = spread(b,DIM=2,NCOPIES=k)
c_t(1:n,1:k) = EOSHIFT(b_t(1:n,1:k),SHIFT=1,DIM=1)
d_t(1:n,1:k) = spread(d,DIM=2,NCOPIES=k) - \&
spread(EVAL_T,DIM=1,NCOPIES=n)
! Start the right-hand side at random values, scaled downward
! to account for the expected 'blowup' in the solution.
y_t=rand (y_t)
! Do two iterations for the eigenvectors.
do i=1, 2
y_t(1:n,1:k) = y_t(1:n,1:k)*epsilon(s_one)
call lin_sol_tri(c_t, d_t, b_t, y_t, \&
iopt=iopt)
iopt(nopt+1) = s_lin_sol_tri_solve_only
end do
! Orthogonalize the eigenvectors. (This is the most
! intensive part of the computing.)
y_t(1:n,1:k) = ORTH(y_t(1:n,1:k))
! See if the performance ratio is smaller than the value one.
! If it is not the code will re-solve the systems using Gaussian
! Elimination. This is an exceptional event. It is a necessary
! complication for achieving reliable results.
res(1:n,1:k) = spread(d,DIM=2,NCOPIES=k)*y_t(1:n,1:k) + \&
spread(b,DIM=2,NCOPIES=k) * \&
EOSHIFT(y_t(1:n,1:k),SHIFT=-1,DIM=1) + \&
EOSHIFT(sp}read(b,DIM=2,NCOPIES=k)*y_t(1:n,1:k),SHIFT=1) \&
- y_t(1:n,1:k)*spread(EVAL_T(1:k),DIM=1,NCOPIES=n)
! If the factorization method is Cyclic Reduction and perf_ratio is
! larger than one, re-solve using Gaussian Elimination. If the
! method is already Gaussian Elimination, the loop exits
! and perf_ratio is checked at the end.
perf_ratio = norm(res(1:n,1:k),1) / \&
norm(EVAL_T(1:k),1) / \&
epsilon(s_one) / (5*n)
if (perf_ratio <= s_one) exit factorization_choice
iopt(nop\overline{t}+1) = s_lin__sol_tri_use_Gauss_elim
end do factorization_choice
if (perf_ratio <= s_one) then
write- (*,*) 'Example 3 for LIN_SOL_TRI (operators) is correct.'
end if

```

> end

\section*{Parallel Example}
```

use linear_operators
use mpi_setup_int
integer, parameter : : N=32, nr=4
real (kind(1.e0)) A(N,N,nr), Q(N,N,nr)
! Setup for MPI
mp_nprocs = mp_setup()
if (mp_rank == 0) then
A = rand (A)
end if
Q = orth(A)
mp_nprocs = mp_setup ('Final')
end

```

\section*{RAND}

Generates a scalar, rank-1, rank-2 or rank-3 array of random numbers.

\section*{Function Return Value}

Scalar, rank-1, rank-2 or rank-3 array of random numbers. The output function value matches the input argument A in type, kind and rank. For complex arguments, the output values will be real and imaginary parts with random values of the same type, kind, and rank. (Output)

\section*{Required Argument}

A - The argument must be a scalar, rank-1, rank-2, or rank-3 array of type single, double, complex, or double complex. Used only to determine the type and rank of the output. (Input)

\section*{Optional Arguments, Packaged Options}

Note: If any of the arrays s_rand_options (:), s_rand_options_once (:), d_rand_options(:), or d_rand_options_once (: ) are allocated, they are passed as arguments to rand_gen using the keyword "iopt=".

The option and derived type names are given in the following table:
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_rand_options (:) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_rand_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

\section*{FORTRAN 90 Interface}

RAND (A)

\section*{Description}

Generates a scalar, rank-1, rank-2 or rank-3 array of random numbers. Each component number is positive and strictly less than one in value.
This function uses rand_gen to obtain the number of values required by the argument. The values are then copied using the RESHAPE intrinsic

\section*{Example}
```

    use show_int
    use rand_int
    implicit none
    ! This is the equivalent of Example 1 for SHOW.
integer, parameter :: n=7, m=3
real(kind(1e0)) s_x(-1:n), s_m(m,n)
real(kind(1d0)) d x(n), d_m(m,n)
complex(kind(le0)) c_x(n), c c_m(m,n)
complex(kind(1d0)) z_x(n), z_m(m,n)
integer i_x(n), i_m(m,n)
type (s options) ōptions(3)
! The data types printed are real(kind(le0)), real(kind(1d0)),
! complex(kind(1e0)), complex(kind(1d0)), and INTEGER. Fill with random
! numbers and then print the contents, in each case with a label.
s_x=rand(s_x); s_m=rand(s_m)
d_x=rand (d_x); d_m=rand (d_m)
c_x=rand(c_x); c_m=rand (c_m)
z_x=rand (z_x); z_m=rand (z_m)
i_x=100*rand(s_x(1:n)); i_m=100*rand(s_m)
call show (s_x, 'Rank-1, REAL')
call show (s_m, 'Rank-2, REAL')
call show (d_x, 'Rank-1, DOUBLE')
call show (d_m, 'Rank-2, DOUBLE')
call show (c_x, 'Rank-1, COMPLEX')
call show (c_m, 'Rank-2, COMPLEX')
call show (z_x, 'Rank-1, DOUBLE COMPLEX')
call show (z m, 'Rank-2, DOUBLE COMPLEX')

```
```

    call show (i x, 'Rank-1, INTEGER')
    call show (i_m, 'Rank-2, INTEGER')
    ! Show 7 digits per number and -1 according to the
! natural or declared size of the array.
options(1)=show_significant_digits_is_7
options(2)=show_starting_index_is
options(3)= -1 ! The starting -1 value.
call show (s x, \&
'Rank-1, REAL with 7 ' digits, natural indexing', IOPT=options)
end

```

\section*{RANK}

Computes the mathematical rank of a matrix.

\section*{Function Return Value}

Integer rank of \(A\). The output function value is an integer with a value equal to the number of singular values that are greater than a tolerance. (Output)

\section*{Required Argument}
\(\boldsymbol{A}\) - Matrix for which the rank is to be computed. The argument must be rank-2 or rank-3 (box) array of type single, double, complex, or double complex. (Input)

\section*{Optional Arguments, Packaged Options}

This function uses LIN_SOL_SVD to compute the singular values of the argument. The singular values are then compared with the value of the tolerance to compute the rank.

The option and derived type names are given in the following tables
\begin{tabular}{|l|c|}
\hline Option Names for RANK & Option Value \\
\hline ?_rank_set_small & 1 \\
?_rank_for_lin_sol_svd & 2 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & Derived Type \\
\hline ?_rank_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_rank_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SOL_SVD located in Chapter 1, "Linear Systems" for the specific options for this routine.

\section*{FORTRAN 90 Interface}

RANK (A)

\section*{Description}

Computes the mathematical rank of a rank-2 or rank-3 array. The output function value is an integer with a value equal to the number of singular values that are greater than a tolerance. The default value for this tolerance is \(\varepsilon^{1 / 2} s_{1}\), where \(\varepsilon\) is machine precision and \(s_{1}\) is the largest singular value of the matrix.

\section*{Examples}
```

use linear_operators
real (kind(1e0)) A (5,5)
A = rand (A)
write (*,*) rank(A)
A=1.0
write (*,*) rank(A)
end

```

\section*{Output}

5
1

\section*{Parallel Example}
```

use linear_operators
use mpi_setup_int
integer, parameter :: N=3, nr=4
integer r(nr)
real (kind(1.e0)) s_mat(N,N), s_box(N,N,nr)
! Setup for MPI
mp_nprocs = mp_setup()
if (mp_rank == 0) then
s_mat = reshape((/1.,0.,0.,epsilon(1.0e0)/), (/n,n/))
s_box = spread(s_mat,dim=3,ncopies=nr)
end i\overline{f}

```
```

    r = rank(s_box)
    mp_nprocs = mp_setup ('Final')
    end
    ```

\section*{SVD}

\section*{岛)MPI}

Computes the singular value decomposition of a matrix, \(A=U S V^{T}\).

\section*{Function Return Value}
\(m x n\) diagonal matrix of singular values, S , from the \(A=U S V^{T}\) decomposition. (Output)

\section*{Required Argument}

A - Array of size \(m x n\) to be decomposed. Must be rank-2 or rank-3 array of type single, double, complex, or double complex. (Input)

\section*{Optional Arguments, Packaged Options}
\(\boldsymbol{U}\) - Array of size \(m \times m\) containing the singular vectors U. (Output)
\(\boldsymbol{V}\) - Array of size \(n \times n\) containing the singular vectors V. (Output)
The option and derived type names are given in the following tables:
\begin{tabular}{|l|c|}
\hline Option Names for SVD & Option Value \\
\hline Options_for_lin_svd & 1 \\
Options_for_lin_sol_svd & 2 \\
skip_error_processing & 5 \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|}
\hline \begin{tabular}{c} 
Name of Unallocated Option Array \\
to Use for Setting Options
\end{tabular} & \multicolumn{1}{c|}{ Use } & \multicolumn{1}{c|}{ Derived Type } \\
\hline ?_svd_options (: ) & \begin{tabular}{l} 
Use when setting options for \\
calls hereafter.
\end{tabular} & ?_options \\
?_svd_options_once (: ) & \begin{tabular}{l} 
Use when setting options for \\
next call only.
\end{tabular} & ?_options \\
\hline
\end{tabular}

For a description on how to use these options, see "Matrix Optional Data Changes". See LIN_SVD and LIN_SOL_SVD located in Chapter 1, "Linear Systems" for the specific options for these routines.

\section*{FORTRAN 90 Interface}
\(\operatorname{SVD}(\mathrm{A}[, \ldots])\)

\section*{Description}

Computes the singular value decomposition of a rank-2 or rank-3 array, \(A=U S V^{T}\).
This function uses one of the routines LIN_SVD and LIN_SOL_SVD. If a complete decomposition is required, LIN_SVD is used. If singular values only, or singular values and one of the right and left singular vectors are required, then LIN_SOL_SVD is called.

\section*{Examples}

\section*{operator_ex14.f90}
```

    use linear_operators
    ```
    implicit nōne
```

! This is the equivalent of Example 2 for LIN_SOL_SVD using operators
! and functions.
integer, parameter :: n=32
real(kind(1d0)) :: one=1d0, zero=0d0
real(kind(1d0)) A(n,n), P(n,n), Q(n,n), \&
S_D(n), U_D(n,n), V_D (n,n)
! Generate a random matrix.
A = rand(A)
! Compute the singular value decomposition.
S_D = SVD (A, U=U_D, V=V_D)
! Compute the (left) orthogonal factor.
P = U_D . xt. V_D
! Compute the (right) self-adjoint factor.
Q = V_D .x. diag(S_D) .xt. V_D
! Check the results.
if (norm( EYE(n) - (P .xt. P)) \&
<= sqrt(epsilon(one))) then
if (norm(A - (P .x. Q))/norm(A) \&
<= sqrt(epsilon(one))) then
write (*,*) 'Example 2 for LIN_SOL_SVD (operators) is correct.'
end if
end if
end

```

\section*{Parallel Example (parallel_ex14.f90)}

Systems of least-squares problems are solved, but now using the SVD () function. A box data type is used. This is an example which uses optional arguments and a generic function overloaded for parallel execution of a box data type. Any number of nodes can be used.
```

    use linear operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 14
! for SVD, .tx. , .x. and NORM.
integer, parameter :: m=128, n=32, nr=4
real(kind(1d0)) :: one=1d0, err(nr)
real(kind(1d0)) A(m,n,nr), b (m,1,nr), x(n,1,nr), U(m,m,nr), \&
V(n,n,nr), S(n,nr), g(m,1,nr)
! Setup for MPI:
mp_nprocs=mp_setup()
if(mp_rank == 0) then
! Generate \overline{a}}\mathrm{ random matrix and right-hand side.
A = rand(A); b = rand(b)
endif
! Compute the least-squares solution matrix of Ax=b.
S = SVD (A, U = U, V = V)
g = U .tx. b
x = V .x. (diag(one/S) .x. g(1:n,:,:))
! Check the results.
err = norm(A .tx. (b - (A .x. x)))/(norm(A) +norm(x))
if (ALL(err <= sqrt(epsilon(one)))) then
if(mp_rank == 0) \&
write (*,*) 'Parallel Example 14 is correct.'
end if
! See to any error messages and quit MPI
mp_nprocs = mp_setup('Final')
end

```

\section*{UNIT}

Normalizes the columns of a matrix so each has Euclidean length of value one.

\section*{Function Return Value}

Matrix containing the normalized values of A. The output function value is an array of the same type and kind as A, where each column of each rank-2 principal section has Euclidean length of value one (Output)

\section*{Required Argument}

A - Matrix to be normalized. The argument must be a rank-2 or rank-3 array of type single, double, complex, or double complex. (Input)

\section*{FORTRAN 90 Interface}

UNIT (A)

\section*{Description}

Normalizes the columns of a rank-2 or rank-3 array so each has Euclidean length of value one.
This function uses a rank-2 Euclidean length subroutine to compute the lengths of the nonzero columns, which are then normalized to have lengths of value one. The subroutine carefully avoids overflow or damaging underflow by rescaling the sums of squares as required.

\section*{Example (operator_ex28.f90)}
```

use linear_operators
implicit none
! This is the equivalent of Example 4 (using operators) for LIN_EIG_SELF.
integer, parameter :: n=64
real(kind(1e0)), parameter :: one=1d0
real(kind(1e0)), dimension(n,n) :: A, B, C, D(n), lambda(n), \&
S(n), vb_d, X, res
! Generate random self-adjoint matrices.
A = rand(A); A = A + .t.A
B = rand(B); B = B + .t. B
! Add a scalar matrix so B is positive definite.
B = B + norm(B)*EYE(n)
! Get the eigenvalues and eigenvectors for B.
S = EIG(B,V=vb_d)
! For full rank problems, convert to an ordinary self-adjoint
! problem. (All of these examples are full rank.)
if (S(n) > epsilon(one)) then
D = one/sqrt(S)
C = diag(D) .x. (vb_d .tx. A .x. vb_d) .x. diag(D)
C = (C + .t.C)/2
! Get the eigenvalues and eigenvectors for C.
lambda = EIG(C,v=X)
! Compute and normalize the generalized eigenvectors.
X = UNIT(vb d .x. diag(D) .x. X)
res = (A .x. X) - (B .x. X .x. diag(lambda))
! Check the results.
if(norm(res)/(norm(A) +norm(B)) <= \&
sqrt(epsilon(one))) then
write (*,*) 'Example 4 for LIN_EIG_SELF (operators) is correct.'
end if

```
```

end if
end

```

\section*{Chapter 11: Utilities}

\section*{Routines}
11.1. ScaLAPACK Utilities
Sets up a processor grid and calculates default values for use
in mapping arrays to the processor grid ScaLAPACK_SETUP ..... 1622
Calculates the array dimensions needed for local arrays ScaLAPACK_GETDIM ..... 1624
Reads matrix data from a file and transmits it into the
two-dimensional block-cyclic form ScaLAPACK_READ1625
Writes the matrix data to a file ScaLAPACK_WRITE ..... 1627
Reads matrix data from an array and transmits it into the two-dimensional block-cyclic form ScaLAPACK_MAP ..... 1636
Writes the matrix data to a global array....... ScaLAPACK_UNMAP ..... 1637
Exits ScaLAPACK usage ScaLAPACK_EXIT ..... 1640
11.2. Print
Prints error messages ERROR_POST ..... 1640
Prints rank-1 or rank-2 arrays of numbers in a readable format ..... SHOW ..... 1643
Real rectangular matrix
with integer row and column labels WRRRN ..... 1647
Real rectangular matrix with given format and labels WRRRL ..... 1649
Integer rectangular matrix with integer row and column labels WRIRN ..... 1653
Integer rectangular matrix with given format and labels WRIRL ..... 1655
Complex rectangular matrix
with row and column labels WRCRN ..... 1658
Complex rectangular matrix with given format and labels WRCRL ..... 1660
Sets or retrieves options for printing a matrix WROPT ..... 1664
Sets or retrieves page width and length PGOPT ..... 1671
11.3. Permute
Elements of a vector PERMU ..... 1673
Rows/columns of a matrix. PERMA ..... 1674
11.4. Sort
Sorts a rank-1 array of real numbers \(x\) so the \(y\) results
are algebraically nondecreasing, \(y_{1} \leq y_{2} \leq \ldots y_{n} \ldots \ldots\). SORT_REAL ..... 1677
Real vector by algebraic value SVRGN ..... 1679
Real vector by algebraic value and permutations returned SVRGP ..... 1681
Integer vector by algebraic value ..... SVIGN ..... 1683
Integer vector by algebraic value and permutations returned SVIGP ..... 1684
Real vector by absolute value SVRBN ..... 1685
Real vector by absolute value and permutations returned SVRBP ..... 1687
Integer vector by absolute value SVIBN ..... 1688
Integer vector by absolute value
and permutations returned SVIBP ..... 1690
11.5. Search
Sorted real vector for a number SRCH ..... 1691
Sorted integer vector for a number ISRCH ..... 1694
Sorted character vector for a string SSRCH ..... 1696
11.6. Character String Manipulation
Gets the character corresponding to agiven ASCII value.ACHAR 1698
Get the integer ASCII value for a given character IACHAR ..... 1699
Gets upper case integer ASCII value for a character .ICASE ..... 1700
Case-insensitive version comparing two strings IICSR ..... 1701
Case-insensitive version of intrinsic function INDEX IIDEX ..... 1703
Converts a character string with digits to an integer CVTSI ..... 1704
11.7. Time, Date, and Version
CPU time CPSEC ..... 1705
Time of day TIMDY ..... 1706
Today's date ..... TDATE ..... 1707
Number of days from January 1, 1900, to the given date ... ..... 1708
Date for the number of days from January 1, 1900 NDYIN ..... 1710
Day of week for given date ..... 1711
IDYWK
Version, system, and serial numbers ..... 1713
11.8. Random Number Generation
Generates a rank-1 array of random numbers. RAND GEN ..... 1714
Retrieves the current value of the seed RNGET ..... 1722
Initializes a random seed RNSET ..... 1723
Selects the uniform \((0,1)\) generator RNOPT ..... 1724
Initializes the 32-bit Merseene Twister generator using an array RNIN32 ..... 1726
Retrieves the current table used in the 32-bit
Mersenne Twister generator RNGE32 ..... 1727
Sets the current table used in the 32-bit
Mersenne Twister generator RNSE32 ..... 1729
Initializes the 32-bit Merseene Twister generator using an array RNIN64 ..... 1729
Retrieves the current table used in the 64-bit Mersenne Twister generator ..... RNGE64 ..... 1730
Sets the current table used in the 64-bit
Mersenne Twister generator ..... RNSE64 ..... 1732
Generates pseudorandom numbers (function form). ..... RNUNF ..... 1732
Generates pseudorandom numbers RNUN ..... 1734
11.9 Low Discrepancy SequencesShuffled Faure sequence initializationFAURE_INIT1736
Frees the structure containing information about the Faure sequence FAURE_FREE ..... 1736
Computes a shuffled Faure sequence FAURE_NEXT ..... 1737
11.10. Options Manager
Gets and puts type integer options IUMAG ..... 1739
Gets and puts type Real options ..... UMAG ..... 1743
Gets and puts type DOUBLE PRECISION options... SUMAG/DUMAG ..... 1746
11.11. Line Printer Graphics
Prints plot of up to 10 sets of points PLOTP ..... 1746
11.12. Miscellaneous
Decomposes an integer into its prime factors .PRIME ..... 1749
Returns mathematical and physical constants CONST ..... 1751
Converts a quantity to different units ..... 1753
Computes \(\sqrt{\mathrm{a}^{2}+\mathrm{b}^{2}}\) without underflow or overflow HYPOT ..... 1757
Initializes or finalizes MPI. MP_SETUP ..... 1758

\section*{Usage Notes for ScaLAPACK Utilities}

For a detailed description of MPI Requirements see "Dense Matrix Parallelism Using MPI" in Chapter 10 of this manual.

This section describes the use of ScaLAPACK, a suite of dense linear algebra solvers, applicable when a single problem size is large. We have integrated usage of IMSL Fortran Library with ScaLAPACK. However, the ScaLAPACK library, including libraries for BLACS and PBLAS, are not part of this Library. To use ScaLAPACK software, the required libraries must be installed on the user's computer system. We adhered to the specification of Blackford, et al. (1997), but use only MPI for communication. The ScaLAPACK library includes certain \(L A P A C K\) routines, Anderson, et al. (1995), redesigned for distributed memory parallel computers. It is written in a Single Program, Multiple Data (SPMD) style using explicit message passing for communication. Matrices are laid out in a two-dimensional block-cyclic decomposition. Using High Performance

Fortran (HPF) directives, Koelbel, et al. (1994), and a static \(p \times q\) processor array, and following declaration of the array, \(A(*, *)\), this is illustrated by:
```

INTEGER, PARAMETER :: N=500, P= 2, Q=3, MB=32, NB=32
!HPF\$ PROCESSORS PROC(P,Q)
!HPF\$ DISTRIBUTE A(cyclic(MB), cyclic(NB)) ONTO PROC

```

Our integration work provides modules that describe the interface to the ScaLAPACK library. We recommend that users include these modules when using ScaLAPACK or ancillary packages, including BLACS and PBLAS. For the job of distributing data within a user's application to the block-cyclic decomposition required by \(S c a L A P A C K\) solvers, we provide a utility that reads data from an external file and arranges the data within the distributed machines for a computational step. Another utility writes the results into an external file. We also provide similar utilities that map/unmap global arrays to/from local arrays. These utilities are used in our ScaLAPACK examples for brevity.

The data types supported for these utilities are integer; single precision, real; double precision, real; single precision, complex; and double precision, complex.

A ScaLAPACK library normally includes routines for:
- the solution of full-rank linear systems of equations,
- general and symmetric, positive-definite, banded linear systems of equations,
- general and symmetric, positive-definite, tri-diagonal, linear systems of equations,
- condition number estimation and iterative refinement for LU and Cholesky factorization,
- matrix inversion,
- full-rank linear least-squares problems,
- orthogonal and generalized orthogonal factorizations,
- orthogonal transformation routines,
- reductions to upper Hessenberg, bidiagonal and tridiagonal form,
- reduction of a symmetric-definite, generalized eigenproblem to standard form,
- the self-adjoint or Hermitian eigenproblem,
- the generalized self-adjoint or Hermitian eigenproblem, and
- the non-symmetric eigenproblem

ScaLAPACK routines are available in four data types: single precision, real; double precision; real, single precision, complex, and double precision, complex. At present, the non-symmetric eigenproblem is only available in single and double precision. More background information and user documentation is available on the World Wide Web at location www.netlib.org/scalapack/slug/scalapack_slug.html.

For users with rank deficiency or simple constraints in their linear systems or least-squares problem, we have routines for:
- full or deficient rank least-squares problems with non-negativity constraints
- full or deficient rank least-squares problems with simple upper and lower bound constraints

These are available in two data types: single precision, real, and double precision, real, and they are not part of ScaLAPACK. The matrices are distributed in a general block-column layout.
We also provide generic interfaces to a number of ScaLAPACK routines through the standard IMSL Library routines. These are listed in Table D in the Introduction of this manual.
The global arrays which are to be distributed across the processor grid for use by the ScaLAPACK routines require that an array descriptor be defined for each of them. We use the ScaLAPACK TOOLS routine DESCINIT to set up array descriptors in our examples. A typical call to DESCINIT:

CALL DESCINIT (DESCA, M, N, MB, NB, IRSRC, ICSRC, ICTXT, LLD, INFO)
Where the arguments in the above call are defined as follows for the matrix being described:
DESCA - An input integer vector of length 9 which is to contain the array descriptor information.
\(\boldsymbol{M}\) - An input integer which indicates the row size of the global array which is being described.
\(N\) - An input integer which indicates the column size of the global array which is being described.

MB - An input integer which indicates the blocking factor used to distribute the rows of the matrix being described.
\(N B\) - An input integer which indicates the blocking factor used to distribute the columns of the matrix being described.

IRSRC - An input integer which indicates the processor grid row over which the first row of the array being described is distributed.

ICSRC - An input integer which indicates the processor grid column over which the first column of the array being described is distributed.

ICTXT - An input integer which indicates the BLACS context handle.
\(\boldsymbol{L L D}\) - An input integer indicating the leading dimension of the local array which is to be used for storing the local blocks of the array being described

INFO - An output integer indicating whether or not the call was successful. INFO =0 indicates a successful exit. \(\mathrm{INFO}=-\mathrm{i}\) indicates the i -th argument had an illegal value.

This call is equivalent to the following assignment statements:
```

DESCA(1) = 1 ! This is the descriptor
! type. In this case, 1.
DESCA(2) = ICTXT
DESCA(3) = M
DESCA(4) = N

```
```

DESCA(5) = MB
DESCA(6) = NB
DESCA(7) = IRSRC
DESCA(8) = ICSRC
DESCA(9) = LLD

```

The IMSL Library routines which interface with ScaLAPACK routines use IRSRC \(=0\) and ICSRC \(=0\) for the internal calls to DESCINIT.

\section*{ScaLAPACK Supporting Modules}

We recommend that users needing routines from ScaLAPACK, PBLAS or BLACS, Version 1.4, use modules that describe the interface to individual codes. This practice, including use of the declaration directive, IMPLICIT NONE, is a reliable way of writing ScaLAPACK application code, since the routines may have lengthy lists of arguments. Using the modules is helpful to avoid the mistakes such as missing arguments or mismatches involving Type, Kind or Rank (TKR). The modules are part of the Fortran Library product. There is a comprehensive module, ScaLAPACK_Support, that includes use of all the modules in the table below. This module decreases the number of lines of code for checking the interface, but at the cost of increasing source compilation time compared with using individual modules.
```

Module Name
ScaLAPACK_Support
ScaLAPACK_Int
PBLAS_Int
BLACS_Int
TOOLS Int
LAPACK_Int
ScaLAPACK_IO_Int
MPI_Node_Int
GRIDINFO_Int
ScaLAPACK_MAP_Int
ScaLAPACK_UNMAP_Int

```

\section*{Contents of the Module}

All of the following modules
All interfaces to ScaLAPACK routines
All interfaces to parallel \(B L A S\), or \(P B L A S\)
All interfaces to basic linear algebra communication routines, or BLACS
Interfaces to ancillary routines used by ScaLAPACK, but not in other packages
All interfaces to LAPACK routines required by ScaLAPACK
All interfaces to ScaLAPACK_READ, ScaLAPACK_WRITE utility routines. See this Chapter.

The module holding data describing the MPI communicator, MP_LIBRARY_WORLD. See Dense Matrix Parallelism Using MPI.
The module holding data describing the processor grid and information required to map the target array to the processors. See the Description section of SCALAPACK_SETUP below.

The interface to the ScaLAPACK_MAP utility routines.
The interface to the SCaLAPACK_UNMAP utility routines.

\section*{ScaLAPACK_SETUP}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine sets up a processor grid and calculates default values for various entities to be used in mapping a global array to the processor grid. All processors in the \(B L A C S\) context call the routine.

\section*{Required Arguments}
\(\boldsymbol{M}\) - The row dimension of the global array for which the local array dimensions are to be calculated. (Input)
\(N\) - The column dimension of the global array for which the local array dimensions are to be calculated. (Input)

NSQUARE -Input logical which indicates whether the block used for mapping the global array to the processor grid must be square. If the block must be square, set NSQUARE to .TRUE., otherwise, set it to .FALSE. (Input)

GRID1D - Input logical which indicates whether the processor grid is to be one dimensional or two dimensional. Set GRID1D to .TRUE. if the grid is to be one dimensional. Otherwise, set GRID1D to .FALSE. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_SETUP (M, N, NSQUARE, GRID1D)

\section*{Description}

Subroutine ScaLAPACK_SETUP creates a processor grid based on the number of processors being used and the GRID1D logical supplied by the user. The argument, NSQUARE, is supplied because some ScaLAPACK routines require that the row and column blocking factors be equal. GRID1D is supplied for those routines which require that the processor grid be one dimensional. ScaLAPACK_SETUP also establishes values for MP_M, MP_N, MP_NPROW, MP_NPCOL, MP_MB, MP_NB, MP_PIGRID, MP_ICTXT, MP_NSQUARE, and MP_GRID1D in the IMSL Fortran Library module GRIDINFO_INT. The above entities are defined as follows:
\(M P \_M\) - The row dimension of the primary array which is to be distributed among the processors.
\(M P \_N\) - The column dimension of the primary array which is to be distributed among the processors.

MP_NPROW - The number of rows in the processor grid.
MP_NPCOL - The number of columns in the processor grid.
\(M P\) _MB - The row blocking factor to be used in distributing the array.
MP_NB — The column blocking factor to be used in distributing the array.
MP_PIGRID - The pointer to the processor grid, MP_IGRID.

MP_ICTXT — The BLACS context ID associated with the processor grid.
MP_NSQUARE - Logical indicating whether or not the block used for mapping the global array to the processor grid must be square.

MP_GRID1D - Logical indicating whether or not the processor grid must be one dimensional.
GRIDINFO_INT is used by MPI_SETUP_INT so users do not need to explicitly use GRIDINFO_INT since they will be using MPI_SETUP_INT when they use MPI.

\section*{Example}

See ScaLAPACK_WRITE.

\section*{ScaLAPACK_GETDIM}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine calculates the row and column dimensions of a local distributed array based on the size of the array to be distributed and the row and column blocking factors to be used. All processors in the BLACS context call the routine.

\section*{Required Arguments}
\(\boldsymbol{M}\) - The row dimension of the global array for which the local array dimensions are to be calculated. (Input)
\(N\) - The column dimension of the global array for which the local array dimensions are to be calculated. (Input)
\(M B\) - The row blocking factor to be used in distributing the array. (Input)
\(N B\) - The column blocking factor to be used in distributing the array. (Input)
MXLDA - The row dimension of the local array. (Output)
MXCOL - The column dimension of the local array. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_GETDIM (M, N, MB, NB, MXLDA, MXCOL)

\section*{Description}

Subroutine ScaLAPACK_GETDIM calculates the row and column dimensions of a local array by using the ScaLAPACK utility numroc.

Note that ScaLAPACK_SETUP must be called prior to calling this routine because ScaLAPACK_GETDIM will use some of the global entities defined by ScaLAPACK_SETUP.

\section*{Example}

See ScaLAPACK_WRITE.

\section*{ScaLAPACK_READ}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine reads matrix data from a file and transmits it into the two-dimensional block-cyclic form required by ScaLAPACK routines. This routine contains a call to a barrier routine so that if one process is writing the file and an alternate process is to read it, the results will be synchronized.

All processors in the BLACS context call the routine.

\section*{Required Arguments}

File_Name - A character variable naming the file containing the matrix data. (Input) This file is opened with STATUS="OLD." If the name is misspelled or the file does not exist, or any access violation occurs, a type = terminal error message will occur. After the contents are read, the file is closed. This file is read with a loop logically equivalent to groups of reads:
```

READ() ((BUFFER(I,J), I=1,M), J=1, NB)
or (optionally):
READ() ((BUFFER(I,J), J=1,N), I=1, MB)

```

DESC_A \(\left(^{*}\right)\) — The nine integer parameters associated with the ScaLAPACK matrix descriptor. Values for NB,MB,LDA are contained in this array. (Input)
\(\boldsymbol{A}\left(\boldsymbol{L D A},{ }^{*}\right)\) - This is an assumed-size array, with leading dimension LDA, that will contain this processor's piece of the block-cyclic matrix. The data type for \(A\left({ }^{*}, *\right)\) is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; and double precision, complex. (Output)

\section*{Optional Arguments}

Format - A character variable containing a format to be used for reading the file containing matrix data. If this argument is not present, an unformatted or list-directed read is used. (Input)
iopt - Derived type array with the same precision as the array A (*,*), used for passing optional data to ScaLAPACK_READ. (Input)
The options are as follows:
\begin{tabular}{|l|l|l|}
\hline \multicolumn{3}{|c|}{ Packaged Options for ScaLAPACK_READ } \\
\hline Option Prefix =? & Option Name & Option Value \\
\hline S_, d_- & ScaLAPACK_READ_UNIT & 1 \\
\hline S_, d_ & ScaLAPACK_READ_FROM_PROCESS & 2 \\
\hline S_, d_ & ScaLAPACK_READ_BY_ROWS & 3 \\
\hline
\end{tabular}
iopt(IO) = ScaLAPACK_READ_UNIT
Sets the unit number to the value in iopt (IO +1 ) \%idummy. The default unit number is the value 11 .
iopt(IO) = ScaLAPACK_READ_FROM_PROCESS
Sets the process number that reads the named file to the value in iopt (IO +1 ) \%idummy. The default process number is the value 0 .
iopt(IO) = ScaLAPACK_READ_BY_ROWS
Read the matrix by rows from the named file. By default the matrix is read by columns.

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_READ (File_Name, DESC_A, A [,...])
Specific: The specific interface names are S_ScaLAPACK_READ and D_ScaLAPACK_READ.

\section*{Description}

Subroutine ScaLAPACK_READ reads columns or rows of a problem matrix so that it is usable by a ScaLAPACK routine. It uses the two-dimensional block-cyclic array descriptor for the matrix to place the data in the desired assumed-size arrays on the processors. The blocks of data are read, then transmitted and received. The block sizes, contained in the array descriptor, determines the data set size for each blocking send and receive pair. The number of these synchronization points is proportional to \(\lceil M \times N /(M B \times N B)\rceil\). A temporary local buffer is allocated for staging the matrix data. It is of size M by NB , when reading by columns, or N by MB , when reading by rows.

\section*{Example}

See ScaLAPACK_WRITE.

\section*{ScaLAPACK_WRITE}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine writes the matrix data to a file. The data is transmitted from the two-dimensional block-cyclic form used by ScaLAPACK. This routine contains a call to a barrier routine so that if one process is writing the file and an alternate process is to read it, the results will be synchronized. All processors in the \(B L A C S\) context call the routine.

\section*{Required Arguments}

File_Name - A character variable naming the file to receive the matrix data. (Input) This file is opened with "STATUS="UNKNOWN." If any access violation happens, a type \(=\) terminal error message will occur. If the file already exists it will be overwritten. After the contents are written, the file is closed. This file is written with a loop logically equivalent to groups of writes:
```

WRITE() ((BUFFER(I,J), I=1,M), J=1, NB)
or (optionally):
WRITE() ((BUFFER(I,J), J=1,N), I=1, MB)

```

DESC_A \(\left(^{*}\right)\) — The nine integer parameters associated with the ScaLAPACK matrix descriptor. Values for \(\mathrm{NB}, \mathrm{MB}, \mathrm{LDA}\) are contained in this array. (Input)
\(\boldsymbol{A}\left(\boldsymbol{L D A},{ }^{*}\right)\) - This is an assumed-size array, with leading dimension LDA, containing this processor's piece of the block-cyclic matrix. The data type for \(A(*, *)\) is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; or double precision, complex. (Input)

\section*{Optional Arguments}

Format -A character variable containing a format to be used for writing the file that receives matrix data. If this argument is not present, an unformatted or list-directed write is used. (Input)
iopt - Derived type array with the same precision as the array \(A(*, *)\), used for passing optional data to ScaLAPACK_WRITE. Use single precision when \(A(*, *)\) is type INTEGER. (Input) The options are as follows:
\begin{tabular}{|l|l|l|}
\hline \multicolumn{3}{|c|}{ Packaged Options for ScaLAPACK_WRITE } \\
\hline Option Prefix =? & Option Name & Option Value \\
\hline S_, d_ & ScaLAPACK_WRITE_UNIT & 1 \\
\hline S_, d_ & ScaLAPACK_WRITE_FROM_PROCESS & 2 \\
\hline S_, d_ & ScaLAPACK_WRITE_BY_ROWS & 3 \\
\hline
\end{tabular}
```

iopt(IO) =ScaLAPACK_WRITE_UNIT

```

Sets the unit number to the integer component of
iopt (IO +1 ) \%idummy. The default unit number is the value 11 .
iopt (IO) = ScaLAPACK_WRITE_FROM_PROCESS
Sets the process number that writes the named file to the integer component of iopt (IO +1 ) \%idummy. The default process number is the value 0 .
iopt(IO) = ScaLAPACK_WRITE_BY_ROWS
Write the matrix by rows to the named file. By default the matrix is written by columns.

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_WRITE (File_Name, DESC_A, A [, ...])
Specific: The specific interface names are S_ScaLAPACK_WRITE and D_ScaLAPACK_WRITE.

\section*{Description}

Subroutine ScaLAPACK_WRITE writes columns or rows of a problem matrix output by a ScaLAPACK routine. It uses the two-dimensional block-cyclic array descriptor for the matrix to extract the data from the assumed-size arrays on the processors. The blocks of data are transmitted and received, then written. The block sizes, contained in the array descriptor, determines the data set size for each blocking send and receive pair. The number of these synchronization points is proportional to \(\lceil M \times N /(M B \times N B)\rceil\). A temporary local buffer is allocated for staging the matrix data. It is of size M by NB , when writing by columns, or N by MB , when writing by rows.

\section*{Example 1: Distributed Transpose of a Matrix, In Place}

The program SCPK_EX1 illustrates an in-situ transposition of a matrix. An \(m \times n\) matrix, \(A\), is written to a file, by rows. The \(n \times m\) matrix, \(B=A^{T}\), overwrites storage for \(A\). Two temporary files are created and deleted. This algorithm for transposing a matrix is not efficient. It is used to illustrate the read and write routines and optional arguments for writing of data by matrix rows.
```

program scpk_exl
! This is Examp}le 1 for ScaLAPACK_READ and ScaLAPACK_WRITE.
! It shows in-situ or in-place transposition of a

```
```

! block-cyclic matrix.
USE ScaLAPACK_SUPPORT
USE ERROR_OPTION_PACKET
USE MPI_SETUP_IN\overline{T}
IMPLICIT NONE
INCLUDE "mpif.h"
INTEGER, PARAMETER :: M=6, N=6, NIN=10
INTEGER DESC_A(9), IERROR, INFO, I, J, K, L, MXLDA, MXCOL
LOGICAL :: GRID1D = .TRUE., NSQUARE = .TRUE.
real(kind(ldO)), allocatable :: A(:,:), AO(:,:)
real(kind(1d0)) ERROR
TYPE(d_OPTIONS) IOPT(1)
MP_NPROCS=MP_SETUP()
! Set up a 1D processor grid and define its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(M, N, NSQUARE, GRID1D)
! Get the array descriptor entities MXLDA, and MXCOL
CALL SCALAPACK_GETDIM(M, N, MP_MB, MP_NB, MXLDA, MXCOL)
! Set up the array descriptor
CALL DESCINIT(DESC_A, M, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
! Allocate space for local arrays
ALLOCATE (A0 (MXLDA,MXCOL))
! A root process is used to create the matrix data for the test.
IF(MP RANK == 0) THEN
ALLOCATE (A (M,N))
! Fill array with a pattern that is easy to recognize.
K=0
DO
K=K+1; IF(10**K > N) EXIT
END DO
DO J=1,N
DO I=1,M
! The values will appear, as decimals I.J, where I is
! the row and J is the column.
A (I,J) =REAL (I) +REAL (J) * 10d0** (-K)
END DO
END DO
OPEN(UNIT=NIN, FILE='test.dat', STATUS='UNKNOWN')
! Write the data by columns.
DO J=1,N,MP_NB
WRITE (NIN, * ) ((A (I,L), I=1,M), L=J,min(N,J+MP_NB-1))
END DO
CLOSE (NIN)
DEALLOCATE (A)
ALLOCATE (A (N,M))
END IF
! Read the matrix into the local arrays.
CALL ScaLAPACK_READ('test.dat', DESC_A, A0)

```
```

! To transpose, write the matrix by rows as the first step.
! This requires an option since the default is to write
! by columns.
IOPT(1)=ScaLAPACK_WRITE_BY_ROWS
CALL ScaLAPACK_WRITE("TEST.DAT", DESC_A, A0, IOPT=IOPT)
! Resize the local storage
DEALLOCATE (A0)
CALL SCALAPACK_GETDIM(N, M, MP_NB, MP_MB, MXLDA, MXCOL)
! Set up the array descriptor
! Reshape the descriptor for the transpose of the matrix.
! The number of rows and columns are swapped.
CALL DESCINIT(DESC_A, N, M, MP_NB, MP_MB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
ALLOCATE (A0 (MXLDA,MXCOL))
! Read the transpose matrix
CALL ScaLAPACK_READ("TEST.DAT", DESC_A, A0)
IF(MP RANK == 0) THEN
! Open the used files and delete when closed.
OPEN(UNIT=NIN, FILE='test.dat', STATUS='OLD')
CLOSE (NIN, STATUS='DELETE')
OPEN(UNIT=NIN, FILE='TEST.DAT', STATUS='OLD')
DO J=1,M,MP MB
READ (NIN, *}) ((A(I,L), I=1,N),L=J,min(M,J+MP_MB-1)
END DO
CLOSE (NIN, STATUS=' DELETE')
DO I=1,N
DO J=1,M
! The values will appear, as decimals I.J, where I is the row
! and J is the column.
A (I, J) =REAL (J) +REAL (I) * 10d0** (-K) - A(I,J)
END DO
END DO
ERROR=SUM (ABS (A))
END IF
! See to any error messages.
call e1pop("Mp_setup")
! Check results on just one process.
IF(ERROR <= SQRT(EPSILON(ERROR)) .and. \&
MP_RANK == 0) THEN
wríte(*,*) " Example 1 for BLACS is correct."
END IF
! Deallocate storage arrays and exit from BLACS.
IF(ALLOCATED(A)) DEALLOCATE (A)
IF(ALLOCATED(A0)) DEALLOCATE (A0)

```
```

! Exit from using this process grid.
CALL SCALAPACK_EXIT( MP_ICTXT )
! Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

Example 1 for BLACS is correct.

\section*{Additional Examples}

\section*{Example 2: Distributed Matrix Product with PBLAS}

The program SCPK_EX2 illustrates computation of the matrix product \(C_{m \times n}=A_{m \times k} B_{k \times n}\). The matrices on the right-hand side are random. Three temporary files are created and deleted. \(B L A C S\) and PBLAS are used. The problem size is such that the results are checked on one process.
```

    program scpk_ex2
    ! This is Example 2 for ScaLAPACK_READ and ScaLAPACK_WRITE.
! The product of two matrices is computed with PBLAS
! and checked for correctness.
USE ScaLAPACK_SUPPORT
USE MPI_SETUP_INT
IMPLICIT NONE
INCLUDE "mpif.h"
INTEGER, PARAMETER :: K=32, M=33, N=34, NIN=10
INTEGER INFO, IA, JA, IB, JB, IC, JC, MXLDA, MXCOL, MXLDB, \&
MXCOLB, MXLDC, MXCOLC, IERROR, I, J, L, \&
DESC_A(9), DESC_B(9), DESC_C(9)
LOGICA\overline{L : : GRID1D }}
real(kind(1d0)) :: ALPHA, BETA, ERROR=1d0, SIZE_C
real(kind(1d0)), allocatable, dimension(:, :) :: A,B,C,X(:),\&
A0, B0, C0
MP_NPROCS=MP_SETUP()
! Set up a 1D processor grid and define its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(M, N, NSQUARE, GRID1D)
! Get the array \overline{descriptor entities}
CALL SCALAPACK_GETDIM(M, K, MP_MB, MP_NB, MXLDA, MXCOL)
CALL SCALAPACK_GETDIM(K, N, MP_NB, MP_MB, MXLDB, MXCOLB)
CALL SCALAPACK_GETDIM(M, N, MP_MB, MP_NB, MXLDC, MXCOLC)
! Set up the array descriptors
CALL DESCINIT(DESC_A, M, K, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
CALL DESCINIT(DESC_B, K, N, MP_NB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDB, INFO)
CALL DESCINIT(DESC_C, M, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDC, INFO)

```
```

ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDB,MXCOLB) , CO (MXLDC,MXCOLC))
! A root process is used to create the matrix data for the test.
IF(MP_RANK == 0) THEN
ALLOCATE (A (M, K), B (K,N), C (M,N), X(M))
CALL RANDOM_NUMBER(A) ; CALL RANDOM_NUMBER (B)
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='UNKNOWN')
! Write the data by columns.
DO J=1,K,MP_NB
WRITE (NIN,*) ((A (I, L), I=1,M),L=J,min(K,J+MP_NB-1))
END DO
CLOSE (NIN)
OPEN(UNIT=NIN, FILE='Btest.dat', STATUS='UNKNOWN')
! Write the data by columns.
DO J=1,N,MP NB
WRITE (NIN,*) ((B (I,L),I=1,K), L=J,min(N,J+MP_NB-1))
END DO
CLOSE (NIN)
END IF
! Read the factors into the local arrays.
CALL ScaLAPACK_READ('Atest.dat', DESC_A, A0)
CALL ScaLAPACK_READ('Btest.dat', DESC_B, B0)
! Compute the distributed product C = A x B.
ALPHA=1d0; BETA=0d0
IA=1; JA=1; IB=1; JB=1; IC=1; JC=1
CO=0
CALL pdGEMM \&
("NO", "NO", M, N, K, ALPHA, AO, IA, JA,\&
DESC_A, BO, IB, JB, DESC_B, BETA, \&
C0, IC, JC, DESC_C )
! Put the product back on the root node.
Call ScaLAPACK_WRITE('Ctest.dat', DESC_C, CO)
IF(MP_RANK == 0) THEN
! Read the residuals and check them for size.
OPEN(UNIT=NIN, FILE='Ctest.dat', STATUS='OLD')
! Read the data by columns.
DO J=1,N,MP_NB
READ (NIN, *) ((C (I,L),I=1,M),L=J,min (N,J+MP_NB-1))
END DO
CLOSE (NIN, STATUS='DELETE')
SIZE_C=SUM(ABS (C)) ; C=C-matmul (A,B)
ERRO\overline{R}=SUM (ABS (C))/SIZE_C
! Open other temporary files and delete them.
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='OLD')

```
```

    CLOSE (NIN,STATUS='DELETE')
    OPEN(UNIT=NIN, FILE='Btest.dat', STATUS='OLD')
    CLOSE (NIN,STATUS='DELETE')
    END IF
! See to any error messages.
call e1pop("Mp_Setup")
! Deallocate storage arrays and exit from BLACS.
IF(ALLOCATED(A)) DEALLOCATE (A)
IF (ALLOCATED(B)) DEALLOCATE (B)
IF(ALLOCATED(C)) DEALLOCATE (C)
IF(ALLOCATED(X)) DEALLOCATE (X)
IF(ALLOCATED(A0)) DEALLOCATE(A0)
IF(ALLOCATED(B0)) DEALLOCATE(B0)
IF(ALLOCATED(C0)) DEALLOCATE(C0)
! Check the results.
IF(ERROR <= SQRT(EPSILON(ALPHA)) . and. \&
MP RANK == 0) THEN
wrīte(*,*) " Example 2 for BLACS and PBLAS is correct."
END IF
! Exit from using this process grid.
CALL SCALAPACK_EXIT( MP_ICTXT )
! Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}
```

Example 2 for BLACS and PBLAS is correct.

```

\section*{Example 3: Distributed Linear Solver with ScaLAPACK}

The program SCPK_EX3 illustrates solving a system of linear-algebraic equations, \(A x=b\) by calling a ScaLAPACK routine directly. The right-hand side is produced by defining \(A\) and \(y\) to have random values. Then the matrix-vector product \(b=A y\) is computed. The problem size is such that the residuals, \(x-y \approx 0\) are checked on one process. Three temporary files are created and deleted. BLACS are used to define the process grid and provide further information identifying each process. Then a ScaLAPACK routine is called directly to compute the approximate solution, \(x\).
```

    program scpk_ex3
    ! This is Example 3 for ScaLAPACK READ and ScaLAPACK WRITE.
! A linear system is solved with ScaLAPACK and checked.
USE ScaLAPACK SUPPORT
USE ERROR OPTİON PACKET
USE MPI_SETUP_INT
IMPLICIT NONE
INCLUDE "mpif.h"

```
```

INTEGER, PARAMETER :: N=9, NIN=10
INTEGER INFO, IA, JA, IB, JB, MXLDA,MXCOL, \&
IERROR, I, J, L, DESC_A(9),\&
DESC_B(9), BUFF(3), RBUF(3)
LOGICAL :: COMMUTE = .TRUE., NSQUARE = .TRUE., GRID1D = .TRUE.
INTEGER, ALLOCATABLE :: IPIVO(:)
real(kind(1d0)) : : ERROR=0d0, SIZE_Y
real(kind(1d0)), allocatable, dimension(:,:) :: A, B(:), \&
X(:), Y(:), A0, B0
MP_NPROCS=MP_SETUP()
! Set up a 1D processor grid and define its context ID, MP_ICTXT
CALL SCALAPACK SETUP(N, N, NSQUARE, GRID1D)
! Get the array descriptor entities
CALL SCALAPACK_GETDIM(N, N, MP MB, MP NB, MXLDA, MXCOL)
! Set up the arrāy descriptors
CALL DESCINIT(DESC_A, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
CALL DESCINIT(DESC_B, N, 1, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
! Allocate local space for each array.
ALLOCATE (A0 (MXLDA,MXCOL), B0 (MXLDA, 1), IPIVO (MXLDA+MP MB))
! A root process is used to create the matrix data for the test.
IF(MP RANK == 0) THEN
ALL\overline{O}CATE (A (N,N), B(N), X(N), Y(N))
CALL RANDOM_NUMBER(A) ; CALL RANDOM_NUMBER(Y)
! Compute the correct result.
B=MATMUL (A,Y); SIZE_Y=SUM(ABS (Y))
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='UNKNOWN')
! Write the data by columns.
DO J=1,N,MP NB
WRITE (NIN,*) ((A (I, L), I=1,N), L=J,min(N,J+MP_NB-1))
END DO
CLOSE (NIN)
OPEN(UNIT=NIN, FILE='Btest.dat', STATUS='UNKNOWN')
! Write the data by columns.
WRITE (NIN, *) (B (I) , I=1,N)
CLOSE (NIN)
END IF
! Read the factors into the local arrays.
CALL ScaLAPACK_READ('Atest.dat', DESC_A, A0)
CALL ScaLAPACK_READ('Btest.dat', DESC_B, B0)
! Compute the distributed product solution to A x = b.
IA=1; JA=1; IB=1; JB=1
CALL pdGESV (N, 1, A0, IA, JA, DESC_A, IPIVO, \&

```
```

B0, IB, JB, DESC B, INFO)
! Put the result on the root node.
Call ScaLAPACK_WRITE('Xtest.dat', DESC_B, B0)
IF(MP_RANK == 0) THEN
! Read the residuals and check them for size.
OPEN(UNIT=NIN, FILE='Xtest.dat', STATUS='OLD')
! Read the approximate solution data.
READ (NIN,*) X
B=X-Y
CLOSE (NIN,STATUS='DELETE')
ERROR=SUM(ABS (B))/SIZE_Y
! Delete temporary files.
OPEN(UNIT=NIN, FILE='Atest.dat', STATUS='OLD')
CLOSE (NIN,STATUS='DELETE')
OPEN(UNIT=NIN, FILE='Btest.dat', STATUS='OLD')
CLOSE(NIN,STATUS='DELETE')
END IF
! See to any error messages.
call e1pop("Mp_Setup")
! Deallocate storage arrays
IF(ALLOCATED(A)) DEALLOCATE (A)
IF(ALLOCATED(B)) DEALLOCATE (B)
IF(ALLOCATED(X)) DEALLOCATE (X)
IF(ALLOCATED(Y)) DEALLOCATE (Y)
IF(ALLOCATED(A0)) DEALLOCATE (A0)
IF(ALLOCATED(B0)) DEALLOCATE(B0)
IF(ALLOCATED(IPIVO)) DEALLOCATE (IPIVO)
IF(ERROR <= SQRT(EPSILON(ERROR)) . and. MP_RANK == 0) THEN
write(*,*) \&
" Example 3 for BLACS and ScaLAPACK solver is correct."
END IF
! Exit from using this process grid.
CALL SCALAPACK_EXIT( MP_ICTXT )
! Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
END

```

\section*{Output}

Example 3 for BLACS and ScaLAPACK is correct.

\section*{ScaLAPACK_MAP}

\section*{匀MPI}

REQUIRED

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine maps array data from a global array to local arrays in the two-dimensional blockcyclic form required by ScaLAPACK routines.

All processors in the \(B L A C S\) context call the routine.

\section*{Required Arguments}

A - Global rank-1 or rank-2 array which is to be mapped to the processor grid. The data type for A is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; double precision, complex. Normally, the user defines A to be valid only on the MP_RANK \(=0\) processor. (Input)

DESC_A - An integer vector containing the nine parameters associated with the ScaLAPACK matrix descriptor for array A. See "Usage Notes for ScaLAPACK Utilities" for a description of the nine parameters. (Input)

A0 - This is a local rank-1 or rank-2 array that will contain this processor's piece of the block-cyclic array. The data type for A0 is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; and double precision, complex. (Output)

\section*{Optional Arguments}
\(\boldsymbol{L D A}\) - Leading dimension of A as specified in the calling program. If this argument is not present, \(\operatorname{SIZE}(A, 1)\) is used. (Input)

COLMAP - Input logical which indicates whether the global array should be mapped in column major form or row major form. COLMAP set to .TRUE. will result in the array being mapped in column- major form while setting COLMAP to .FALSE. will result in the array being mapped in row major form. The default value of COLMAP is .TRUE. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_MAP (A, DESC_A, A0 [, ...])

\section*{Description}

Subroutine ScaLAPACK_MAP maps columns or rows of a global array on MP_RANK \(=0\) to local distributed arrays so that the problem array is usable by a ScaLAPACK routine. It uses the two-dimensional block-cyclic array descriptor for the matrix to place the data in the desired assumed-size arrays on the processors. The block sizes, contained in the array descriptor, determine the data set size for each blocking send and receive pair. The number of these synchronization points is proportional to \(\lceil M \times N /(M B \times N B)\rceil\). A temporary local buffer is allocated for staging the array data. It is of size m by NB , when mapping by columns, or N by MB , when mapping by rows.

\section*{Example}

See ScaLAPACK_UNMAP.

\section*{ScaLAPACK_UNMAP}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine unmaps array data from local distributed arrays to a global array. The data in the local arrays must have been stored in the two-dimensional block-cyclic form required by ScaLAPACK routines. All processors in the BLACS context call the routine.

\section*{Required Arguments}
\(\boldsymbol{A 0}\) - This is a local rank-1 or rank-2 array that contains this processor's piece of the blockcyclic array. The data type for A0 is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; or double precision, complex. (Input)

DESC_A - An integer vector containing the nine parameters associated with the ScaLAPACK matrix descriptor for array A. See "Usage Notes for ScaLAPACK Utilities" for a description of the nine parameters. (Input)

A - Global rank-1 or rank-2 array which is to receive the array which had been mapped to the processor grid. The data type for A is any of five Fortran intrinsic types: integer; single precision, real; double precision, real; single precision, complex; or double precision, complex. A is only valid on MP_RANK \(=0\) after ScaLAPACK_UNMAP has been called. (Output)

\section*{Optional Arguments}
\(\boldsymbol{L D A}\) - Leading dimension of A as specified in the calling program. If this argument is not present, \(\operatorname{SIZE}(A, 1)\) is used. (Input)

COLMAP - Input logical which indicates whether the global array should be mapped in column major form or row major form. COLMAP set to .TRUE. will result in the array being mapped in column major form while setting COLMAP to .FALSE. will result in the array being mapped in row major form. The default value of COLMAP is .TRUE. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_UNMAP (A0, DESC_A, A [,...])

\section*{Description}

Subroutine ScaLAPACK_UNMAP unmaps columns or rows of local distributed arrays to a global array on MP_RANK \(=0\). It uses the two-dimensional block-cyclic array descriptor for the matrix to retrieve the data from the assumed-size arrays on the processors. The block sizes, contained in the array descriptor, determine the data set size for each blocking send and receive pair. The number of these synchronization points is proportional to \(\lceil M \times N /(M B \times N B)\rceil\). A temporary local buffer is allocated for staging the array data. It is of size M by NB, when mapping by columns, or N by MB, when mapping by rows.

\section*{Example: Distributed Linear Solver with IMSL ScaLAPACK Interface}

The program SCPKMP_EX1 illustrates solving a system of linear-algebraic equations, \(A x=b\) by calling routine LSLRG, an IMSL routine which interfaces with a ScaLAPACK routine. The righthand side is produced by defining \(A\) and \(y\) to have random values. Then the matrix-vector product \(b=A y\) is computed. The problem size is such that the residuals, \(x-y \approx 0\) are checked on MP_RANK \(=0\). IMSL routine ScaLAPACK_SETUP is called to define the process grid and provide further information identifying each process. IMSL routine ScaLAPACK_MAP is called to map the global arrays to local distributed arrays. Then LSLRG is called to compute the approximate solution, \(x\).
```

program scpkmp_ex1
! This is Example 1 for ScaLAPACK_MAP and ScaLAPACK_UNMAP.
! A linear system is solved with an IMSL routine which
! interfaces with ScaLAPACK and is checked.
USE ScaLAPACK_SUPPORT
USE ERROR OPTION PACKET
USE MPI SETUP INT
USE LSLRG_INT
IMPLICIT NONE
INCLUDE "mpif.h"
INTEGER, PARAMETER :: N=9

```
```

INTEGER MXLDA, MXCOL, INFO, DESC_A(9), DESC_X(9)
LOGICAL :: GRID1D = .TRUE., NSQUARE = .TRUE.
real(kind(1d0)) :: ERROR=0d0, SIZE_Y
real(kind(1d0)), allocatable, dimension(:,:) :: A, B(:), \&
X(:), Y(:), A0, BO(:), XO(:)
MP_NPROCS=MP_SETUP()
! Set up a 1D processor grid and define its context ID, MP_ICTXT
CALL SCALAPACK_SETUP(N, N, NSQUARE, GRID1D)
! Get the array \overline{descriptor entities MXLDA, and MXCOL}
CALL SCALAPACK_GETDIM(N, N, MP_MB, MP_NB, MXLDA, MXCOL)
! Set up the array descriptors
CALL DESCINIT(DESC_A, N, N, MP_MB, MP_NB, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
CALL DESCINIT(DESC_X, N, 1, MP_MB, 1, 0, 0, MP_ICTXT, \&
MXLDA, INFO)
! Allocate space for local arrays
ALLOCATE (A0 (MXLDA,MXCOL), BO (MXLDA), X0 (MXLDA))
! A root process is used to create the matrix data for the test.
IF(MP RANK == O) THEN
ALLOCATE (A (N,N), B(N), X(N), Y(N))
CALL RANDOM_NUMBER(A); CALL RANDOM_NUMBER(Y)
! Compute the correct result.
B=MATMUL (A,Y) ; SIZE_Y=SUM(ABS (Y))
END IF
! Map the input arrays to the processor grid
CALL SCALAPACK_MAP(A, DESC_A, A0)
CALL SCALAPACK_MAP(B, DESC_X, B0)
! Compute the distributed product solution to A x = b.
CALL LSLRG(A0, BO, XO)
! Put the result on the root node.
Call ScaLAPACK UNMAP(X0, DESC X, X)
IF(MP RANK == O) THEN
! Check the residuals for size.
B=X-Y
ERROR=SUM(ABS (B))/SIZE_Y
END IF
! See to any error messages.
call e1pop("Mp_Setup")
IF(ERROR <= SQRT(EPSILON(ERROR)) . and. MP_RANK == 0) THEN
write(*,*) \&
" Example 1 for ScaLAPACK MAP and ScaLAPACK UNMAP is correct."
END IF
! Deallocate storage arrays.
IF (MP_RANK == 0) DEALLOCATE (A, B, X, Y)
DEALLOCATE(A0, B0, XO)

```
```

! Exit from using this process grid.
CALL SCALAPACK_EXIT( MP_ICTXT )
! Shut down MPI
MP_NPROCS = MP_SETUP('FINAL')
EN\overline{D}

```

\section*{Output}

Example 1 for ScaLAPACK_MAP and ScaLAPACK_UNMAP is correct.

\section*{ScaLAPACK_EXIT}

For a detailed description of MPI Requirements see "Using ScaLAPACK Enhanced Routines" in the Introduction of this manual.

This routine exits ScaLAPACK mode for the IMSL Library routines. All processors in the BLACS context call the routine.

\section*{Required Arguments}

ICTXT - The BLACS context ID to which the processor grid is associated. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL ScaLAPACK_EXIT (ICTXT)

\section*{Description}

Subroutine ScaLAPACK_EXIT exits ScaLAPACK mode for the IMSL Library routines. The following actions occur when this routine is called:
- BLACS_GRIDEXIT is called with the input BLACS context ID.
- The pointer to the grid ID, MP_PIGRID is nullified.
- If the grid, MP_IGRID, has been allocated, it is deallocated.
- MP_ICTXT is reset to its default value, \(\operatorname{HUGE}(1)\).

\section*{ERROR_POST}

Prints error messages that are generated by IMSL routines using EPACK.

\section*{Required Argument}

EPACK - (Input [/Output]) Derived type array of size \(p\) containing the array of message numbers and associated data for the messages. The definition of this derived type is packaged within the modules used as interfaces for each suite of routines. The declaration is:
```

type ?_error
integer idummy; real(kind(? )) rdummy
end type

```

The choice of "? " is either "s_" or "d_" depending on the accuracy of the data. This array gets additional messages and data from each routine that uses the "epack=" optional argument, provided \(p\) is large enough to hold data for a new message. The value \(p=8\) is sufficient to hold the longest single terminal, fatal, or warning message that an IMSL Fortran Library routine generates.

The location at entry epack (1) \%idummy contains the number of data items for all messages. When the error_post routine exits, this value is set to zero. Locations in array positions (2:) \%idummy contain groups of integers consisting of a message number, the error severity level, then the required integer data for the message. Floating-point data, if required in the message, is passed in locations(: ) \(\%\) rdummy matched with the starting point for integer data. The extent of the data for each message is determined by the requirements of the larger of each group of integer or floating-point values.

\section*{Optional Arguments}
new_unit = nunit (Input)
Unit number, of type integer, associated for reading the direct-access file of error messages for the IMSL Fortran 90 routines.
Default: nunit \(=4\)
new_path = path (Input)
Pathname in the local file space, of type character*64, needed for reading the directaccess file of error messages. Default string for path is defined during the installation procedure for certain IMSL Fortran Library routines.

\section*{FORTRAN 90 Interface}

Generic: CALL ERROR_POST (EPACK [,...])
Specific: The specific interface names are S_ERROR_POST and D_ERROR_POST.

\section*{Description}

A default direct-access error message file (.daf file) is supplied with this product. This file is read by error_post using the contents of the derived type argument epack, containing the message number, error severity level, and associated data. The message is converted into character strings accepted by the error processor and then printed. The number of pending messages that print
depends on the settings of the parameters PRINT and STOP in the Reference Material in the IMSL MATH/LIBRARY User's Manual. These values are initialized to defaults such that any Level 5 or Level 4 message causes a STOP within the error processor after a print of the text. To change these defaults so that more than one error message prints, use the routine ERSET documented and illustrated with examples in the Reference Material in the IMSL MATH/LIBRARY User's Manual. The method of using a message file to store the messages is required to support "sharedmemory parallelism."

\section*{Managing the Message File}

For most applications of this product, there will be no need to manage this file. However, there are a few situations which may require changing or adding messages:
- New system-wide messages have been developed for applications using this Library.
- All or some of the existing messages need to be translated to another language
- A subset of users need to add a specific message file for their applications using this Library.

Following is information on changing the contents of the message file, and information on how to create and access a message file for a private application.

\section*{Changing Messages}

In order to change messages, two files are required:
- An editable message glossary, messages.gls, supplied with this product.
- A source program, prepmess.f, used to generate an executable which builds messages.daf from messages.gls.

To change messages, first make a backup copy of messages.gls. Use a text editor to edit messages.gls. The format of this file is a series of pairs of statements:
- message_number=<nnnn>
- message='message string'
(Note that neither of these lines should begin with a tab.)
The variable <nnnn> is an integer message number (see below for ranges and reserved message numbers).

The 'message string' is any valid message string not to exceed 255 characters. If a message line is too long for a screen, the standard Fortran 90 concatenation operator // with the line continuation character \& may be used to wrap the text.
Most strings have substitution parameters embedded within them. These may be in the following forms:
- \(\%(\mathrm{i}<\mathrm{n}>)\) for an integer substitution, where n is the n th integer output in this message.
- \(\%(\mathrm{r}<\mathrm{n}>)\) for single precision real number substitution, where n is the n th real number output in this message.
- \(\quad \%(\mathrm{~d}<\mathrm{n}>)\) for double precision real number substitution, where n is the n th double precision number output in this message.
New messages added to the system-wide error message file should be placed at the end of the file. Message numbers 5000 through 10000 have been reserved for user-added messages. Currently, messages 1 through 1400 are used by IMSL. Gaps in message number ranges are permitted; however, the message numbers must be in ascending order within the file. The message numbers used for each IMSL Fortran Library subroutine are documented in this manual and in online help.
If existing messages are being edited or translated, make sure not to alter the message_number lines. (This prevents conflicts with any new messages.gls file supplied with future versions of this Library.)

\section*{Building a New Direct-access Message File}

The prepmess executable must be available to complete the message changing process. For information on building the prepmess executable from prepmess.f, consult the installation guide for this product.

Once new messages have been placed in the messages.gls file, make a backup copy of the messages.daf file. Then remove messages.daf from the current directory. Now enter the following command:
```

prepmess > prepmess_output

```

A new messages.daf file is created. Edit the prepmess_output file and look near the end of the file for the new error messages. The prepmess program processes each message through the error message system as a validity check. There should be no FATAL error announcement within the prepmess_output file.

\section*{Private Message Files}

Users can create a private message file within their own messages. This file would generally be used by an application that calls this Library. Follow the steps outlined above to create a private messages.gls file. The user should then be given a copy of the prepmess executable. In the application code, call the error_post subprogram with the new_unit/new_path optional arguments. The new path should point to the directory in which the private messages. daf file resides.

\section*{SHOW}

Prints rank-1 or rank-2 arrays of numbers in a readable format.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Rank-1 or rank-2 array containing the numbers to be printed. (Input)

\section*{Optional Arguments}
```

text = CHARACTER (Input)
CHARACTER (LEN=*) string used for labeling the array.

```
image = buffer (Output)
CHARACTER (LEN=*) string used for an internal write buffer. With this argument present the output is converted to characters and packed. The lines are separated by an end-of-line sequence. The length of buffer is estimated by the line width in effect, time the number of lines for the array.
iopt \(=\) iopt (:) (Input)
Derived type array with the same precision as the input array; used for passing optional data to the routine. Use the REAL (KIND (1E0) ) precision for output of INTEGER arrays. The options are as follows:
\begin{tabular}{|l|l|l|}
\hline \multicolumn{3}{|c|}{ Packaged Options for SHOW } \\
\hline Prefix is blank & Option Name & Option Value \\
\hline & show_significant_digits_is_4 & 1 \\
\hline & show_significant_digits_is_7 & 2 \\
\hline & show_significant_digits_is_16 & 3 \\
\hline & show_line_width_is_44 & 4 \\
\hline & show_line_width_is_72 & 5 \\
\hline & show_line_width_is_128 & 6 \\
\hline & show_end_of_line_sequence_is & 7 \\
\hline & show_starting_index_is & 8 \\
\hline & show_starting_row_index_is & 9 \\
\hline & show_starting_col_index_is & 10 \\
\hline
\end{tabular}
```

iopt(IO) = show_significant_digits_is_4
iopt(IO) = show_significant_digits_is_7
iopt(IO) = show_significant_digits_is_16

```

These options allow more precision to be displayed. The default is 4D for each value. The other possible choices display 7D or 16D.
```

iopt(IO) = show_line_width_is_44
iopt(IO) = show_line_width_is_72
iopt(IO) = show_line_width_is_128

```

These options allow varying the output line width. The default is 72 characters per line. This allows output on many work stations or terminals to be read without wrapping of lines.
```

iopt(IO) = show_end-of_line_sequence_is

```

The sequence of characters ending a line when it is placed into the internal character buffer corresponding to the optional argument 'IMAGE \(=\) buffer'.

The value of iopt (IO 1 ) \%idummy is the number of characters. These are followed, starting at iopt (IO+2) \%idummy, by the ASCII codes of the characters themselves. The default is the single character, ASCII value 10 or New Line.
```

iopt(IO) = show_starting_index_is

```

This are used to reset the starting index for a rank-1 array to a value different from the default value, which is 1 .
```

iopt(IO) = show_starting_row_index_is
iopt(IO) = show_starting_col_index_is

```

These are used to reset the starting row and column indices to values different from their defaults, each 1.

\section*{FORTRAN 90 Interface}

Generic: CALL SHOW (X [, ...])

Specific: The specific interface names are S_SHOW and D_SHOW.

\section*{Description}

The show routine is a generic subroutine interface to separate low-level subroutines for each data type and array shape. Output is directed to the unit number IUNIT. That number is obtained with the subroutine UMACH, IMSL MATH/LIBRARY User's Manual. Thus the user must open this unit in the calling program if it desired to be different from the standard output unit. If the optional argument 'IMAGE = buffer' is present, the output is not sent to a file but to a character string within buffer. These characters are available to output or be used in the application.

\section*{Fatal and Terminal Error Messages}

See the messages.gls file for error messages for SHOW. These error messages are numbered 601-606; 611-617; 621-627; 631-636; 641-646.

\section*{Example 1: Printing an Array}

Array of random numbers for all the intrinsic data types are printed. For REAL (KIND (1E0)) rank-1 arrays, the number of displayed digits is reset from the default value of 4 to the value 7 and the subscripts for the array are reset so they match their declared extent when printed. The output is not shown.
```

    use show_int
    use rand_int
    implicit none
    ! This is Example 1 for SHOW.
integer, parameter :: n=7, m=3

```
```

        real(kind(1e0)) s x(-1:n), s m(m,n)
        real(kind(1d0)) d_x(n), d_m(m,n)
        complex(kind(1e0)) c_x(n), c_m(m,n)
        complex(kind(1d0)) z_x(n),z_m(m,n)
        integer i_x(n), i_m(m,n)
            type (s_options) options(3)
    ! The data types printed are real(kind(le0)), real(kind(1d0)),
! complex(kind(1e0)), complex(kind(1d0)), and INTEGER.
! Fill with randsom numbers and then print the contents,
! in each case with a label.
s x=rand(s_x); s m=rand (s m)
d_x=rand (d_x); d_m=rand (d_m)
c_x=rand (c_x); c_m=rand (c_m)
z_x=rand(z_x); z_m=rand(z_m)
i_x=100*rand(s_x(1:n)); i_m=100*rand(s_m)
call show (s x, 'Rank-1, REAL')
call show (s_m, 'Rank-2, REAL')
call show (d x, 'Rank-1, DOUBLE')
call show (d_m, 'Rank-2, DOUBLE')
call show (c_x, 'Rank-1, COMPLEX')
call show (c m, 'Rank-2, COMPLEX')
call show (z_x, 'Rank-1, DOUBLE COMPLEX')
call show (z m, 'Rank-2, DOUBLE COMPLEX')
call show (i-x, 'Rank-1, INTEGER')
call show (i_m, 'Rank-2, INTEGER')
! Show 7 digits per number and according to the
! natural or declared size of the array.
options(1)=show_significant_digits_is_7
options(2)=show_starting_index_is
options(3)= -1 ! The starting value.
call show (s x, \&
'Rank-1, REAL with 7 digits, natural indexing', IOPT=options)
end

```

\section*{Output}

Example 1 for SHOW is correct.

\section*{Additional Examples}

\section*{Example 2: Writing an Array to a Character Variable}

This example prepares a rank-1 array for further processing, in this case delayed writing to the standard output unit. The indices and the amount of precision are reset from their defaults, as in Example 1. An end-of-line sequence of the characters CR-NL \((A S C I I 10,13)\) is used in place of the standard ASCII 10. This is not required for writing this array, but is included for an illustration of the option.
```

use show int
use rand_int

```
```

    implicit none
    ! This is Example 2 for SHOW.
integer, parameter :: n=7
real(kind(1e0)) s_x(-1:n)
type (s_options) options(7)
CHARACTER (LEN= (72+2)*4) BUFFER
! The data types printed are real(kind(1e0)) random numbers.
s_x=rand(s_x)
! Show 7 digits per number and according to the
! natural or declared size of the array.
! Prepare the output lines in array BUFFER.
! End each line with ASCII sequence CR-NL.
options(1)=show_significant_digits_is_7
options(2)=show starting_index_is
options(3)= -1 ! The starting value.
options(4)=show_end_of_line_sequence_is
options(5)= 2 ! Use 2 EOL characters.
options(6)= 10 ! The ASCII code for CR.
options(7)= 13 ! The ASCII code for NL.
BUFFER= ' ' ! Blank out the buffer.
! Prepare the output in BUFFER.
call show (s_x, \&
'Rank-1, REAL with 7 digits, natural indexing '//\&
'internal BUFFER, CR-NL EOLs.',\&
IMAGE=BUFFER, IOPT=options)
! Display BUFFER as a CHARACTER array. Discard blanks
! on the ends.
WRITE(*,'(1x,A)') TRIM(BUFFER)
end

```

\section*{Output}

Example 2 for SHOW is correct.

\section*{WRRRN}

Prints a real rectangular matrix with integer row and column labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input)
TITLE set equal to a blank character(s) suppresses printing of the title. Use " \(\% /\) /" within the title to create a new line. Long titles are automatically wrapped.
\(\boldsymbol{A}\) - NRA by NCA matrix to be printed. (Input)

\section*{Optional Arguments}

NRA - Number of rows. (Input)
Default: NRA \(=\operatorname{SIZE}(\mathrm{A}, 1)\).
NCA - Number of columns. (Input)
Default: NCA \(=\operatorname{SIZE}(\mathrm{A}, 2)\).
\(\boldsymbol{L D A}\) - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA \(=\operatorname{SIZE}(A, 1)\).
ITRING — Triangle option. (Input)
Default: ITRING \(=0\).

\section*{ITRING Action}
\(0 \quad\) Full matrix is printed.
1 Upper triangle of \(A\) is printed, including the diagonal.
2 Upper triangle of A excluding the diagonal of A is printed.
-1 Lower triangle of \(A\) is printed, including the diagonal.
-2 Lower triangle of A excluding the diagonal of \(A\) is printed.

\section*{FORTRAN 90 Interface}

Generic: CALL WRRRN (TITLE, A \([, \ldots]\) )
Specific: The specific interface names are S_WRRRN and D_WRRRN for two dimensional arrays, and S_WRRRN1D and D_WRRRN1D for one dimensional arrays.

\section*{FORTRAN 77 Interface}

Single: CALL WRRRN (TITLE, NRA, NCA, A, LDA, ITRING)
Double: The double precision name is DWRRRN.

\section*{Description}

Routine WRRRN prints a real rectangular matrix with the rows and columns labeled \(1,2,3\), and so on. WRRRN can restrict printing to the elements of the upper or lower triangles of matrices via the ITRING option. Generally, ITRING \(\neq 0\) is used with symmetric matrices.

In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRA to the length of the array and set NCA \(=1\). For a row vector, set NRA \(=1\) and set NCA to the length of the array. In both cases, set LDA \(=\) NRA and set \(I T R I N G=0\).

\section*{Comments}
1. A single \(D, E\), or \(F\) format is chosen automatically in order to print 4 significant digits for the largest element of A in absolute value. Routine WROPT can be used to change the default format.
2. Horizontal centering, a method for printing large matrices, paging, printing a title on each page, and many other options can be selected by invoking WROPT.
3. A page width of 78 characters is used. Page width and page length can be reset by invoking PGOPT.
4. Output is written to the unit specified by UMACH (see the Reference Material).

\section*{Example}

The following example prints all of a \(3 \times 4\) matrix \(A\) where \(a_{i j}=i+j / 10\). USE WRRRN_INT

IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=10, NCA=4, NRA=3)
\(!\)
INTEGER I, J
REAL A (LDA, NCA)
!
DO 20 I=1, NRA DO \(10 \mathrm{~J}=1\), NCA \(A(I, J)=I+J * 0.1\) CONTINUE
CONTINUE
! Write A matrix.
CALL WRRRN ('A', A, NRA=NRA)
END

\section*{Output}
\begin{tabular}{rrrrr} 
& \multicolumn{4}{c}{ A } \\
& 1 & 2 & 3 & 4 \\
1 & 1.100 & 1.200 & 1.300 & 1.400 \\
2 & 2.100 & 2.200 & 2.300 & 2.400 \\
3 & 3.100 & 3.200 & 3.300 & 3.400
\end{tabular}

\section*{WRRRL}

Print a real rectangular matrix with a given format and labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input) TITLE set equal to a blank character(s) suppresses printing of the title.
\(\boldsymbol{A}\) - NRA by NCA matrix to be printed. (Input)
RLABEL — CHARACTER * (*) vector of labels for rows of A. (Input)
If rows are to be numbered consecutively \(1,2, \ldots\), NRA, use \(\operatorname{RLABEL}(1)={ }^{\prime}\) NUMBER' \(^{\prime}\). If no row labels are desired, use RLABEL \((1)={ }^{\prime}\) NONE \(^{\prime}\). Otherwise, RLABEL is a vector of length NRA containing the labels.

CLABEL — CHARACTER * (*) vector of labels for columns of A. (Input)
If columns are to be numbered consecutively \(1,2, \ldots\), NCA, use
\(\operatorname{CLABEL}(1)=\) ' \(\operatorname{NUMBER}\) '. If no column labels are desired, use CLABEL(1)= \({ }^{\prime} \mathrm{NONE}^{\prime}\). Otherwise, CLABEL(1) is the heading for the row labels, and either CLABEL(2) must be ' NUMBER' or 'NONE', or CLABEL must be a vector of length NCA +1 with \(\operatorname{CLABEL}(1+j)\) containing the column heading for the \(j\)-th column.

\section*{Optional Arguments}

NRA - Number of rows. (Input)
Default: NRA \(=\operatorname{SIZE}(A, 1)\).

NCA - Number of columns. (Input)
Default: NCA \(=\operatorname{SIZE}(A, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA \(=\operatorname{SIZE}(A, 1)\).
ITRING - Triangle option. (Input)
Default: ITRING \(=0\).

\section*{ITRING Action}
\(0 \quad\) Full matrix is printed.
1 Upper triangle of \(A\) is printed, including the diagonal.
2 Upper triangle of A excluding the diagonal of \(A\) is printed.
-1 Lower triangle of A is printed, including the diagonal.
-2 Lower triangle of A excluding the diagonal of A is printed.

FMT - Character string containing formats. (Input)
If FMT is set to a blank character(s), the format used is specified by WROPT. Otherwise, FMT must contain exactly one set of parentheses and one or more edit descriptors. For example, \(\mathrm{FMT}={ }^{\prime}(\mathrm{F} 10.3)^{\prime}\) specifies this F format for the entire matrix.
FMT \(={ }^{\prime}(2 E 10.3,3 F 10.3)^{\prime}\) specifies an \(E\) format for columns 1 and 2 and an \(F\) format for columns 3, 4 and 5. If the end of FMT is encountered and if some columns of the matrix remain, format control continues with the first format in FMT. Even though the matrix A is real, an I format can be used to print the integer part of matrix elements of \(A\). The most useful formats are special formats, called the \(v\) and \(w\) formats, that can be used to specify pretty formats automatically. Set FMT = '(V10.4)' if you want a single \(D, E\), or \(F\) format selected automatically with field width 10 and with 4 significant digits. Set \(\mathrm{FMT}={ }^{\prime}(\mathrm{W} 10.4)^{\prime}\) if you want a single D, E, F, or I format selected automatically with field width 10 and with 4 significant digits. While the v format prints trailing zeroes and a trailing decimal point, the w format does not. See Comment 4 for general descriptions of the v and w formats. FMT may contain only \(\mathrm{D}, \mathrm{E}\), F, G, I, V, or w edit descriptors, e.g., the x descriptor is not allowed.
Default: \(\operatorname{FMT}=\) ' .

\section*{FORTRAN 90 Interface}

Generic: CALL WRRRL (TITLE, A, RLABEL, CLABEL [ , ...])
Specific: The specific interface names are S_WRRRL and D_WRRRL for two dimensional arrays, and S_WRRRL1D and D_WRRRL1D for one dimensional arrays.

\section*{FORTRAN 77 Interface}

Single: CALL WRRRL (TITLE, NRA, NCA, A, LDA, ITRING, FMT, RLABEL, CLABEL)

Double: The double precision name is DWRRRL.

\section*{Description}

Routine WRRRL prints a real rectangular matrix (stored in \(A\) ) with row and column labels (specified by RLABEL and CLABEL, respectively) according to a given format (stored in FMT). WRRRL can restrict printing to the elements of upper or lower triangles of matrices via the ITRING option. Generally, ITRING \(\neq 0\) is used with symmetric matrices.

In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRA to the length of the array and set NCA \(=1\). For a row vector, set NRA \(=1\) and set NCA to the length of the array. In both cases, set LDA \(=\) NRA, and set \(\operatorname{ITRING}=0\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of W2RRL / DW2RRL. The reference is:
```

CALL W2RRL (TITLE, NRA, NCA, A, LDA, ITRING, FMT, RLABEL, CLABEL, CHWK)

```

The additional argument is:
\(\boldsymbol{C H W K}\) - CHARACTER * 10 work vector of length NCA. This workspace is referenced only if all three conditions indicated at the beginning of this comment are met. Otherwise, CHWK is not referenced and can be a CHARACTER * 10 vector of length one.
2. The output appears in the following form:

\section*{TITLE}
\begin{tabular}{llll} 
CLABEL (1) & CLABEL (2) & CLABEL (3) & CLABEL (4) \\
RLABEL(1) & Xxxxx & Xxxxx & Xxxxx \\
RLABEL(2) & Xxxxx & Xxxxx & Xxxxx
\end{tabular}
3. Use "\% /" within titles or labels to create a new line. Long titles or labels are automatically wrapped.
4. For printing numbers whose magnitudes are unknown, the G format in FORTRAN is useful; however, the decimal points will generally not be aligned when printing a column of numbers. The V and w formats are special formats used by this routine to select a \(D, E, F\), or \(I\) format so that the decimal points will be aligned. The V and w formats are specified as Vn.d and Wn.d. Here, \(n\) is the field width and \(d\) is the number of significant digits generally printed. Valid values for \(n\) are \(3,4, \ldots, 40\). Valid values for \(d\) are \(1,2, \ldots, n-2\). If FMT specifies one format and that format is a v or w format, all elements of the matrix A are examined to determine one FORTRAN format for printing. If FMT specifies more than one format, FORTRAN formats are generated separately from each \(v\) or \(w\) format.
5. A page width of 78 characters is used. Page width and page length can be reset by invoking PGOPT.
6. Horizontal centering, method for printing large matrices, paging, method for printing NaN (not a number), printing a title on each page, and many other options can be selected by invoking WROPT.
7. Output is written to the unit specified by UMACH (see Reference Material).

\section*{Example}

The following example prints all of a \(3 \times 4\) matrix \(A\) where \(a_{i j}=(i+j / 10) 10^{j-3}\).
```

USE WRRRL_INT
IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=10, NCA=4, NRA=3)

```
!
```

    INTEGER I, J
    REAL A(LDA,NCA)
    CHARACTER CLABEL(5)*5, FMT*8, RLABEL(3)*5
    !
DATA FMT/'(W10.6)'/
DATA CLABEL/' ', 'Col 1', 'Col 2', 'Col 3', 'Col 4'/
DATA RLABEL/'Row 1', 'Row 2', 'Row 3'/
!
DO 20 I=1, NRA
DO 10 J=1, NCA
A(I,J) = (I+J*0.1)*10.0**(J-3)
CONTINUE
CONTINUE
CALL WRRRL ('A', A, RLABEL, CLABEL, NRA=NRA, FMT=FMT)
END

```

\section*{Output}
\begin{tabular}{llllr} 
& & \multicolumn{2}{c}{\({ }^{\text {A }}\)} & Col 3
\end{tabular}

\section*{WRIRN}

Prints an integer rectangular matrix with integer row and column labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input)
TITLE set equal to a blank character(s) suppresses printing of the title. Use "\% / " within the title to create a new line. Long titles are automatically wrapped.

MAT - NRMAT by NCMAT matrix to be printed. (Input)

\section*{Optional Arguments}

NRMAT - Number of rows. (Input)
Default: NRMAT \(=\) SIZE (MAT, 1 ).
NCMAT - Number of columns. (Input) Default: NCMAT \(=\operatorname{SIZE}(\) MAT,2 \()\).

LDMAT - Leading dimension of MAT exactly as specified in the dimension statement in the calling program. (Input)
Default: LDMAT \(=\operatorname{SIZE}(\) MAT, 1\()\).

ITRING - Triangle option. (Input)
Default: \(\operatorname{ITRING}=0\).

\section*{ITRING Action}
\(0 \quad\) Full matrix is printed.
1 Upper triangle of MAT is printed, including the diagonal.
2 Upper triangle of MAT excluding the diagonal of MAT is printed.
-1 Lower triangle of MAT is printed, including the diagonal.
-2 Lower triangle of MAT excluding the diagonal of MAT is printed.

\section*{FORTRAN 90 Interface}

Generic: CALL WRIRN (TITLE, MAT [,...])
Specific: The specific interface name is S_WRIRN.

\section*{FORTRAN 77 Interface}

Single: CALL WRIRN (TITLE, NRMAT, NCMAT, MAT, LDMAT, ITRING)

\section*{Description}

Routine WRIRN prints an integer rectangular matrix with the rows and columns labeled \(1,2,3\), and so on. WRIRN can restrict printing to elements of the upper and lower triangles of matrices via the ITRING option. Generally, ITRING \(\neq 0\) is used with symmetric matrices.

In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRMAT to the length of the array and set NCMAT \(=1\). For a row vector, set NRMAT \(=1\) and set NCMAT to the length of the array. In both cases, set LDMAT \(=\) NRMAT and set ITRING \(=0\).

\section*{Comments}
1. All the entries in MAT are printed using a single I format. The field width is determined by the largest absolute entry.
2. Horizontal centering, a method for printing large matrices, paging, printing a title on each page, and many other options can be selected by invoking WROPT.
3. A page width of 78 characters is used. Page width and page length can be reset by invoking PGOPT.
4. Output is written to the unit specified by UMACH (see Reference Material).

\section*{Example}

The following example prints all of a \(3 \times 4\) matrix \(A=\) MAT where \(a_{i j}=10 i+j\).
```

USE WRIRN_INT
IMPLICIT NONE
INTEGER ITRING, LDMAT, NCMAT, NRMAT
PARAMETER (ITRING=0, LDMAT=10, NCMAT=4, NRMAT=3)
INTEGER I, J, MAT(LDMAT,NCMAT)
DO 20 I=1, NRMAT
DO 10 J=1, NCMAT
MAT (I,J) = I*10 + J
CONTINUE
2 0 ~ C O N T I N U E ~
! Write MAT matrix.
CALL WRIRN ('MAT', MAT, NRMAT=NRMAT)
END

```
!
\(!\)
    10

\section*{Output}
\begin{tabular}{rrrrr}
\multicolumn{5}{c}{ MAT } \\
& 1 & 2 & 3 & 4 \\
1 & 11 & 12 & 13 & 14 \\
2 & 21 & 22 & 23 & 24 \\
3 & 31 & 32 & 33 & 34
\end{tabular}

\section*{WRIRL}

Print an integer rectangular matrix with a given format and labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input)
TITLE set equal to a blank character(s) suppresses printing of the title.
MAT - NRMAT by NCMAT matrix to be printed. (Input)
\(\boldsymbol{R L A B E L}\) - CHARACTER * (*) vector of labels for rows of MAT. (Input)
If rows are to be numbered consecutively \(1,2, \ldots\), NRMAT, use
\(\operatorname{RLABEL}(1)=\) ' NUMBER' . If no row labels are desired, use RLABEL(1) \(='\) NONE'.
Otherwise, RLABEL is a vector of length NRMAT containing the labels.
CLABEL — CHARACTER * (*) vector of labels for columns of MAT. (Input)
If columns are to be numbered consecutively \(1,2, \ldots\), NCMAT, use \(\operatorname{CLABEL}(1)='\) NUMBER'. If no column labels are desired, use \(\operatorname{CLABEL}(1)='\) NONE'.
Otherwise, CLABEL(1) is the heading for the row labels, and either CLABEL(2) must be ' NUMBER' or ' \({ }^{\prime}\) NONE' , or CLABEL must be a vector of length

NCMAT +1 with CLABEL \((1+j)\) containing the column heading for the \(j\)-th column.

\section*{Optional Arguments}

NRMAT - Number of rows. (Input)
Default: NRMAT \(=\) SIZE (MAT, 1 ).
NCMAT - Number of columns. (Input)
Default: NCMAT \(=\) SIZE (MAT, 2).

LDMAT - Leading dimension of MAT exactly as specified in the dimension statement in the calling program. (Input)
Default: LDMAT \(=\) SIZE (MAT, 1 ).
ITRING - Triangle option. (Input)
Default: \(\operatorname{ITRING}=0\).

\section*{ITRING Action}
\(0 \quad\) Full matrix is printed.

1 Upper triangle of MAT is printed, including the diagonal.
2 Upper triangle of MAT excluding the diagonal of MAT is printed.
-1 Lower triangle of MAT is printed, including the diagonal.
-2 Lower triangle of MAT excluding the diagonal of MAT is printed.
FMT - Character string containing formats. (Input)
If FMT is set to a blank character(s), the format used is a single I format with field width determined by the largest absolute entry. Otherwise, FMT must contain exactly one set of parentheses and one or more I edit descriptors. For example,
\(\mathrm{FMT}={ }^{\prime}(\mathrm{I} 10)^{\prime}\) specifies this I format for the entire matrix. \(\mathrm{FMT}={ }^{\prime}(2 \mathrm{I} 10,3 \mathrm{I} 5)^{\prime}\) specifies an I10 format for columns 1 and 2 and an I5 format for columns 3, 4 and 5 . If the end of FMT is encountered and if some columns of the matrix remain, format control continues with the first format in FMT. FMT may only contain the I edit descriptor, e.g., the x edit descriptor is not allowed.
Default: FMT \(=\quad\) ' .

\section*{FORTRAN 90 Interface}

Generic: CALL WRIRL (TITLE, MAT, RLABEL, CLABEL [,...])

Specific: The specific interface name is S_WRIRL.

\section*{FORTRAN 77 Interface}

Single: CALL WRIRL (TITLE, NRMAT, NCMAT, MAT, LDMAT, ITRING, FMT, RLABEL, CLABEL)

\section*{Description}

Routine WRIRL prints an integer rectangular matrix (stored in MAT) with row and column labels (specified by RLABEL and CLABEL, respectively), according to a given format (stored in FMT). WRIRL can restrict printing to the elements of upper or lower triangles of matrices via the ITRING option. Generally, ITRING \(\neq 0\) is used with symmetric matrices. In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRMAT to the length of the array and set NCMAT \(=1\). For a row vector, set NRMAT \(=1\) and set NCMAT to the length of the array. In both cases, set LDMAT \(=\) NRMAT, and set \(\operatorname{ITRING}=0\).

\section*{Comments}
1. The output appears in the following form:

TITLE
\begin{tabular}{llll} 
CLABEL(1) & CLABEL(2) & CALBEL(3) & CLABEL 4) \\
RLABEL(1) & Xxxxx & xxxxx & xxxxx \\
RLABEL(2) & Xxxxx & xxxxx & xxxxx
\end{tabular}
2. Use "\%/" within titles or labels to create a new line. Long titles or labels are automatically wrapped.
3. A page width of 78 characters is used. Page width and page length can be reset by invoking PGOPT.
4. Horizontal centering, a method for printing large matrices, paging, printing a title on each page, and many other options can be selected by invoking WROPT.
5. Output is written to the unit specified by UMACH (see the Reference Material).

\section*{Example}

The following example prints all of a \(3 \times 4\) matrix \(A=\) MAT where \(a_{i j}=10 i+j\).
```

USE WRIRL INT
IMPLICIT NONE
INTEGER ITRING, LDMAT, NCMAT, NRMAT
PARAMETER (ITRING=0, LDMAT=10, NCMAT=4, NRMAT=3)
INTEGER I, J, MAT(LDMAT,NCMAT)
CHARACTER CLABEL(5)*5, FMT*8, RLABEL (3)*5
DATA FMT/'(I2)'/

```
!
```

    DATA CLABEL/' ', 'Col 1', 'Col 2', 'Col 3', 'Col 4'/
    DATA RLABEL/'Row 1', 'Row 2', 'Row 3'/
    !
DO 20 I=1, NRMAT
DO 10 J=1, NCMAT
MAT(I,J) = I*10 + J
CONTINUE
10rONTI
! Write MAT matrix.
CALL WRIRL ('MAT', MAT, RLABEL, CLABEL, NRMAT=NRMAT)
END

```

\section*{Output}
\begin{tabular}{cccccc} 
& & \multicolumn{5}{c}{ MAT } \\
Row 1 & Col 1 & Col 2 & Col 3 & Col 4 \\
Row 2 & 11 & 12 & 13 & 14 \\
Row 3 & 21 & 22 & 23 & 24 \\
& 31 & 32 & 33 & 34
\end{tabular}

\section*{WRCRN}

Prints a complex rectangular matrix with integer row and column labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input)
TITLE set equal to a blank character(s) suppresses printing of the title. Use "\% /" within the title to create a new line. Long titles are automatically wrapped.
\(\boldsymbol{A}\) - Complex NRA by NCA matrix to be printed. (Input)

\section*{Optional Arguments}

NRA - Number of rows. (Input)
Default: NRA \(=\operatorname{SIZE}(\mathrm{A}, 1)\).
NCA - Number of columns. (Input)
Default: NCA \(=\operatorname{SIZE}(A, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA \(=\operatorname{SIZE}(A, 1)\).
ITRING - Triangle option. (Input)
Default: \(\operatorname{ITRING}=0\).

ITRING Action

1 Upper triangle of A is printed, including the diagonal.
2 Upper triangle of A excluding the diagonal of A is printed.
-1 Lower triangle of A is printed, including the diagonal.
-2 Lower triangle of A excluding the diagonal of \(A\) is printed.

\section*{FORTRAN 90 Interface}

Generic: CALL WRCRN (TITLE,A \([, \ldots]\) )
Specific: The specific interface names are S_WRCRN and D_WRCRN for two dimensional arrays, and S_WRCRN1D and D_WRCRN1D for one dimensional arrays.

\section*{FORTRAN 77 Interface}

Single: CALL WRCRN (TITLE, NRA, NCA, A, LDA, ITRING)
Double: The double precision name is DWRCRN.

\section*{Description}

Routine WRCRN prints a complex rectangular matrix with the rows and columns labeled \(1,2,3\), and so on. WRCRN can restrict printing to the elements of the upper or lower triangles of matrices via the ITRING option. Generally, ITRING \(\neq 0\) is used with Hermitian matrices.

In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRA to the length of the array, and set NCA \(=1\). For a row vector, set NRA \(=1\), and set NCA to the length of the array. In both cases, set LDA \(=\) NRA, and set \(I T R I N G=0\).

\section*{Comments}
1. A single \(D, E\), or \(F\) format is chosen automatically in order to print 4 significant digits for the largest real or imaginary part in absolute value of all the complex numbers in A. Routine WROPT can be used to change the default format.
2. Horizontal centering, a method for printing large matrices, paging, method for printing NaN (not a number), and printing a title on each page can be selected by invoking WROPT.
3. A page width of 78 characters is used. Page width and page length can be reset by invoking subroutine PGOPT.
4. Output is written to the unit specified by UMACH (see Reference Material).

\section*{Example}

This example prints all of a \(3 \times 4\) complex matrix \(A\) with elements
\[
a_{m n}=m+n i, \text { where } i=\sqrt{-1}
\]

USE WRCRN_INT
IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=10, NCA=4, NRA=3)
!
INTEGER I, J
COMPLEX A(LDA,NCA), CMPLX
INTRINSIC CMPLX
!
DO 20 I=1, NRA DO \(10 \mathrm{~J}=1\), NCA
\(A(I, J)=\operatorname{CMPLX}(I, J)\)
CONTINUE
10
20 CONTINUE
! Write A matrix.
CALL WRCRN ('A', A, NRA=NRA)
END

\section*{Output}


\section*{WRCRL}

Prints a complex rectangular matrix with a given format and labels.

\section*{Required Arguments}

TITLE - Character string specifying the title. (Input) TITLE set equal to a blank character(s) suppresses printing of the title.
\(\boldsymbol{A}\) - Complex NRA by NCA matrix to be printed. (Input)
\(\boldsymbol{R L A B E L}\) - CHARACTER * (*) vector of labels for rows of A. (Input)
If rows are to be numbered consecutively \(1,2, \ldots\), NRA, use RLABEL \((1)={ }^{\prime}\) NUMBER' . If no row labels are desired, use RLABEL \((1)={ }^{\prime}\) NONE \(^{\prime}\). Otherwise, RLABEL is a vector of length NRA containing the labels.

CLABEL - CHARACTER * (*) vector of labels for columns of A. (Input) If columns are to be numbered consecutively \(1,2, \ldots\), NCA, use

CLABEL(1) \(=\) ' \(\operatorname{NUMBER}\) ' . If no column labels are desired, use CLABEL(1) = ' NONE' . Otherwise, CLABEL(1) is the heading for the row labels, and either CLABEL(2) must be ' NUMBER' or ' \({ }^{\text {NONE' }}\), or CLABEL must be a vector of length NCA +1 with \(\operatorname{CLABEL}(1+j)\) containing the column heading for the \(j\)-th column.

\section*{Optional Arguments}

NRA - Number of rows. (Input)
Default: NRA \(=\operatorname{SIZE}(\mathrm{A}, 1)\).
NCA - Number of columns. (Input)
Default: NCA \(=\operatorname{SIZE}(A, 2)\).
LDA - Leading dimension of A exactly as specified in the dimension statement in the calling program. (Input)
Default: LDA \(=\operatorname{SIZE}(\mathrm{A}, 1)\).
ITRING - Triangle option. (Input)
Default: \(\operatorname{ITRING}=0\).

\section*{ITRING}

Action

0
1

2
\(-1\)
-2 Lower triangle of A excluding the diagonal of \(A\) is printed.
FMT - Character string containing formats. (Input)
If FMT is set to a blank character(s), the format used is specified by WROPT. Otherwise, FMT must contain exactly one set of parentheses and one or more edit descriptors. Because a complex number consists of two parts (a real and an imaginary part), two edit descriptors are used for printing a single complex number. \(\mathrm{FMT}={ }^{\prime}(\mathrm{E} 10.3\), F10.3)' specifies an \(E\) format for the real part and an \(F\) format for the imaginary part. \(F M T={ }^{\prime}(\text { F10.3 })^{\prime}\) uses an \(F\) format for both the real and imaginary parts. If the end of FMT is encountered and if all columns of the matrix have not been printed, format control continues with the first format in FMT. Even though the matrix A is complex, an I format can be used to print the integer parts of the real and imaginary components of each complex number. The most useful formats are special formats, called the "V and \(w\) formats," that can be used to specify pretty formats automatically. Set FMT \(={ }^{\prime}(\mathrm{V} 10.4)^{\prime}\) if you want a single \(D, E\), or \(F\) format selected automatially with field width 10 and with 4 significant digits. Set FMT \(=^{\prime}(\text { W10.4 })^{\prime}\) if you want a single D, E, F, or I format selected automatically with field width 10 and with 4 significant
digits. While the v format prints trailing zeroes and a trailing decimal point, the w format does not. See Comment 4 for general descriptions of the \(v\) and \(w\) formats. FMT may contain only \(D, E, F, G, I, V\), or \(W\) edit descriptors, e.g., the \(X\) descriptor is not allowed.
Default: FMT \(=\) ' \(\quad\).

\section*{FORTRAN 90 Interface}

Generic: CALL WRCRL (TITLE, A, RLABEL, CLABEL[,...])
Specific: The specific interface names are S_WRCRL and D_WRCRL for two dimensional arrays, and S_WRCRL1D and D_WRCRL1D for one dimensional arrays.

\section*{FORTRAN 77 Interface}

Single: CALL WRCRL (TITLE, NRA, NCA, A, LDA, ITRING, FMT, RLABEL, CLABEL)

Double: \(\quad\) The double precision name is DWRCRL.

\section*{Description}

Routine WRCRL prints a complex rectangular matrix (stored in \(A\) ) with row and column labels (specified by RLABEL and CLABEL, respectively) according to a given format (stored in FMT). Routine WRCRL can restrict printing to the elements of upper or lower triangles of matrices via the ITRING option. Generally, the ITRING \(\neq 0\) is used with Hermitian matrices.

In addition, one-dimensional arrays can be printed as column or row vectors. For a column vector, set NRA to the length of the array, and set NCA \(=1\). For a row vector, set NRA \(=1\), and set NCA to the length of the array. In both cases, set LDA \(=\) NRA, and set \(\operatorname{ITRING}=0\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of \(W 2\) CRL/DW2CRL. The reference is:

CALL W2CRL (TITLE, NRA, NCA, A, LDA, ITRING, FMT, RLABEL, CLABEL, CHWK)
The additional argument is:
CHWK - CHARACTER * 10 work vector of length \(2 *\) NCA. This workspace is referenced only if all three conditions indicated at the beginning of this comment are met. Otherwise, CHWK is not referenced and can be a CHARACTER * 10 vector of length one.
2. The output appears in the following form:
```

CLABEL (1)
RLABEL (1) (xxxxx, xxxxx) (xxxxx, xxxxx) (xxxxx, xxxxx)
CLABEL(2) CLABEL(3) CLABEL(4)
RLABEL (2) (xxxxx, xxxxx) (xxxxx, xxxxx) (xxxxx, xxxxx)

```
3. Use " \(\% /\) /" within titles or labels to create a new line. Long titles or labels are automatically wrapped.
4. For printing numbers whose magnitudes are unknown, the G format in FORTRAN is useful; however, the decimal points will generally not be aligned when printing a column of numbers. The v and w formats are special formats used by this routine to select a \(D, E, F\), or I format so that the decimal points will be aligned. The \(V\) and \(W\) formats are specified as Vn.d and Wn.d. Here, \(n\) is the field width, and \(d\) is the number of significant digits generally printed. Valid values for \(n\) are \(3,4, \ldots, 40\). Valid values for \(d\) are \(1,2, \ldots, n-2\). If FMT specifies one format and that format is a v or w format, all elements of the matrix A are examined to determine one FORTRAN format for printing. If FMT specifies more than one format, FORTRAN formats are generated separately from each \(v\) or \(w\) format.
5. A page width of 78 characters is used. Page width and page length can be reset by invoking PGOPT.
6. Horizontal centering, a method for printing large matrices, paging, method for printing NaN (not a number), printing a title on each page, and may other options can be selected by invoking WROPT.
7. Output is written to the unit specified by UMACH (see the Reference Material).

\section*{Example}

The following example prints all of a \(3 \times 4\) matrix \(A\) with elements
\[
a_{m n}=(m+.123456)+n i, \text { where } i=\sqrt{-1}
\]
```

USE WRCRL INT
IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=10, NCA=4, NRA=3)
INTEGER I, J
COMPLEX A(LDA,NCA), CMPLX
CHARACTER CLABEL(5)*5, FMT*8, RLABEL (3)*5
INTRINSIC CMPLX
DATA FMT/'(W12.6)'/
DATA CLABEL/' ', 'Col 1', 'Col 2', 'Col 3', 'Col 4'/
DATA RLABEL/'Row 1', 'Row 2', 'Row 3'/
DO 20 I=1, NRA

```
\(!\)
```

        DO 10 J=1, NCA
                A(I,J)=CMPLX(I,J) + 0.123456
    CONTINUE
    CONTINUE
    ! Write A matrix.
CALL WRCRL ('A', A, RLABEL, CLABEL, NRA=NRA, FMT=FMT)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline & & & & A & & & \\
\hline & & & & Col 1 & & & Col 2 \\
\hline Row & 1 & 1 & 1.12346, & 1.00000) & 1 & 1.12346, & \(2.00000)\) \\
\hline Row & 2 & ( & 2.12346, & 1.00000) & \((\) & 2.12346, & \(2.00000)\) \\
\hline Row & 3 & ( & 3.12346, & \(1.00000)\) & ( & 3.12346 , & \(2.00000)\) \\
\hline & & & & Col 3 & & & Col 4 \\
\hline Row & 1 & ( & 1.12346, & \(3.00000)\) & \((\) & 1.12346, & 4.00000 ) \\
\hline Row & 2 & ( & 2.12346, & \(3.00000)\) & ( & 2.12346, & \(4.00000)\) \\
\hline Row & 3 & ( & 3.12346 , & 3.00000 ) & ( & 3.12346 , & \(4.00000)\) \\
\hline
\end{tabular}

\section*{WROPT}

Sets or retrieves an option for printing a matrix.

\section*{Required Arguments}

IOPT - Indicator of option type. (Input)

\section*{IOPT Description of Option Type}
\(-1,1 \quad\) Horizontal centering or left justification of matrix to be printed
\(-2,2 \quad\) Method for printing large matrices
\(-3,3 \quad\) Paging
-4, 4 Method for printing NaN (not a number), and negative and positive machine infinity.
\(-5,5\) Title option
\(-6,6\) Default format for real and complex numbers
\(-7,7 \quad\) Spacing between columns
-8, 8 Maximum horizontal space reserved for row labels
\(-9,9\) Indentation of continuation lines for row labels
\(-10,10\) Hot zone option for determining line breaks for row labels
-11, 11 Maximum horizontal space reserved for column labels
\(-12,12\) Hot zone option for determining line breaks for column labels
-13, 13 Hot zone option for determining line breaks for titles
\(-14,14\) Option for the label that appears in the upper left hand corner that can be used as a heading for the row numbers or a label for the column headings for \(W R * * N\) routines
\(-15,15\) Option for skipping a line between invocations of \(W R * * N\) routines, provided a new page is not to be issued
\(-16,16\) Option for vertical alignment of the matrix values relative to the associated row labels that occupy more than one line
\(0 \quad\) Reset all the current settings saved in internal variables back to their last setting made with an invocation of WROPT with ISCOPE \(=1\). (This option is used internally by routines printing a matrix and is not useful otherwise.)

If IOPT is negative, ISETNG and ISCOPE are input and are saved in internal variables. If IOPT is positive, ISETNG is output and receives the currently active setting for the option (if ISCOPE \(=0\) ) or the last global setting for the option (if ISCOPE \(=1\) ). If IOPT \(=0\), ISETNG and ISCOPE are not referenced.

ISETNG - Setting for option selected by IOPT. (Input, if IOPT is negative; output, if IOPT is positive; not referenced if \(I O P T=0\) )
\begin{tabular}{lll} 
IOPT & ISETNG & Meaning \\
\(-1,1\) & 0 & Matrix is left justified \\
\(-2,2\) & 1 & Matrix is centered horizontally on page \\
& 0 & \begin{tabular}{l} 
A complete row is printed before the next row is \\
printed. Wrapping is used if necessary.
\end{tabular}
\end{tabular}
\[
-3
\]
\(-3,3 \quad-2\)

Here, \(m\) is a positive integer. Let \(n_{1}\) be the maximum number of columns beginning with column 1 that fit across the page (as determined by the widths of the printing formats). First, columns 1 through \(n_{1}\) are printed for rows 1 through \(m\). Let \(n_{2}\) be the maximum number of columns beginning with column \(n_{1}+1\) that fit across the page.
Second, columns \(n_{1}+1\) through \(n_{1}+n_{2}\) are printed for rows 1 through m . This continues until the last columns are printed for rows 1 through m . Printing continues in this fashion for the next m rows, etc.
Printing begins on the next line, and no paging occurs.

Paging is on. Every invocation of a WR*** routine begins on a new page, and paging occurs within each invocation as is needed

Paging is on. The first invocation of a WR*** routine begins on a new page, and subsequent paging occurs as is needed. With this option, every invocation of a WR*** routine ends with a call to WROPT to reset this option to k , a positive integer giving the number of lines printed on the current page.
Here, k is a positive integer. Paging is on, and k lines have been printed on the current page. If \(k\) is less than the page length IPAGE (see PGOPT), then IPAGE - \(k\) lines are printed before a new page instruction is issued. If \(k\) is greater than or equal to IPAGE, then the first invocation of a WR*** routine begins on a new page. In any case, subsequent paging occurs as is needed. With this option, every invocation of a WR*** routine ends with a call to WROPT to reset the value of \(k\).
NaN is printed as a series of decimal points, negative machine infinity is printed as a series of minus signs, and positive machine infinity is printed as a series of plus signs.
NaN is printed as a series of blank characters, negative machine infinity is printed as a series of minus signs, and positive machine infinity is printed as a series of plus signs.
NaN is printed as " NaN ," negative machine infinity is printed as "-Inf" and positive machine infinity is printed as "Inf."
\begin{tabular}{|c|c|c|}
\hline & 3 & NaN is printed as a series of blank characters, negative machine infinity is printed as "-Inf," and positive machine infinity is printed as "Inf." \\
\hline \multirow[t]{2}{*}{\(-5,5\)} & 0 & Title appears only on first page. \\
\hline & 1 & Title appears on the first page and all continuation pages. \\
\hline \multirow[t]{6}{*}{-6, 6} & 0 & Format is (W10.4). See Comment 2. \\
\hline & 1 & Format is (W12.6). See Comment 2. \\
\hline & 2 & Format is (1PE12.5). \\
\hline & 3 & Format is Vn. 4 where the field width n is determined. See Comment 2. \\
\hline & 4 & Format is \(V \mathrm{n} .6\) where the field width \(n\) is determined. Comment 2. \\
\hline & 5 & Format is 1 PEn. d where \(\mathrm{n}=\mathrm{d}+7\), and \(\mathrm{d}+1\) is the maximum number of significant digits. \\
\hline -7, 7 & \(\mathrm{K}_{1}\) & Number of characters left blank between columns. \(\mathrm{K}_{1}\) must be between 0 and 5 , inclusively. \\
\hline \(-8,8\) & \(\mathrm{K}_{2}\) & Maximum width (in characters) reserved for row labels. \(\mathrm{K}_{2}=0\) means use the default. \\
\hline -9, 9 & \(\mathrm{K}_{3}\) & Number of characters used to indent continuation lines for row labels. \(\mathrm{k}_{3}\) must be between 0 and 10 , inclusively. \\
\hline \(-10,10\) & K \({ }_{4}\) & Width (in characters) of the hot zone where line breaks in row labels can occur. \(\mathrm{k}_{4}=0\) means use the default. k must not exceed 50 . \\
\hline \(-11,11\) & K 5 & Maximum width (in characters) reserved for column labels. \(\mathrm{K}_{5}=0\) means use the default. \\
\hline \(-12,12\) & \(\mathrm{K}_{6}\) & Width (in characters) of the hot zone where line breaks in column labels can occur. \(\mathrm{k}_{6}=0\) means use the default. \(\mathrm{k}_{6}\) must not exceed 50. \\
\hline \(-13,13\) & \(\mathrm{K}_{7}\) & Width (in characters) of the hot zone where line breaks in titles can occur. \(\mathrm{k}_{7}\) must be between 1 and 50 , inclusively. \\
\hline \multirow[t]{2}{*}{-14} & 0 & There is no label in the upper left hand corner. \\
\hline & 1 & The label in the upper left hand corner is "Component" if a row vector or column vector is printed; the label is "Row/Column" if both the number of rows and columns are greater than one; otherwise, there is no label. \\
\hline
\end{tabular}
\(\left.\begin{array}{lll}-15 & 0 & \begin{array}{l}\text { A blank line is printed on each invocation of a } \\
\text { WR } * * N \text { r routine before the matrix title provided a } \\
\text { new page is not to be issued. }\end{array} \\
\text { A blank line is not printed on each invocation of a } \\
\text { WR**N routine before the matrix title. }\end{array}\right\}\)\begin{tabular}{l} 
The matrix values are aligned vertically with the \\
last line of the associated row label for the case \\
IOPT = 2 and ISET is positive.
\end{tabular}

ISCOPE - Indicator of the scope of the option. (Input if IOPT is nonzero; not referenced if IOPT \(=0\) )

\section*{ISCOPE Action}

0 Setting is temporarily active for the next invocation of a WR*** matrix printing routine.

1 Setting is active until it is changed by another invocation of WROPT.

\section*{FORTRAN 90 Interface}

Generic: CALL WROPT (IOPT, ISETNG, ISCOPE)
Specific: The specific interface name is wropt.

\section*{FORTRAN 77 Interface}

Single: CALL WROPT (IOPT, ISETNG, ISCOPE)

\section*{Description}

Routine WROPT allows the user to set or retrieve an option for printing a matrix. The options controlled by WROPT include the following: horizontal centering, a method for printing large matrices, paging, method for printing NaN (not a number) and positive and negative machine infinities, printing titles, default formats for numbers, spacing between columns, maximum widths reserved for row and column labels, indentation of row labels that continue beyond one line, widths of hot zones for breaking of labels and titles, the default heading for row labels, whether to print a blank line between invocations of routines, and vertical alignment of matrix entries with respect to row labels continued beyond one line. ( NaN and positive and negative machine infinities can be retrieved by АМАСН and DMACH that are documented in the section "MachineDependent Constants" in the Reference Material.) Options can be set globally (ISCOPE =1) or temporarily for the next call to a printing routine ( \(\mathrm{ISCOPE}=0\) ).

\section*{Comments}
1. This program can be invoked repeatedly before using a \(W R * * *\) routine to print a matrix. The matrix printing routines retrieve these settings to determine the printing options. It is not necessary to call WROPT if a default value of a printing option is desired. The defaults are as follows.
\begin{tabular}{lll} 
IOPT & \begin{tabular}{c} 
Default \\
Value for \\
ISET
\end{tabular} &
\end{tabular} \begin{tabular}{l} 
Meaning \\
1
\end{tabular}

\section*{IOPT Default Meaning Value for ISET}

160
The matrix values are aligned vertically with the last line of the associated row label.

For \(I O P T=8\), the default depends on the current value for the page width, IPAGEW (see PGOPT).
2. The \(v\) and \(w\) formats are special formats that can be used to select a D, E, F, or I format so that the decimal points will be aligned. The v and w formats are specified as Vn.d and \(W n . d\). Here, \(n\) is the field width and \(d\) is the number of significant digits generally printed. Valid values for \(n\) are \(3,4, \ldots, 40\). Valid values for \(d\) are \(1,2, \ldots, n-2\). While the v format prints trailing zeroes and a trailing decimal point, the w format does not.

\section*{Example}

The following example illustrates the effect of WROPT when printing a \(3 \times 4\) real matrix \(A\) with WRRRN where \(a_{i j}=i+j / 10\). The first call to WROPT sets horizontal printing so that the matrix is first printed horizontally centered on the page. In the next invocation of WRRRN, the left-justification option has been set via routine WROPT so the matrix is left justified when printed. Finally, because the scope of left justification was only for the next call to a printing routine, the last call to WRRRN results in horizontally centered printing.
```

USE WROPT INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=10, NCA=4, NRA=3)
INTEGER I, IOPT, ISCOPE, ISETNG, J
REAL A(LDA,NCA)
DO 20 I=1, NRA
DO 10 J=1, NCA
A(I,J) = I + J*0.1
10 CONTINUE
2 0 ~ C O N T I N U E ~
IOPT = -1
ISETNG = 1
ISCOPE = 1
CALL WROPT (IOPT, ISETNG, ISCOPE)
Write A matrix.
CALL WRRRN ('A', A, NRA=NRA)
Activate left justification.
Scope is local.

```
!
```

ISETNG = 0
ISCOPE = 0
CALL WROPT (IOPT, ISETNG, ISCOPE)
CALL WRRRN ('A', A, NRA=NRA)
CALL WRRRN ('A', A, NRA=NRA)
END

```

Output
\begin{tabular}{rrrrr} 
& & A & & \\
& 1 & 2 & 3 & 4 \\
1 & 1.100 & 1.200 & 1.300 & 1.400 \\
2 & 2.100 & 2.200 & 2.300 & 2.400 \\
3 & 3.100 & 3.200 & 3.300 & 3.400
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{A} \\
\hline & 1 & 2 & 3 & 4 & & & \\
\hline 1 & 1.100 & 1.200 & 1.300 & 1.400 & & & \\
\hline 2 & 2.100 & 2.200 & 2.300 & 2.400 & & & \\
\hline 3 & 3.100 & 3.200 & 3.300 & 3.400 & & & \\
\hline & & & \multicolumn{5}{|c|}{A} \\
\hline & & & & 1 & 2 & 3 & 4 \\
\hline & & & 1 & 1.100 & 1.200 & 1.300 & 1.400 \\
\hline & & & 2 & 2.100 & 2.200 & 2.300 & 2.400 \\
\hline & & & 3 & 3.100 & 3.200 & 3.300 & 3.400 \\
\hline
\end{tabular}

\section*{PGOPT}

Sets or retrieves page width and length for printing.

\section*{Required Arguments}

IOPT - Page attribute option. (Input)

\section*{IOPT Description of Attribute}
\(-1,1 \quad\) Page width.
\(-2,2 \quad\) Page length.
Negative values of IOPT indicate the setting IPAGE is input. Positive values
of IOPT indicate the setting IPAGE is output.
IPAGE - Value of page attribute. (Input, if IOPT is negative; output, if IOPT is positive.)

\section*{IOPT Description of Attribute Settings for IPAGE}
\(-1,1\) Page width (in characters) \(10,11, \ldots\)
\[
-2,2 \text { Page length (in lines) } \quad 10,11, \ldots
\]

\section*{FORTRAN 90 Interface}

Generic: CALL PGOPT (IOPT, IPAGE)
Specific: The specific interface name is PGOPT.

\section*{FORTRAN 77 Interface}

Single: CALL PGOPT (IOPT, IPAGE)

\section*{Description}

Routine PGOPT is used to set or retrieve the page width or the page length for routines that perform printing.

\section*{Example}

The following example illustrates the use of PGOPT to set the page width at 20 characters. Routine WRRRN is then used to print a \(3 \times 4\) matrix \(A\) where \(a_{i j}=i+j / 10\).
```

USE PGOPT INT
USE WRRRN_INT
IMPLICIT NONE
INTEGER ITRING, LDA, NCA, NRA
PARAMETER (ITRING=0, LDA=3, NCA=4, NRA=3)
INTEGER I, IOPT, IPAGE, J
REAL A(LDA,NCA)
DO 20 I=1, NRA
DO 10 J=1, NCA
A(I,J) = I + J*0.1
CONTINUE
CONTINUE
IOPT = -1
IPAGE = 20
CALL PGOPT (IOPT, IPAGE)
CALL WRRRN ('A', A)
END

```
!
!

\section*{Output}

A
\(1 \quad 2\)
\(1 \quad 1.100 \quad 1.200\)
\(2.100 \quad 2.200\)
\begin{tabular}{rrr}
3 & 3.100 & 3.200 \\
& 3 & 4 \\
1 & 1.300 & 1.400 \\
2 & 2.300 & 2.400 \\
3 & 3.300 & 3.400
\end{tabular}

\section*{PERMU}

Rearranges the elements of an array as specified by a permutation.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Real vector of length N containing the array to be permuted. (Input)

IPERMU - Integer vector of length N containing a permutation IPERMU(1), \(\ldots, \operatorname{IPERMU}(\mathrm{N})\) of the integers \(1, \ldots, \mathrm{~N}\). (Input)

XPERMU - Real vector of length N containing the array x permuted. (Output) If \(X\) is not needed, \(X\) and XPERMU can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Length of the arrays X and XPERMU. (Input)
Default: N = SIZE (IPERMU,1).
IPATH - Integer flag. (Input)
Default: IPATH \(=1\).
IPATH \(=1\) means IPERMU represents a forward permutation, i.e., \(X(\operatorname{IPERMU}(I))\) is moved to \(\operatorname{XPERMU}(I)\). IPATH \(=2\) means IPERMU represents a backward permutation, i.e., \(X(I)\) is moved to XPERMU (IPERMU(I)).

\section*{FORTRAN 90 Interface}

Generic: CALL PERMU (X, IPERMU, XPERMU [, ...])

Specific: The specific interface names are S_PERMU and D_PERMU.

\section*{FORTRAN 77 Interface}

Single: CALL PERMU (N, X, IPERMU, IPATH, XPERMU)
Double: The double precision name is DPERMU.

\section*{Description}

Routine PERMU rearranges the elements of an array according to a permutation vector. It has the option to do both forward and backward permutations.

\section*{Example}

This example rearranges the array \(X\) using IPERMU; forward permutation is performed.
```

USE PERMU_INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER IPATH, N
PARAMETER (IPATH=1,N=4)
!
INTEGER IPERMU(N), J, NOUT
REAL X(N), XPERMU (N)
Set values for X, IPERMU
X = ( 5.0 6.0 1.0 4.0 )
IPERMU =(($$
\begin{array}{llll}{3}&{1}&{4}&{2}\end{array}
$$)
DATA X/5.0, 6.0, 1.0, 4.0/, IPERMU/3, 1, 4, 2/
Permute X into XPERMU
CALL PERMU (X, IPERMU, XPERMU)
Get output unit number
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) (XPERMU (J), J=1,N)
!
99999 FORMAT (' The output vector is:', /, 10(1X,F10.2))
END

```

\section*{Output}

The Output vector is:
\begin{tabular}{llll}
1.00 & 5.00 & 4.00 & 6.00
\end{tabular}

\section*{PERMA}

Permutes the rows or columns of a matrix.

\section*{Required Arguments}
\(\boldsymbol{A}\) - NRA by NCA matrix to be permuted. (Input)
IPERMU - Vector of length \(K\) containing a permutation IPERMU(1), ..., I PERMU(K) of the integers \(1, \ldots, K\) where \(K=\) NRA if the rows of \(A\) are to be permuted and \(K=N C A\) if the columns of A are to be permuted. (Input)

APER - NRA by NCA matrix containing the permuted matrix. (Output)
If A is not needed, A and APER can share the same storage locations.

\section*{Optional Arguments}

NRA - Number of rows. (Input)
Default: NRA \(=\operatorname{SIZE}(A, 1)\).

NCA - Number of columns. (Input)
Default: NCA \(=\operatorname{SIZE}(A, 2)\).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\) SIZE (A, 1).
IPATH - Option parameter. (Input)
IPATH \(=1\) means the rows of A will be permuted. IPATH \(=2\) means the columns of A will be permuted.
Default: IPATH \(=1\).

LDAPER - Leading dimension of APER exactly as specified in the dimension statement of the calling program. (Input)
Default: LDAPER = SIZE (APER,1).

\section*{FORTRAN 90 Interface}

Generic: CALL PERMA (A, IPERMU, APER [, ...])
Specific: The specific interface names are S_PERMA and D_PERMA.

\section*{FORTRAN 77 Interface}

Single: CALL PERMA (NRA, NCA, A, LDA, IPERMU, IPATH, APER, LDAPER)

Double: The double precision name is DPERMA.

\section*{Description}

Routine PERMA interchanges the rows or columns of a matrix using a permutation vector such as the one obtained from routines SVRBP or SVRGP.

The routine PERMA permutes a column (row) at a time by calling PERMU. This process is continued until all the columns (rows) are permuted. On completion, let \(B=\operatorname{APER}\) and \(p_{i}=\operatorname{IPERMU}(\mathrm{I})\), then
\[
B_{i j}=A_{p_{i} j}
\]
for all \(i, j\).

\section*{Comments}
1. Workspace may be explicitly provided, if desired, by use of P2RMA/DP2RMA. The reference is:

CALL P2RMA (NRA, NCA, A, LDA, IPERMU, IPATH, APER, LDAPER, WORK)
The additional argument is:
WORK - Real work vector of length NCA.

\section*{Example}

This example permutes the columns of a matrix \(A\).
```

    USE PERMA_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER IPATH, LDA, LDAPER, NCA, NRA
    PARAMETER (IPATH=2, LDA=3, LDAPER=3, NCA=5, NRA=3)
    INTEGER I, IPERMU(5), J, NOUT
    REAL A(LDA,NCA), APER(LDAPER,NCA)
                            Set values for A, IPERMU
                            A =( 3.0 5.0 1.0 2.0 4.0 )
                                    ( 3.0 5.0 1.0 2.0 4.0 )
                                    ( 3.0 5.0 1.0 2.0 4.0)
            IPERMU =( (\begin{array}{lllll}{3}&{4}&{1}&{5}&{2}\end{array})
    DATA A/3*3.0, 3*5.0, 3*1.0, 3*2.0, 3*4.0/, IPERMU/3, 4, 1, 5, 2/
            Perform column permutation on A,
                    giving APER
    CALL PERMA (A, IPERMU, APER, IPATH=IPATH)
            Get output unit number
    CALL UMACH (2, NOUT)
    WRITE (NOUT,99999) ((APER(I, J), J=1,NCA),I=1,NRA)
    !
99999 FORMAT (' The output matrix is:', /, 3(5F8.1,/))
END

```

\section*{Output}
\begin{tabular}{ccccc} 
The Output matrix is: & & \\
1.0 & 2.0 & 3.0 & 4.0 & 5.0 \\
1.0 & 2.0 & 3.0 & 4.0 & 5.0 \\
1.0 & 2.0 & 3.0 & 4.0 & 5.0
\end{tabular}

\section*{SORT_REAL}

Sorts a rank-1 array of real numbers \(x\) so the \(y\) results are algebraically nondecreasing, \(y_{1} \leq y_{2} \leq \ldots y_{n}\).

\section*{Required Arguments}
\(\boldsymbol{X}\) - Rank-1 array containing the numbers to be sorted. (Output)
\(\boldsymbol{Y}\) - Rank-1 array containing the sorted numbers. (Output)

\section*{Optional Arguments}

NSIZE = n (Input)
Uses the sub-array of size n for the numbers.
Default value: \(\mathrm{n}=\operatorname{SIZE}(\mathrm{x})\)
IPERM \(=\) iperm (Input/Output)
Applies interchanges of elements that occur to the entries of iperm(:). If the values iperm ( i\()=\mathrm{i}, \mathrm{i}=1, \mathrm{n}\) are assigned prior to call, then the output array is moved to its proper order by the subscripted array assignment \(\mathrm{y}=\mathrm{x}(\mathrm{iperm}(1: \mathrm{n}))\).

ICYCLE = icycle (Output)
Permutations applied to the input data are converted to cyclic interchanges. Thus, the output array y is given by the following elementary interchanges, where \(:=:\) denotes a swap:
j = icycle(i)
\(y(j):=: y(i), i=1, n\)
IOPT = iopt (:) (Input)
Derived type array with the same precision as the input matrix; used for passing optional data to the routine. The options are as follows:
\begin{tabular}{|l|l|l|}
\hline \multicolumn{3}{|c|}{ Packaged Options for SORT_REAL } \\
\hline Option Prefix =? & Option Name & Option Value \\
\hline s_, d_ & Sort_real_scan_for_NaN & 1 \\
\hline
\end{tabular}
iopt(IO) = ?_options(?_sort_real_scan_for_NaN, ?_dummy)
Examines each input array entry to find the first value such that
isNaN(x(i)) == .true.
See the isNan() function, Chapter 10.
Default: Does not scan for NaNs.

\section*{FORTRAN 90 Interface}

Generic: CALL SORT_REAL (X, Y [, ...])

Specific: The specific interface names are \(S_{\text {_S }}\) SORT_REAL and D_SORT_REAL.

\section*{Description}

For a detailed description, see the "Description" section of routine SVRGN, which appears later in this chapter.

\section*{Fatal and Terminal Error Messages}

See the messages.gls file for error messages for SORT_REAL. These error messages are numbered 561-567; 581-587.

\section*{Example 1: Sorting an Array}

An array of random numbers is obtained. The values are sorted so they are nondecreasing.
```

    use sort real int
    use rand_gen_int
    implicit none
    ! This is Example 1 for SORT_REAL.
integer, parameter :: n=100
real(kind(1e0)), dimension(n) :: x, y
! Generate random data to sort.
call rand gen(x)
! Sort the data so it is non-decreasing.
call sort_real(x, y)
! Check that the sorted array is not decreasing.
if (count (y(1:n-1) > y(2:n)) == 0) then
write (*,*) 'Example 1 for SORT_REAL is correct.'
end if
end

```

\section*{Output}

Example 1 for SORT_REAL is correct.

\section*{Additional Examples}

\section*{Example 2: Sort and Final Move with a Permutation}

A set of \(n\) random numbers is sorted so the results are nonincreasing. The columns of an \(n \times n\) random matrix are moved to the order given by the permutation defined by the interchange of the entries. Since the routine sorts the results to be algebraically nondecreasing, the array of negative values is used as input. Thus, the negative value of the sorted output order is nonincreasing. The optional argument "iperm=" records the final order and is used to move the matrix columns to
that order. This example illustrates the principle of sorting record keys, followed by direct movement of the records to sorted order.
```

    use sort_real_int
    use rand_gen_int
    implicit none
    ! This is Example 2 for SORT_REAL.
integer i
integer, parameter :: n=100
integer ip(n)
real(kind(le0)) a(n,n), x(n), y(n), temp(n*n)
! Generate a random array and matrix of values.
call rand_gen(x)
call rand_gen(temp)
a = reshape(temp,(/n,n/))
! Initialize permutation to the identity.
do i=1, n
ip(i) = i
end do
! Sort using negative values so the final order is
! non-increasing.
call sort_real(-x, y, iperm=ip)
! Final movement of keys and matrix columns.
y = x(ip (1:n))
a = a(:,ip(1:n))
! Check the results.
if (count (y(1:n-1) < y(2:n)) == 0) then
write (*,*) 'Example 2 for SORT_REAL is correct.'
end if
end

```

\section*{Output}

Example 2 for SORT_REAL is correct.

\section*{SVRGN}

Sorts a real array by algebraically increasing value.

\section*{Required Arguments}
\(\boldsymbol{R A}\) - Vector of length N containing the array to be sorted. (Input)
\(\boldsymbol{R B}\) - Vector of length \(N\) containing the sorted array. (Output)
If RA is not needed, \(R A\) and \(R B\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{RA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVRGN (RA, RB \([, \ldots]\) )
Specific: The specific interface names are S_SVRGN and D_SVRGN.

\section*{FORTRAN 77 Interface}

Single: CALL SVRGN (N, RA, RB)
Double: The double precision name is DSVRGN.

\section*{Description}

Routine SVRGN sorts the elements of an array, \(A\), into ascending order by algebraic value. The array \(A\) is divided into two parts by picking a central element \(T\) of the array. The first and last elements of \(A\) are compared with \(T\) and exchanged until the three values appear in the array in ascending order. The elements of the array are rearranged until all elements greater than or equal to the central element appear in the second part of the array and all those less than or equal to the central element appear in the first part. The upper and lower subscripts of one of the segments are saved, and the process continues iteratively on the other segment. When one segment is finally sorted, the process begins again by retrieving the subscripts of another unsorted portion of the array. On completion, \(A_{j} \leq A_{i}\) for \(j<i\). For more details, see Singleton (1969), Griffin and Redish (1970), and Petro (1970).

\section*{Example}

This example sorts the 10-element array RA algebraically.
```

USE SVRGN INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER N, NOUT, J
PARAMETER (N=10)
REAL RA(N), RB(N)
RA=( (-1.0 2.0 -3.0 4.0 -5.0 6.0 -7.0 8.0 -9.0
DATA RA/-1.0, 2.0, -3.0, 4.0, -5.0, 6.0, -7.0, 8.0, -9.0, 10.0/

```
```

! Sort RA by algebraic value into RB
!
CALL SVRGN (RA, RB)
CALL UMACH (2,NOUT)
WRITE (NOUT, 99999) (RB(J),J=1,N)
!
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ o u t p u t ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , F 5 . 1 ) )
END

```

\section*{Output}
```

The Output vector is:

```
\(\begin{array}{llllllllllllllll}-9.0 & -7.0 & -5.0 & -3.0 & -1.0 & 2.0 & 4.0 & 6.0 & 8.0 & 10.0\end{array}\)

\section*{SVRGP}

Sorts a real array by algebraically increasing value and return the permutation that rearranges the array.

\section*{Required Arguments}
\(\boldsymbol{R A}\) - Vector of length N containing the array to be sorted. (Input)
\(\boldsymbol{R B}\) - Vector of length N containing the sorted array. (Output)
If \(R A\) is not needed, \(R A\) and \(R B\) can share the same storage locations.
IPERM — Vector of length N . (Input/Output)
On input, IPERM should be initialized to the values \(1,2, \ldots\), N. On output, IPERM contains a record of permutations made on the vector RA.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\) IPERM, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVRGP (RA, RB, IPERM [,...])
Specific: The specific interface names are S_SVRGP and D_SVRGP.

\section*{FORTRAN 77 Interface}

Single: CALL SVRGP (N, RA, RB, IPERM)
Double: The double precision name is DSVRGP.

\section*{Description}

Routine SVRGP sorts the elements of an array, \(A\), into ascending order by algebraic value, keeping a record in \(P\) of the permutations to the array \(A\). That is, the elements of \(P\) are moved in the same manner as are the elements in \(A\) as \(A\) is being sorted. The routine SVRGP uses the algorithm discussed in SvRgn. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Comments}

For wider applicability, integers \((1,2, \ldots, N)\) that are to be associated with RA(I) for \(I=1,2, \ldots, N\) may be entered into \(\operatorname{IPERM}(\mathrm{I})\) in any order. Note that these integers must be unique.

\section*{Example}

This example sorts the 10 -element array RA algebraically.
```

USE SVRGP_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N, NOUT, J
PARAMETER (N=10)
REAL RA(N), RB(N)
INTEGER IPERM(N)
RA =( (10.0 -9.0 8.0 -7.0 6.0 5.0 4.0 -3.0 -2.0 -1.0)
IPERM =($$
\begin{array}{lllllllllll}{1}&{2}&{3}&{4}&{5}&{6}&{7}&{8}&{9}&{10}\end{array}
$$)
DATA RA/10.0, -9.0, 8.0, -7.0, 6.0, 5.0, 4.0, -3.0, -2.0, -1.0/
DATA IPERM/1, 2, 3, 4, 5, 6, 7, 8, 9, 10/
Sort RA by algebraic value into RB
CALL SVRGP (RA, RB, IPERM)
Print results
CALL UMACH (2,NOUT)
WRITE (NOUT, 99998) (RB(J), J=1,N)
WRITE (NOUT, 99999) (IPERM(J),J=1,N)
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ T h e ~ o u t p u t ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , F 5 . 1 ) ) ~
99999 FORMAT (' The permutation vector is:', /, 10(1X,I5))
END

```
!

\section*{Output}
```

The output vector is:
-9.0 -7.0 -3.0 -2.0 -1.0 4.0 5.0 6.0 8.0 10.0
The permutation vector is:
2

```

\section*{SVIGN}

Sorts an integer array by algebraically increasing value.

\section*{Required Arguments}

IA - Integer vector of length N containing the array to be sorted. (Input)

IB - Integer vector of length N containing the sorted array. (Output)
If IA is not needed, IA and IB can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{IA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVIGN (IA, IB [, ...])

Specific: The specific interface name is S_SVIGN .

\section*{FORTRAN 77 Interface}

Single: CALL SVIGN ( \(\mathrm{N}, \mathrm{IA}, \mathrm{IB}\) )

\section*{Description}

Routine SVIGN sorts the elements of an integer array, \(A\), into ascending order by algebraic value.
The routine SVIGN uses the algorithm discussed in SVRGN. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Example}

This example sorts the 10-element array IA algebraically.
```

USE SVIGN INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N, NOUT, J
PARAMETER (N=10)
INTEGER IA(N), IB(N)
IA =($$
\begin{array}{llllllllllll}{-1}&{2}&{-3}&{4}&{-5}&{6}&{-7}&{8}&{-9}&{10}\end{array}
$$)
DATA IA/-1, 2, -3, 4, -5, 6, -7, 8, -9, 10/
CALL SVIGN (IA, IB)

```
\(!\)
```

! Print results
CALL UMACH (2,NOUT)
WRITE (NOUT, 99999) (IB(J),J=1,N)
!
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ o u t p u t ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , I 5 ) ) ,
END

```

\section*{Output}
```

| -9 | -7 | -5 | -3 | -1 | 2 | 4 | 6 | 8 | 10 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

```

\section*{SVIGP}

Sorts an integer array by algebraically increasing value and return the permutation that rearranges the array.

\section*{Required Arguments}

IA - Integer vector of length N containing the array to be sorted. (Input)

IB - Integer vector of length N containing the sorted array. (Output)
If IA is not needed, IA and IB can share the same storage locations.
IPERM - Vector of length N. (Input/Output)
On input, IPERM should be initialized to the values \(1,2, \ldots\), \(N\). On output, IPERM contains a record of permutations made on the vector IA.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}\) (IPERM,1).

\section*{FORTRAN 90 Interface}

Generic: CALL SVIGP (IA, IB, IPERM [, ...])
Specific: The specific interface name is S_SVIGP.

\section*{FORTRAN 77 Interface}

Single: CALL SVIGP (N, IA, IB, IPERM)

\section*{Description}

Routine SVIGP sorts the elements of an integer array, \(A\), into ascending order by algebraic value, keeping a record in \(P\) of the permutations to the array \(A\). That is, the elements of \(P\) are moved in
the same manner as are the elements in \(A\) as \(A\) is being sorted. The routine SVIGP uses the algorithm discussed in SVRGN. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Comments}

For wider applicability, integers \((1,2, \ldots, N)\) that are to be associated with \(\operatorname{IA}(I)\) for \(I=1,2, \ldots, N\) may be entered into IPERM(I) in any order. Note that these integers must be unique.

\section*{Example}

This example sorts the 10 -element array IA algebraically.
```

USE SVIGP INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER N, J, NOUT
PARAMETER (N=10)
INTEGER IA(N), IB(N), IPERM(N)
Set values for IA and IPERM
IA =( ($$
\begin{array}{llllllllllll}{10}&{-9}&{8}&{-7}&{6}&{5}&{4}&{-3}&{-2}&{-1}\end{array}
$$)
IPERM = ( 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 )
DATA IA/10, -9, 8, -7, 6, 5, 4, -3, -2, -1/
DATA IPERM/1, 2, 3, 4, 5, 6, 7, 8, 9, 10/
Sort IA by algebraic value into IB
CALL SVIGP (IA, IB, IPERM)
Print results
CALL UMACH (2,NOUT)
WRITE (NOUT, 99998) (IB(J),J=1,N)
WRITE (NOUT, 99999) (IPERM(J),J=1,N)
99998 FORMAT (' The output vector is:', /, 10(1X,I5))
99999 FORMAT (' The permutation vector is:', /, 10(1X,I5))
END

```
!

\section*{Output}
\(\begin{array}{llllllllll}-9 & -7 & -3 & -2 & -1 & 4 & 5 & 6 & 8 & 10\end{array}\)
The permutation vector is:
\begin{tabular}{llllllllll}
2 & 4 & 8 & 9 & 10 & 7 & 6 & 5 & 3 & 1
\end{tabular}

\section*{SVRBN}

Sorts a real array by nondecreasing absolute value.

\section*{Required Arguments}
\(\boldsymbol{R A}\) - Vector of length N containing the array to be sorted. (Input)
\(\boldsymbol{R B}\) - Vector of length n containing the sorted array. (Output)
If RA is not needed, \(R A\) and \(R B\) can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{RA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVRBN (RA, RB [, ...])

Specific: The specific interface names are S_SVRBN and D_SVRBN.

\section*{FORTRAN 77 Interface}

Single: CALL SVRBN (N, RA, RB)
Double: \(\quad\) The double precision name is DSVRBN.

\section*{Description}

Routine SVRBN sorts the elements of an array, \(A\), into ascending order by absolute value. The routine SVRBN uses the algorithm discussed in SVRGN. On completion, \(\left|A_{j}\right| \leq\left|A_{i}\right|\) for \(j<i\).

\section*{Example}

This example sorts the 10-element array RA by absolute value.
```

USE SVRBN_INT
USE UMACH_}\mp@subsup{}{-}{-
IMPLICIT NONE
INTEGER N, J, NOUT
PARAMETER (N=10)
REAL RA(N), RB(N)
RA =( -1.0 3.0 -4.0 2.0 -1.0
DATA RA/-1.0, 3.0, -4.0, 2.0, -1.0, 0.0, -7.0, 6.0, 10.0, -7.0/
CALL SVRBN (RA, RB)
CALL UMACH (2,NOUT)
WRITE (NOUT, 99999) (RB(J),J=1,N)

```
!
99999 FORMAT (' The output vector is :', /, \(10(1 \mathrm{X}, \mathrm{F} 5.1)\) )
    END

\section*{Output}

The Output vector is :
\(0.0-1.0 \quad-1.0 \quad 2.0 \quad 3.0 \quad-4.0 \quad 6.0 \quad-7.0 \quad-7.0 \quad 10.0\)

\section*{SVRBP}

Sorts a real array by nondecreasing absolute value and return the permutation that rearranges the array.

\section*{Required Arguments}
\(\boldsymbol{R A}-\) Vector of length N containing the array to be sorted. (Input)
\(\boldsymbol{R B}\) - Vector of length N containing the sorted array. (Output)
If RA is not needed, RA and RB can share the same storage locations.
IPERM - Vector of length N. (Input/Output)
On input, IPERM should be initialized to the values \(1,2, \ldots\), \(N\). On output, IPERM contains a record of permutations made on the vector IA.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: N = SIZE (IPERM,1).

\section*{FORTRAN 90 Interface}

Generic: CALL SVRBP (RA, RB, IPERM [,...])

Specific: The specific interface names are S_SVRBP and D_SVRBP.

\section*{FORTRAN 77 Interface}

Single: CALL SVRBP (N, RA, RB, IPERM)
Double: The double precision name is DSVRBP.

\section*{Description}

Routine SVRBP sorts the elements of an array, \(A\), into ascending order by absolute value, keeping a record in \(P\) of the permutations to the array \(A\). That is, the elements of \(P\) are moved in the same
manner as are the elements in \(A\) as \(A\) is being sorted. The routine SVRBP uses the algorithm discussed in SVRGN. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Comments}

For wider applicability, integers \((1,2, \ldots, N)\) that are to be associated with \(\operatorname{RA}(I)\) for \(I=1,2, \ldots, N\) may be entered into IPERM(I) in any order. Note that these integers must be unique.

\section*{Example}

This example sorts the 10 -element array RA by absolute value.
```

    USE SVRBP INT
    USE UMACH_INT
    IMPLICIT NONE
!
INTEGER N, J, NOUT, I
PARAMETER (N=10)
REAL RA(N), RB(N)
INTEGER IPERM(N)
RA =( (10.0 9.0 8.0 7.0 6.0 5.0 -4.0 3.0 -2.0 1.0 )
IPERM = ($$
\begin{array}{lllllllllll}{1}&{2}&{3}&{4}&{5}&{6}&{7}&{8}&{9}&{10}\end{array}
$$)
DATA RA/10.0, 9.0, 8.0, 7.0, 6.0, 5.0, -4.0, 3.0, -2.0, 1.0/
DATA IPERM/1, 2, 3, 4, 5, 6, 7, 8, 9, 10/
Sort RA by absolute value into RB
CALL SVRBP (RA, RB, IPERM)
Print results
CALL UMACH (2,NOUT)
WRITE (NOUT, 99998) (RB(J),J=1,N)
WRITE (NOUT, 99999) (IPERM(I), I=1,N)
!
9 9 9 9 8 ~ F O R M A T ~ ( ' ~ T h e ~ o u t p u t ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , F 5 . 1 ) )
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ p e r m u t a t i o n ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , I 5 ) )
END

```

\section*{Output}

The output vector is:
\(1.0-2.0 \quad 3.0 \quad-4.0 \quad 5.0 \quad 6.0 \quad 7.0 \quad 8.0 \quad 9.0 \quad 10.0\)
The permutation vector is:
\(\begin{array}{llllllllll}10 & 9 & 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1\end{array}\)

\section*{SVIBN}

Sorts an integer array by nondecreasing absolute value.

\section*{Required Arguments}
\(\boldsymbol{I A}\) - Integer vector of length N containing the array to be sorted. (Input)
IB - Integer vector of length N containing the sorted array. (Output) If IA is not needed, IA and IB can share the same storage locations.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input) Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{IA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVIBN (IA, IB \([, \ldots]\) )

Specific: The specific interface name is S_SVIBN.

\section*{FORTRAN 77 Interface}

Single: CALL SVIBN (N, IA, IB)

\section*{Description}

Routine SVIBN sorts the elements of an integer array, \(A\), into ascending order by absolute value. This routine SVIBN uses the algorithm discussed in SVRGN. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Example}

This example sorts the 10 -element array IA by absolute value.
```

USE SVIBN INT
USE UMACH_INT
IMPLICIT NONE
! Declare variables
INTEGER I, J, NOUT, N
PARAMETER (N=10)
INTEGER IA(N), IB(N)
IA =( (-1 [10
DATA IA/-1, 3, -4, 2, -1, 0, -7, 6, 10, -7/
CALL SVIBN (IA, IB)
CALL UMACH (2,NOUT)
WRITE (NOUT, 99999) (IB(J),J=1,N)
99999 FORMAT (' The output vector is:', /, 10(1X,I5))

```
!

Output
The Output vector is:
\(\begin{array}{llllllllll}0 & -1 & -1 & 2 & 3 & -4 & 6 & -7 & -7 & 10\end{array}\)

\section*{SVIBP}

Sorts an integer array by nondecreasing absolute value and return the permutation that rearranges the array.

\section*{Required Arguments}

IA - Integer vector of length N containing the array to be sorted. (Input)

IB - Integer vector of length N containing the sorted array. (Output)
If IA is not needed, IA and IB can share the same storage locations.
IPERM - Vector of length N. (Input/Output)
On input, IPERM should be initialized to the values \(1,2, \ldots\), N. On output, IPERM
contains a record of permutations made on the vector IA.

\section*{Optional Arguments}
\(N\) - Number of elements in the array to be sorted. (Input)
Default: \(\mathrm{N}=\operatorname{SIZE}(\mathrm{IA}, 1)\).

\section*{FORTRAN 90 Interface}

Generic: CALL SVIBP (IA, IB, IPERM [,...])

Specific: The specific interface name is S_SVIBP.

\section*{FORTRAN 77 Interface}

Single: CALL SVIBP (N, IA, IB, IPERM)

\section*{Description}

Routine SVIBP sorts the elements of an integer array, \(A\), into ascending order by absolute value, keeping a record in \(P\) of the permutations to the array \(A\). That is, the elements of \(P\) are moved in the same manner as are the elements in \(A\) as \(A\) is being sorted. The routine SVIBP uses the algorithm discussed in SVRGN. On completion, \(A_{j} \leq A_{i}\) for \(j<i\).

\section*{Comments}

For wider applicability, integers \((1,2, \ldots, \mathrm{~N})\) that are to be associated with \(\operatorname{IA}(\mathrm{I})\) for \(\mathrm{I}=1,2, \ldots, \mathrm{~N}\) may be entered into \(\operatorname{IPERM}(\mathrm{I})\) in any order. Note that these integers must be unique.

\section*{Example}

This example sorts the 10 -element array IA by absolute value.
```

    USE SVIBP_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER N, U, NOUT, J
    PARAMETER (N=10)
    INTEGER IA (N), IB(N), IPERM(N)
        Set values for IA
    IA =(\begin{array}{llllllllllllll}{10}&{9}&{8}&{7}&{6}&{5}&{-4}&{3}&{-2}&{1}\end{array})
    IPERM =( ($$
\begin{array}{llllllllllll}{1}&{2}&{3}&{4}&{5}&{6}&{7}&{8}&{9}&{10}\end{array}
$$)
DATA IA/10, 9, 8, 7, 6, 5, -4, 3, -2, 1/
DATA IPERM/1, 2, 3, 4, 5, 6, 7, 8, 9, 10/
CALL SVIBP (IA, IB, IPERM)
CALL UMACH (2,NOUT)
WRITE (NOUT, 99998) (IB(J),J=1,N)
WRITE (NOUT, 99999) (IPERM(J),J=1,N)
!
99998 FORMAT (' The output vector is:', /, 10(1X,I5))
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ T h e ~ p e r m u t a t i o n ~ v e c t o r ~ i s : ' , ~ / , ~ 1 0 ( 1 X , I 5 ) )
END

```

\section*{Output}

The Output vector is:
\begin{tabular}{llllllllll}
1 & -2 & 3 & -4 & 5 & 6 & 7 & 8 & 9 & 10
\end{tabular}

The permutation vector is:
\begin{tabular}{llllllllll}
10 & 9 & 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1
\end{tabular}

\section*{SRCH}

Searches a sorted vector for a given scalar and return its index.

\section*{Required Arguments}

VALUE - Scalar to be searched for in Y. (Input)
\(\boldsymbol{X}\) - Vector of length N * INCX. (Input)
\(Y\) is obtained from \(X\) for \(I=1,2, \ldots, N\) by \(Y(I)=X(1+(I-1) * I N C X) . Y(1), Y(2), \ldots\), \(\mathrm{Y}(\mathrm{N})\) must be in ascending order.

INDEX - Index of Y pointing to VALUE. (Output)
If INDEX is positive, VALUE is found in Y. If INDEX is negative, VALUE is not found in Y.

INDEX
Location of VALUE

1 thru N
VALUE \(=Y(\) INDEX \()\)
\(-1 \quad\) VALUE \(<\mathrm{Y}(1)\) or \(\mathrm{N}=0\)
-N thru \(-2 \quad \mathrm{Y}(-\operatorname{INDEX}-1)<\) VALUE \(<\mathrm{Y}(\) INDEX \()\)
\(-(\mathrm{N}+1) \quad\) VALUE \(>\mathrm{Y}(\mathrm{N})\)

\section*{Optional Arguments}
\(N\) - Length of vector Y. (Input)
Default: \(\mathrm{N}=(\operatorname{SIZE}(\mathrm{X}, 1)) / \operatorname{INCX}\).
\(\boldsymbol{I N C X}\) - Displacement between elements of X . (Input)
INCX must be greater than zero.
Default: \(\operatorname{INCX}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL SRCH (VALUE, X, INDEX [, ...])
Specific: The specific interface names are S_SRCH and D_SRCH.

\section*{FORTRAN 77 Interface}

Single: CALL SRCH (N, VALUE, X, INCX, INDEX)
Double: The double precision name is DSRCH.

\section*{Description}

Routine SRCH searches a real vector \(x\) (stored in X ), whose \(n\) elements are sorted in ascending order for a real number \(c\) (stored in VALUE). If \(c\) is found in \(x\), its index \(i\) (stored in INDEX) is returned so that \(x_{i}=c\). Otherwise, a negative number \(i\) is returned for the index. Specifically,
\[
\begin{array}{ll}
\text { if } 1 \leq i \leq n & \text { then } x_{i}=c \\
\text { if } i=-1 & \text { then } c<x_{1} \text { or } n=0 \\
\text { if }-n \leq I \leq-2 & \text { then } x_{-i-1}<c<x_{-i} \\
\text { if } i=-(n+1) & \text { then } c>x_{n}
\end{array}
\]

The argument INCX is useful if a row of a matrix, for example, row number I of a matrix \(X\), must be searched. The elements of row I are assumed to be in ascending order. In this case, set INCX equal to the leading dimension of x exactly as specified in the dimension statement in the calling program. With X declared
```

REAL X(LDX,N)

```
the invocation
```

CALL SRCH (N, VALUE, X(I, 1), LDX, INDEX)

```
returns an index that will reference a column number of x .
Routine SRCH performs a binary search. The routine is an implementation of algorithm \(B\) discussed by Knuth (1973, pages 407-411).

\section*{Example}

This example searches a real vector sorted in ascending order for the value 653.0. The problem is discussed by Knuth (1973, pages 407-409).
```

USE SRCH INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=16)
INTEGER INDEX, NOUT
REAL VALUE, X(N)
DATA X/61.0, 87.0, 154.0, 170.0, 275.0, 426.0, 503.0, 509.0, \&
512.0, 612.0, 653.0, 677.0, 703.0, 765.0, 897.0, 908.0/
VALUE = 653.0
CALL SRCH (VALUE, X, INDEX)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'INDEX = ', INDEX
END

```
\(!\)
!
\(!\)
\(!\)

\section*{Output}

INDEX \(=11\)

\section*{ISRCH}

Searches a sorted integer vector for a given integer and return its index.

\section*{Required Arguments}

IVALUE - Scalar to be searched for in IY. (Input)
\(\boldsymbol{I} \boldsymbol{X}\) - Vector of length N * INCX. (Input)
IY is obtained from \(I X\) for \(I=1,2, \ldots, N\) by
\(\operatorname{IY}(\mathrm{I})=\operatorname{IX}(1+(\mathrm{I}-1) * \operatorname{INCX}) . \operatorname{IY}(1), \operatorname{IY}(2), \ldots, I Y(\mathrm{~N})\) must be in ascending order.
INDEX - Index of IY pointing to IVALUE. (Output)
If INDEX is positive, IVALUE is found in IY. If INDEX is negative, IVALUE is not found in IY.

INDEX
Location of VALUE

1 thru N IVALUE \(=\operatorname{IY}(\) INDEX \()\)
\(-1\)
IVALUE \(<\operatorname{IY}(1)\) or \(\mathrm{N}=0\)
-N thru -2
\(\operatorname{IY}(-I N D E X-1)<\operatorname{IVALUE}<\operatorname{IY}(-I N D E X)\)
\(-(\mathrm{N}+1)\)
IVALUE \(>\mathrm{Y}(\mathrm{N})\)

\section*{Optional Arguments}
\(N\) - Length of vector IY. (Input)
Default: \(N=\) SIZE (IX,1) / INCX.

INCX - Displacement between elements of IX. (Input)
INCX must be greater than zero.
Default: \(\operatorname{INCX}=1\).

\section*{FORTRAN 90 Interface}

Generic: CALL ISRCH (IVALUE, IX, INDEX [,...])
Specific: The specific interface name is \(S_{-}\)ISRCH.

\section*{FORTRAN 77 Interface}

Single: CALL ISRCH (N, IVALUE, IX, INCX, INDEX)

\section*{Description}

Routine ISRCH searches an integer vector \(x\) (stored in IX), whose \(n\) elements are sorted in ascending order for an integer \(c\) (stored in IVALUE). If \(c\) is found in \(x\), its index \(i\) (stored in INDEX) is returned so that \(x_{i}=c\). Otherwise, a negative number \(i\) is returned for the index. Specifically,
```

if $1 \leq i \leq n \quad$ Then $x_{i}=c$
if $i=-1 \quad$ Then $c<x_{1}$ or $n=0$
if $-n \leq i \leq-2 \quad$ Then $x_{-i-1}<c<x_{-i}$
if $i=-(n+1) \quad$ Then $c>x_{n}$

```

The argument INCX is useful if a row of a matrix, for example, row number I of a matrix IX, must be searched. The elements of row I are assumed to be in ascending order. Here, set INCX equal to the leading dimension of IX exactly as specified in the dimension statement in the calling program. With IX declared

InTEGER IX(LDIX,N)
the invocation
```

CALL ISRCH (N,IVALUE,IX(I,1),LDIX,INDEX)

```
returns an index that will reference a column number of IX.
The routine ISRCH performs a binary search. The routine is an implementation of algorithm \(B\) discussed by Knuth (1973, pages 407-411).

\section*{Example}

This example searches an integer vector sorted in ascending order for the value 653 . The problem is discussed by Knuth (1973, pages 407-409).
```

USE ISRCH INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=16)
INTEGER INDEX, NOUT
INTEGER IVALUE, IX(N)
DATA IX/61, 87, 154, 170, 275, 426, 503, 509, 512, 612, 653, 677, \&
703, 765, 897, 908/
IVALUE = 653
CALL ISRCH (IVALUE, IX, INDEX)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'INDEX = ', INDEX
END

```

\section*{Output}

INDEX \(=11\)

\section*{SSRCH}

Searches a character vector, sorted in ascending ASCII order, for a given string and return its index.

\section*{Required Arguments}
\(N\) - Length of vector CHY. (Input)
Default: N = SIZE (CHX,1) / INCX.

STRING - Character string to be searched for in CHY. (Input)
\(\boldsymbol{C H X}\) - Vector of length N * INCX containing character strings. (Input)
CHY is obtained from CHX for \(I=1,2, \ldots, N\) by \(\operatorname{CHY}(I)=\operatorname{CHX}(1+(I-1) * I N C X)\). CHY(1), CHY(2), ..., CHY(N) must be in ascending ASCII order.

INCX - Displacement between elements of CHX. (Input)
INCX must be greater than zero.
Default: \(\operatorname{INCX}=1\).

INDEX - Index of CHY pointing to STRING. (Output)
If INDEX is positive, STRING is found in CHY. If INDEX is negative, STRING is not found in CHY.

INDEX

1 thru N
\(-1\)
-N thru -2
\(-(\mathrm{N}+1) \quad\) STRING \(>\mathrm{CHY}(\mathrm{N})\)

\section*{FORTRAN 90 Interface}

Generic: CALL SSRCH (N, STRING, CHX, INCX, INDEX)

Specific: The specific interface name is SSRCH.

\section*{FORTRAN 77 Interface}

Single: CALL SSRCH (N, STRING, CHX, INCX, INDEX)

\section*{Description}

Routine SSRCH searches a vectorof character strings \(x\) (stored in CHX), whose \(n\) elements are sorted in ascending ASCII order, for a character string \(c\) (stored in STRING). If \(c\) is found in \(x\), its index \(i\) (stored in INDEX) is returned so that \(x_{i}=c\). Otherwise, a negative number \(i\) is returned for the index. Specifically,
\[
\begin{array}{ll}
\text { if } 1 \leq i \leq n & \text { Then } x_{i}=c \\
\text { if } i=-1 & \text { Then } c<x_{1} \text { or } n=0 \\
\text { if }-n \leq i \leq-2 & \text { Then } x_{-i-1}<c<x_{-i} \\
\text { if } i=-(n+1) & \text { Then } c>x_{n}
\end{array}
\]

Here, " \(<\) " and " \(>\) " are in reference to the ASCII collating sequence. For comparisons made between character strings \(c\) and \(x_{i}\) with different lengths, the shorter string is considered as if it were extended on the right with blanks to the length of the longer string. (SSRCH uses FORTRAN intrinsic functions LLT and LGT.)

The argument INCX is useful if a row of a matrix, for example, row number I of a matrix CHX, must be searched. The elements of row I are assumed to be in ascending ASCII order. In this case, set INCX equal to the leading dimension of CHX exactly as specified in the dimension statement in the calling program. With CHX declared
```

CHARACTER * }7\mathrm{ CHX(LDCHX,N)

```
the invocation
CALL SSRCH (N, STRING, CHX (I, 1), LDCHX, INDEX)
returns an index that will reference a column number of CHX .
The routine SSRCH performs a binary search. The routine is an implementation of algorithm \(B\) discussed by Knuth (1973, pages 407-411).

\section*{Example}

This example searches a CHARACTER * 2 vector containing 9 character strings, sorted in ascending ASCII order, for the value ' \(\mathrm{CC}^{\prime}\).
```

USE SSRCH_INT
USE UMACH INT
IMPLICIT NONE
INTEGER N, INCX
PARAMETER (N=9)

```
INTEGER INDEX, NOUT
CHARACTER CHX(N)*2, STRING*2
DATA CHX/'AA', 'BB', 'CC', 'DD', 'EE', 'FF', 'GG', 'HH', \&
'II'/
\(\operatorname{INCX}=1\)
```

STRING = 'CC'
CALL SSRCH (N, STRING, CHX, INCX, INDEX)
!
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'INDEX = ', INDEX
END

```

\section*{Output}

INDEX = 3

\section*{ACHAR}

This function returns a character given its ASCII value.

\section*{Function Return Value}

ACHAR - CHARACTER * 1 string containing the character in the I-th position of the ASCII collating sequence. (Output)

\section*{Required Arguments}

I- Integer ASCII value of the character desired. (Input) I must be greater than or equal to zero and less than or equal to 127 .

\section*{FORTRAN 90 Interface}

Generic: ACHAR (I)
Specific: The specific interface name is ACHAR.

\section*{FORTRAN 77 Interface}

Single: ACHAR (I)

\section*{Description}

Routine ACHAR returns the character of the input ASCII value. The input value should be between 0 and 127. If the input value is out of range, the value returned in ACHAR is machine dependent.

\section*{Example}

This example returns the character of the ASCII value 65 .
```

USE ACHAR INT
USE UMACH_INT
! IMPLICIT NONE
INTEGER I, NOUT

```
```

!
CALL UMACH (2, NOUT)
! Get character for ASCII value
I = 65
WRITE (NOUT, 99999) I, ACHAR(I)
!
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ F o r ~ t h e ~ A S C I I ~ v a l u e ~ o f ~ ' , ~ I 2 , ~ ' , ~ t h e ~ c h a r a c t e r ~ i s ~ : ~ ' , ~ \& ~
A1)
END

```

\section*{Output}

For the ASCII value of 65, the character is : A

\section*{IACHAR}

This function returns the integer ASCII value of a character argument.

\section*{Function Return Value}

IACHAR - Integer ASCII value for CH. (Output) The character CH is in the IACHAR-th position of the ASCII collating sequence.

\section*{Required Arguments}
\(\boldsymbol{C H}\) - Character argument for which the integer ASCII value is desired. (Input)

\section*{FORTRAN 90 Interface}

Generic: IACHAR (CH)

Specific: The specific interface name is IACHAR.

\section*{FORTRAN 77 Interface}

\section*{Description}

Routine IACHAR returns the ASCII value of the input character.
Single: IACHAR (CH)

\section*{Example}

This example gives the ASCII value of character A.
USE IACHAR_INT
IMPLICIT NONE
```

    INTEGER NOUT
    CHARACTER CH
    !
CALL UMACH (2, NOUT)
! Get ASCII value for the character
CH = 'A'
WRITE (NOUT,99999) CH, IACHAR(CH)
!
99999 FORMAT (' For the character ', A1, ' the ASCII value is : ', \&
I3)
END

```

\section*{Output}
```

For the character A the ASCII value is : }6

```

\section*{ICASE}

This function returns the ASCII value of a character converted to uppercase.

\section*{Function Return Value}

ICASE - Integer ASCII value for CH without regard to the case of ch. (Output)
Routine ICASE returns the same value as IACHAR for all but lowercase letters. For these, it returns the \(I A C H A R\) value for the corresponding uppercase letter.

\section*{Required Arguments}
\(\mathbf{C H}\) - Character to be converted. (Input)

\section*{FORTRAN 90 Interface}

Generic: ICASE (CH)
Specific: The specific interface name is ICASE.

\section*{FORTRAN 77 Interface}

Single: ICASE (CH)

\section*{Description}

Routine ICASE converts a character to its integer ASCII value. The conversion is case insensitive; that is, it returns the ASCII value of the corresponding uppercase letter for a lowercase letter.

\section*{Example}

This example shows the case insensitive conversion.
```

    USE ICASE_INT
    USE UMACH_INT
    IMPLICIT NONE
    INTEGER NOUT
    CHARACTER CHR
    ! Get output unit number
CALL UMACH (2, NOUT)
! Get ASCII value for the character
CHR = 'a'
WRITE (NOUT,99999) CHR, ICASE (CHR)
!
9 9 9 9 9 ~ F O R M A T ~ ( ' ~ F o r ~ t h e ~ c h a r a c t e r ~ ' , ~ A 1 , ~ ' ~ t h e ~ I C A S E ~ v a l u e ~ i s ~ : ~ ' , ~ \& ~
I3)
END

```

\section*{Output}
```

For the character a the ICASE value is : 65

```

\section*{IICSR}

This function compares two character strings using the ASCII collating sequence but without regard to case.

\section*{Function Return Value}

IICSR - Comparison indicator. (Output)
Let USTR1 and USTR2 be the uppercase versions of STR1 and STR2, respectively. The following table indicates the relationship between USTR1 and USTR2 as determined by the ASCII collating sequence.

IICSR Meaning
-1 USTR1 precedes USTR2
0 USTR1 equals USTR2
1 USTR1 follows USTR2

\section*{Required Arguments}

STR1 - First character string. (Input)
STR2 - Second character string. (Input)

\section*{FORTRAN 90 Interface}

Generic: IICSR (STR1, STR2)
Specific: The specific interface name is IICSR.

\section*{FORTRAN 77 Interface}
```

Single: IICSR (STR1, STR2)

```

\section*{Description}

Routine IICSR compares two character strings. It returns -1 if the first string is less than the second string, 0 if they are equal, and 1 if the first string is greater than the second string. The comparison is case insensitive.

\section*{Comments}

If the two strings, STR1 and STR2, are of unequal length, the shorter string is considered as if it were extended with blanks to the length of the longer string.

\section*{Example}

This example shows different cases on comparing two strings.
```

USE IICSR_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
CHARACTER STR1*6, STR2*6
CALL UMACH (2, NOUT)
STR1 = 'ABC 1'
STR2 = ' '
WRITE (NOUT,99999) STR1, STR2, IICSR(STR1,STR2)
STR1 = 'AbC'
STR2 = 'ABC'
WRITE (NOUT,99999) STR1, STR2, IICSR(STR1,STR2)
String1 is 'smaller' than String2
STR1 = 'ABC'
STR2 = 'aBC 1'
WRITE (NOUT,99999) STR1, STR2, IICSR(STR1,STR2)
99999 FORMAT (' For String1 = ', A6, 'and String2 = ', A6, \&
' IICSR = ', I2, /)
END

```
\(!\)
!
!

\section*{Output}
```

For String1 = ABC 1 and String2 = IICSR = 1
For String1 = AbC and String2 = ABC IICSR = 0
For String1 = ABC and String2 = aBC 1 IICSR = -1

```

\section*{IIDEX}

This funcion determines the position in a string at which a given character sequence begins without regard to case.

\section*{Function Return Value}

IIDEX - Position in ChRSTR where KEY begins. (Output)
If KEY occurs more than once in CHRSTR, the starting position of the first occurrence is returned. If KEY does not occur in CHRSTR, then IIDEX returns a zero.

\section*{Required Arguments}

CHRSTR - Character string to be searched. (Input)
\(\boldsymbol{K E Y}\) - Character string that contains the key sequence. (Input)

\section*{FORTRAN 90 Interface}

Generic: IIDEX (CHRSTR, KEY)

Specific: The specific interface name is IIDEX.

\section*{FORTRAN 77 Interface}

Single: IIDEX (CHRSTR, KEY)

\section*{Description}

Routine IIDEX searches for a key string in a given string and returns the index of the starting element at which the key character string begins. It returns 0 if there is no match. The comparison is case insensitive. For a case-sensitive version, use the FORTRAN 77 intrinsic function INDEX.

\section*{Comments}

If the length of KEY is greater than the length CHRSTR, IIDEX returns a zero.

\section*{Example}

This example locates a key string.
```

        USE IIDEX_INT
    USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
CHARACTER KEY*5, STRING*10
! Get output unit number
CALL UMACH (2, NOUT)
STRING = 'a1b2c3d4e5'
KEY = 'C3d4E'
WRITE (NOUT,99999) STRING, KEY, IIDEX(STRING,KEY)
!
KEY = 'F'
WRITE (NOUT,99999) STRING, KEY, IIDEX(STRING,KEY)
!
99999 FORMAT (' For STRING = ', A10, ' and KEY = ', A5, ' IIDEX = ', I2, \&
/)
END

```

\section*{Output}
For STRING \(=\) a1b2c3d4e5 and \(K E Y=C 3 d 4 E\) IIDEX \(=5\)
For STRING \(=\) a1b2c3d4e5 and KEY \(=\mathrm{F} \quad\) IIDEX \(=0\)

\section*{CVTSI}

Converts a character string containing an integer number into the corresponding integer form.

\section*{Required Arguments}

STRING - Character string containing an integer number. (Input)
NUMBER - The integer equivalent of STRING. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL CVTSI (STRING, NUMBER)
Specific: The specific interface name is CVTSI.

\section*{FORTRAN 77 Interface}

Single: CALL CVTSI (STRING, NUMBER)

\section*{Description}

Routine CVTSI converts a character string containing an integer to an INTEGER variable. Leading and trailing blanks in the string are ignored. If the string contains something other than an integer, a terminal error is issued. If the string contains an integer larger than can be represented by an INTEGER variable as determined from routine IMACH (see the Reference Material), a terminal error is issued.

\section*{Example}

The string " 12345 " is converted to an INTEGER variable.
```

USE CVTSI INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT, NUMBER
CHARACTER STRING*10
DATA STRING/'12345'/
CALL CVTSI (STRING, NUMBER)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'NUMBER = ', NUMBER
END

```
\(!\)
\(!\)
\(!\)
Output
NUMBER \(=12345\)

\section*{CPSEC}

This fuction returns CPU time used in seconds.

\section*{Function Return Value}

CPSEC - CPU time used (in seconds) since first call to CPSEC. (Output)

\section*{Required Arguments}

None

\section*{FORTRAN 90 Interface}

Generic: CPSEC ()

Specific: The specific interface name is CPSEC.

\section*{FORTRAN 77 Interface}

Single: CPSEC (1)

\section*{Comments}
1. The first call to CPSEC returns 0.0 .
2. The accuracy of this routine depends on the hardware and the operating system. On some systems, identical runs can produce timings differing by more than 10 percent.

\section*{TIMDY}

Gets time of day.

\section*{Required Arguments}

IHOUR - Hour of the day. (Output) IHOUR is between 0 and 23 inclusive.

MINUTE - Minute within the hour. (Output)
MINUTE is between 0 and 59 inclusive.
ISEC - Second within the minute. (Output)
ISEC is between 0 and 59 inclusive.

\section*{FORTRAN 90 Interface}

Generic: CALL tIMDY (IHOUR, MINUTE, ISEC)
Specific: \(\quad\) The specific interface name is \(\operatorname{TIMDY}\).

\section*{FORTRAN 77 Interface}

Single: CALL TIMDY (IHOUR, MINUTE, ISEC)

\section*{Description}

Routine TIMDY is used to retrieve the time of day.

\section*{Example}

The following example uses TIMDY to return the current time. Obviously, the output is dependent upon the time at which the program is run.
```

USE TIMDY INT
USE UMACH_INT

```
```

IMPLICIT NONE
INTEGER IHOUR, IMIN, ISEC, NOUT
!
CALL TIMDY (IHOUR, IMIN, ISEC)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'Hour:Minute:Second = ', IHOUR, ':', IMIN, \&
':', ISEC
IF (IHOUR .EQ. O) THEN
WRITE (NOUT,*) 'The time is ', IMIN, ' minute(s), ', ISEC, \&
' second(s) past midnight.'
ELSE IF (IHOUR .LT. 12) THEN
WRITE (NOUT,*) 'The time is ', IMIN, ' minute(s), ', ISEC, \&
' second(s) past ', IHOUR, ' am.'
ELSE IF (IHOUR .EQ. 12) THEN
WRITE (NOUT,*) 'The time is ', IMIN, ' minute(s), ', ISEC, \&
' second(s) past noon.'
ELSE
WRITE (NOUT,*) 'The time is ', IMIN, ' minute(s), ', ISEC, \&
' second(s) past ', IHOUR-12, ' pm.'
END IF
END

```

\section*{Output}

Hour:Minute:Second \(=14\) : 34 : 30
The time is 34 minute(s), 30 second(s) past 2 pm .

\section*{TDATE}

Gets today's date.

\section*{Required Arguments}

IDAY - Day of the month. (Output) IDAY is between 1 and 31 inclusive.

MONTH - Month of the year. (Output) MONTH is between 1 and 12 inclusive.

IYEAR - Year. (Output)
For example, IYEAR \(=1985\).

\section*{FORTRAN 90 Interface}

Generic: CALL TDATE (IDAY, MONTH, IYEAR)
Specific: \(\quad\) The specific interface name is TDATE.

\section*{FORTRAN 77 Interface}

Single: CALL TDATE (IDAY, MONTH, IYEAR)

\section*{Description}

Routine TDATE is used to retrieve today's date. Obviously, the output is dependent upon the date the program is run.

\section*{Example}

The following example uses TDATE to return today's date.
```

USE TDATE INT
USE UMACH_
IMPLICIT NONE
INTEGER IDAY, IYEAR, MONTH, NOUT
CALL TDATE (IDAY, MONTH, IYEAR)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'Day-Month-Year = ', IDAY, '-', MONTH, \&
END

```
!

Output
Day-Month-Year \(=7-7-2006\)

\section*{NDAYS}

This function computes the number of days from January 1, 1900, to the given date.

\section*{Function Return Value}

NDAYS - Function value. (Output)
If NDAYS is negative, it indicates the number of days prior to January 1, 1900.

\section*{Required Arguments}

IDAY - Day of the input date. (Input)
MONTH - Month of the input date. (Input)
IYEAR - Year of the input date. (Input)
1950 would correspond to the year 1950 A.D. and 50 would correspond to year 50 A.D.

\section*{FORTRAN 90 Interface}

Generic: NDAYS (IDAY, MONTH, IYEAR)
Specific: The specific interface name is NDAYS.

\section*{FORTRAN 77 Interface}

Single: NDAYS (IDAY, MONTH, IYEAR)

\section*{Description}

Function NDAYS returns the number of days from January 1, 1900, to the given date. The function NDAYS returns negative values for days prior to January 1, 1900. A negative IYEAR can be used to specify B.C. Input dates in year 0 and for October 5, 1582, through October 14, 1582, inclusive, do not exist; consequently, in these cases, NDAYS issues a terminal error.

\section*{Comments}
1. Informational error
\begin{tabular}{ccl} 
Type & Code & \\
1 & 1 & \begin{tabular}{l} 
The Julian calendar, the first modern calendar, went into use in 45 \\
B.C. No calendar prior to 45 B.C. was as universally used nor as \\
accurate as the Julian. Therefore, it is assumed that the Julian \\
calendar was in use prior to 45 B.C.
\end{tabular}
\end{tabular}
2. The number of days from one date to a second date can be computed by two references to NDAYS and then calculating the difference.
3. The beginning of the Gregorian calendar was the first day after October 4, 1582, which became October 15, 1582. Prior to that, the Julian calendar was in use. NDAYS makes the proper adjustment for the change in calendars.

\section*{Example}

The following example uses NDAYS to compute the number of days from January 15, 1986, to February 28, 1986:
```

USE NDAYS_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER IDAY, IYEAR, MONTH, NDAY0, NDAY1, NOUT
IDAY = 15
MONTH = 1
IYEAR = 1986
NDAYO = NDAYS(IDAY,MONTH,IYEAR)
IDAY = 28
MONTH = 2

```
```

IYEAR = 1986
NDAY1 = NDAYS(IDAY,MONTH,IYEAR)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'Number of days = ', NDAY1 - NDAYO
END

```

\section*{Output}

Number of days \(=44\)

\section*{NDYIN}

Gives the date corresponding to the number of days since January 1, 1900.

\section*{Required Arguments}

NDAYS - Number of days since January 1, 1900. (Input)
IDAY - Day of the input date. (Output)
MONTH - Month of the input date. (Output)
IYEAR - Year of the input date. (Output)
1950 would correspond to the year 195 A.D. and -50 would correspond to year 50 B.C.

\section*{FORTRAN 90 Interface}

Generic: CALL NDYIN (NDAYS, IDAY, MONTH, IYEAR)

Specific: The specific interface name is NDYIN.

\section*{FORTRAN 77 Interface}

Single: CALL NDYIN (NDAYS, IDAY, MONTH, IYEAR)

\section*{Description}

Routine NDYIN computes the date corresponding to the number of days since January 1, 1900. For an input value of NDAYS that is negative, the date computed is prior to January 1, 1900. The routine NDYIN is the inverse of NDAYS.

\section*{Comments}

The beginning of the Gregorian calendar was the first day after October 4, 1582, which became October 15, 1582. Prior to that, the Julian calendar was in use. Routine NDY IN makes the proper adjustment for the change in calendars.

\section*{Example}

The following example uses NDYIN to compute the date for the 100th day of 1986. This is accomplished by first using NDAYS to get the "day number" for December 31, 1985.
```

USE NDYIN_INT
USE NDAYS INT
USE UMACH_INT
IMPLICIT NONE
INTEGER IDAY, IYEAR, MONTH, NDAYO, NOUT, NDAYO
NDAYO = NDAYS (31,12,1985)
CALL NDYIN (NDAY0+100, IDAY, MONTH, IYEAR)
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'Day 100 of 1986 is (day-month-year) ', IDAY, \&
'-', MONTH, '-', IYEAR
END

```
!

\section*{Output}

Day 100 of 1986 is (day-month-year) 10- 4- 1986

\section*{IDYWK}

This function computes the day of the week for a given date.

\section*{Function Return Value}

IDYWK - Function value. (Output)
The value of IDYWK ranges from 1 to 7, where 1 corresponds to Sunday and 7 corresponds to Saturday.

\section*{Required Arguments}

IDAY - Day of the input date. (Input)
MONTH - Month of the input date. (Input)
IYEAR - Year of the input date. (Input)
1950 would correspond to the year 1950 A.D. and 50 would correspond to year 50 A.D.

\section*{FORTRAN 90 Interface}

Generic: IDYWK (IDAY, MONTH, IYEAR)

Specific: The specific interface name is IDYWK.

\section*{FORTRAN 77 Interface}

Single: IDYWK (IDAY, MONTH, IYEAR)

\section*{Description}

Function IDYWK returns an integer code that specifies the day of week for a given date. Sunday corresponds to 1 , Monday corresponds to 2 , and so forth.

A negative IYEAR can be used to specify B.C. Input dates in year 0 and for October 5, 1582, through October 14, 1582, inclusive, do not exist; consequently, in these cases, IDYWK issues a terminal error.

\section*{Comments}
1. Informational error
\begin{tabular}{ccl} 
Type & Code \\
1 & 1 & \begin{tabular}{l} 
The Julian calendar, the first modern calendar, went into use in 45
\end{tabular} \\
\begin{tabular}{l} 
B.C. No calendar prior to 45 B.C. was as universally used nor as \\
accurate as the Julian. Therefore, it is assumed that the Julian \\
calendar was in use prior to 45 B.C.
\end{tabular}
\end{tabular}
2. The beginning of the Gregorian calendar was the first day after October 4, 1582, which became October 15, 1582. Prior to that, the Julian calendar was in use. Function IDYWK makes the proper adjustment for the change in calendars.

\section*{Example}

The following example uses IDYWK to return the day of the week for February 24, 1963.
```

USE IDYWK_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER IDAY, IYEAR, MONTH, NOUT
IDAY = 24
MONTH = 2
IYEAR = 1963
CALL UMACH (2, NOUT)
WRITE (NOUT,*) 'IDYWK (index for day of week) = ', \&
IDYWK(IDAY,MONTH,IYEAR)
END

```
!

\section*{Output}

IDYWK (index for day of week) = 1

\section*{VERML}

This function obtains IMSL MATH/LIBRARY-related version, system and serial numbers.

\section*{Function Return Value}

VERML — CHARACTER string containing information. (Output)

\section*{Required Arguments}

ISELCT - Option for the information to retrieve. (Input)

\section*{ISELCT VERML}

1 IMSL MATH/LIBRARY version number

2 Operating system (and version number) for which the library was produced.
3 Fortran compiler (and version number) for which the library was produced.
4
IMSL MATH/LIBRARY serial number

\section*{FORTRAN 90 Interface}

Generic: VERML (ISELCT)
Specific: The specific interface name is VERML.

\section*{FORTRAN 77 Interface}

Single: VERML (ISELCT)

\section*{Example}

In this example, we print all of the information returned by VERML on a particular machine. The output is omitted because the results are system dependent.
```

USE UMACH INT
USE VERML_INT
IMPLICIT NONE
INTEGER ISELCT, NOUT
CHARACTER STRING(4)*50, TEMP*32
STRING(1) = '('' IMSL MATH/LIBRARY Version Number: '', A)'
STRING(2) = '('' Operating System ID Number: '', A)'
STRING(3) = '('' Fortran Compiler Version Number: '', A)'
STRING(4) = '('' IMSL MATH/LIBRARY Serial Number: '', A)'
Print the versions and numbers.

```
!
```

    CALL UMACH (2, NOUT)
    DO 10 ISELCT=1, 4
        TEMP = VERML(ISELCT)
        WRITE (NOUT,STRING(ISELCT)) TEMP
    10 CONTINUE
END

```

\section*{Output}
```

IMSL MATH/LIBRARY Version Number: IMSL Fortran Numerical Library, Version 6.0.0
Operating System ID Number: Solaris Version 10
Fortran Compiler Version Number: Sun Fortran 95 8.1 2005/01/07 (Workshop 10.0)
IMSL MATH/LIBRARY Serial Number: 999999

```

\section*{RAND_GEN}

Generates a rank-1 array of random numbers. The output array entries are positive and less than 1 in value.

\section*{Required Argument}
\(\boldsymbol{X}\) - Rank-1 array containing the random numbers. (Output)

\section*{Optional Arguments}

IRND = IRND (Output)
Rank-1 integer array. These integers are the internal results of the Generalized Feedback Shift Register (GFSR) algorithm. The values are scaled to yield the floatingpoint array \(x\). The output array entries are between 1 and \(2^{31}-1\) in value.

ISTATE_IN = ISTATE_IN (Input)
Rank-1 integer array of size \(3 p+2\), where \(p=521\), that defines the ensuing state of the GFSR generator. It is used to reset the internal tables to a previously defined state. It is the result of a previous use of the "ISTATE_OUT=" optional argument.

ISTATE_OUT = ISTATE_OUT (Output)
Rank-1 integer array of size \(3 p+2\) that describes the current state of the GFSR generator. It is normally used to later reset the internal tables to the state defined following a return from the GFSR generator. It is the result of a use of the generator without a user initialization, or it is the result of a previous use of the optional argument "ISTATE_IN=" followed by updates to the internal tables from newly generated values. Example 2 illustrates use of ISTATE_IN and ISTATE_OUT for setting and then resetting RAND_GEN so that the sequence of integers, irnd, is repeatable.

IOPT = IOPT (: ) (Input[/Output])
Derived type array with the same precision as the array x ; used for passing optional data to rand_gen. The options are as follows:
\begin{tabular}{|l|l|l|}
\hline \multicolumn{2}{|c|}{ Packaged Options for RAND_GEN } \\
\hline Option Prefix =? & Option Name & Option Value \\
\hline s_, d_ & Rand_gen_generator_seed & 1 \\
\hline s_, d_ & Rand_gen_LCM_modulus & 2 \\
\hline s_, d_ & Rand_gen_use_Fushimi_start & 3 \\
\hline
\end{tabular}
\(\operatorname{IOPT}(I O)=\) ?_options(?_rand_gen_generator_seed, ?_dummy)
Sets the initial values for the GFSR. The present value of the seed, obtained by default from the real-time clock as described below, swaps places with
iopt (IO + 1) \%idummy. If the seed is set before any current usage of RAND_GEN, the exchanged value will be zero.

IOPT(IO) = ?_options(?_rand_gen_LCM_modulus, ?_dummy)
IOPT(IO+1) = ? options(modulus, ? dummy)
Sets the initial values for the GFSR. The present value of the LCM, with default value \(k=16807\), swaps places with iopt (IO +1 ) \%idummy.

IOPT(IO) = ?_options(?_rand_gen_use_Fushimi_start, ?_dummy)
Starts the \(\overline{\text { GFSR }}\) sequence as suggested by Fushimi \(\overline{(1990) . ~ T h e ~ d e f a u l t ~ s t a r t i n g ~}\) sequence is with the LCM recurrence described below.

\section*{FORTRAN 90 Interface}

Generic: CALL RAND_GEN (X [, ...])
Specific: The specific interface names are S_RAND_GEN and D_RAND_GEN.

\section*{Description}

This GFSR algorithm is based on the recurrence
\[
x_{t}=x_{t-3 p} \oplus x_{t-3 p}
\]
where \(a \oplus b\) is the exclusive OR operation on two integers \(a\) and \(b\). This operation is performed until SIZE \((X)\) numbers have been generated. The subscripts in the recurrence formula are computed modulo \(3 p\). These numbers are converted to floating point by effectively multiplying the positive integer quantity
\[
x_{t} \cup 1
\]
by a scale factor slightly smaller than \(1 . /(\operatorname{huge}(1))\). The values \(p=521\) and \(q=32\) yield a sequence with a period approximately
\[
2^{p}>10^{156.8}
\]

The default initial values for the sequence of integers \(\left\{x_{t}\right\}\) are created by a congruential generator starting with an odd integer seed
\[
m=v+\mid \text { count } \cap\left(2^{\text {bit_size(1) }}-1\right) \mid \cup 1
\]
obtained by the Fortran 90 real-time clock routine:
```

CALL SYSTEM_CLOCK(COUNT=count,CLOCK_RATE=CLRATE)

```

An error condition is noted if the value of CLRATE \(=0\). This indicates that the processor does not have a functioning real-time clock. In this exceptional case a starting seed must be provided by the user with the optional argument "iopt=" and option number ?_rand_generator_seed. The value \(v\) is the current clock for this day, in milliseconds. This value is obtained using the date routine:
```

CALL DATE_AND_TIME(VALUES=values)

```
and converting values (5:8) to milliseconds.
The LCM generator initializes the sequence \(\left\{x_{t}\right\}\) using the following recurrence:
\[
m \leftarrow m \times k, \bmod (\text { huge }(1) / 2)
\]

The default value of \(k=16807\). Using the optional argument "iopt=" and the packaged option number ?_rand_gen_LCM_modulus, \(k\) can be given an alternate value. The option number ?_rand_gen_generator_seed can be used to set the initial value of \(m\) instead of using the asynchronous value given by the system clock. This is illustrated in Example 2. If the default choice of \(m\) results in an unsatisfactory starting sequence or it is necessary to duplicate the sequence, then it is recommended that users set the initial seed value to one of their own choosing. Resetting the seed complicates the usage of the routine.
This software is based on Fushimi (1990), who gives a more elaborate starting sequence for the \(\{\mathrm{xt}\}\). The starting sequence suggested by Fushimi can be used with the option number ?_rand_gen_use_Fushimi_start. Fushimi's starting process is more expensive than the default method, and it is equivalent to starting in another place of the sequence with period \(2^{p}\).

\section*{Fatal and Terminal Error Messages}

See the messages.gls file for error messages for RAND_GEN. These error messages are numbered 521-528; 541-548.

\section*{Example 1: Running Mean and Variance}

An array of random numbers is obtained. The sample mean and variance are computed. These values are compared with the same quantities computed using a stable method for the running means and variances, sequentially moving through the data. Details about the running mean and variance are found in Henrici (1982, pp. 21-23).
```

    use rand_gen_int
    implicit none
    ! This is Example 1 for RAND_GEN.
integer i
integer, parameter :: n=1000

```
```

    real(kind(1e0)), parameter :: one=1e0, zero=0e0
    real(kind(1e0)) x(n), mean_1(0:n), mean_2(0:n), s_1(0:n), s_2(0:n)
    ! Obtain random numbers.
call rand_gen(x)
! Calculate each partial mean.
do i=1,n
mean 1(i) = sum(x(1:i))/i
end do
! Calculate each partial variance.
do i=1,n
s_1(i)=sum((x(1:i)-mean_1(i))**2)/i
end do
mean_2(0)=zero
mean-2(1)=x(1)
s_2(0:1)=zero
! Alternately calculate each running mean and variance,
! handling the random numbers once.
do i=2,n
mean_2(i)=((i-1)*mean_2(i-1)+x(i))/i
s_2(\overline{i})=(i-1)*s_2(\overline{i}-1)/i+(mean_2(i)-x(i))**2/(i-1)
end do
! Check that the two sets of means and variances agree.
if (maxval(abs(mean_1(1:)-mean_2(1:))/mean_1(1:)) <= \&
sqrt(epsilon(one))) then
if (maxval(abs(s_1(2:)-s_2(2:))/s_1(2:)) <= \&
sqrt(epsilon(one))) then
write (*,*) 'Example 1 for RAND_GEN is correct.'
end if
end if
end

```

\section*{Output}

Example 1 for RAND GEN is correct.

\section*{Additional Examples}

\section*{Example 2: Seeding, Using, and Restoring the Generator}
```

    use rand_gen_int
    implicit none
    ! This is Example 2 for RAND GEN.
integer i
integer, parameter :: n=34, p=521

```
```

    real(kind(1e0)), parameter :: one=1.0e0, zero=0.0e0
    integer irndi(n), i_out(3*p+2), hidden_message(n)
    real(kind(le0)) x(n), y(n)
    type(s_options) :: iopti(2)=s_options(0,zero)
    character*34 message, returne\overline{d_message}
    ! This is the message to be hidden.
message = 'SAVE YOURSELF. WE ARE DISCOVERED!''
! Start the generator with a known seed.
iopti(1) = s_options(s_rand_gen_generator_seed,zero)
iopti(2) = s_options(1\overline{2}3,ze\overline{ro)}
call rand_ge\overline{n}(x, iopt=iopti)
! Save the state of the generator.
call rand_gen(x, istate_out=i_out)
! Get random integers.
call rand_gen(y, irnd=irndi)
! Hide text using collating sequence subtracted from integers.
do i=1, n
hidden message(i) = irndi(i) - ichar(message(i:i))
end do
! Reset generator to previous state and generate the previous
! random integers.
call rand_gen(x, irnd=irndi, istate_in=i_out)
! Subtract hidden text from integers and convert to character.
do i=1, n
returned_message(i:i) = char(irndi(i) - hidden_message(i))
end do
! Check the results.
if (returned_message == message) then
write (*,*) 'Example 2 for RAND_GEN is correct.'
end if
end

```

\section*{Output}

Example 2 for RAND_GEN is correct.

\section*{Example 3: Generating Strategy with a Histogram}

We generate random integers but with the frequency as in a histogram with \(n_{\text {bins }}\) slots. The generator is initially used a large number of times to demonstrate that it is making choices with the same shape as the histogram. This is not required to generate samples. The program next generates a summary set of integers according to the histogram. These are not repeatable and are
representative of the histogram in the sense of looking at 20 integers during generation of a large number of samples.
use rand_gen_int
use show_int
implicit none
! This is Example 3 for RAND_GEN.
integer i, i_bin, i_map, i_left, i_right
integer, parāmeter \(\overline{:}: \mathrm{n}\) wor \(\mathrm{r}=1000\)
integer, parameter :: n_bins=10
integer, parameter :: scale=1000
integer, parameter :: total_counts=100
integer, parameter :: n_samples=total_counts*scale
integer, dimension(n_bins) :: histogram= \&
(/4, 6, 8, 14, 20 \(, 17,12, ~ 9, ~ 7, ~ 3 /)\)
integer, dimension(n_work) :: working=0
integer, dimension(n_bins) :: distribution=0
integer break_points(0:n_bins)
real(kind(1e0)) rn(n_samples)
real(kind(1e0)), parāmeter : : tolerance=0.005
integer, parameter :: n_samples_20=20
integer rand_num_20(n_samples_20)
real (kind(1ē)) \(\overline{\mathrm{r}} \mathrm{n}_{\mathrm{l}} 20\) (n_samples_20)
! Compute the normalized cumulative distribution.
break_points \((0)=0\)
do i=1,n_bins
break_points(i)=break_points(i-1)+histogram(i)
end do
break_points=break_points*n_work/total_counts
! Obtain uniform random numbers.
call rand_gen(rn)
! Set up the secondary mapping array.
do i_bin=1,n_bins
i_Ieft=breāk_points(i_bin-1)+1
i_right=break_points(i_bin)
do i=i_left, i_right
working(i)=i-bin
end do
end do
! Map the random numbers into the 'distribution' array.
! This is made approximately proportional to the histogram.
do \(i=1, n \_\)samples
i_map=nint(rn(i)*(n_work-1)+1)
distribution(working(i map))) \& distribution(working(i_map)) +1
```

    end do
    ! Check the agreement between the distribution of the
! generated random numbers and the original histogram.
write (*, '(A)', advance='no') 'Original: '
write (*, '(10I6)') histogram*scale
write (*, '(A)', advance='no') 'Generated:'
write (*, '(10I6)') distribution
if (maxval(abs(histogram(1:)*scale-distribution(1:))) \&
<= tolerance*n_samples) then
write(*, '(A/)') 'Example 3 for RAND_GEN is correct.
end if
! Generate 20 integers in 1, 10 according to the distribution
! induced by the histogram.
call rand_gen(rn_20)
! Map from the uniform distribution to the induced distribution.
do i=1,n_samples_20
i_map=\overline{n}int(rn_\overline{20}(i) *(n_work-1)+1)
rand_num_20(i)=working(i_map)
end do
call show(rand_num_20,\&
'Twenty integers generāted according to the histogram:')
end

```

\section*{Output}

Example 3 for RAND_GEN is correct.

\section*{Example 4: Generating with a Cosine Distribution}

We generate random numbers based on the continuous distribution function
\[
p(x)=(1+\cos (x)) / 2 \pi,-\pi \leq x \leq \pi
\]

Using the cumulative
\[
q(x)=\int_{-\pi}^{x} p(t) d t=1 / 2+(x+\sin (x)) / 2 \pi
\]
we generate the samples by obtaining uniform samples \(u, 0<u<1\) and solve the equation
\[
q(x)-u=0,-\pi<x<\pi
\]

These are evaluated in vector form, that is all entries at one time, using Newton's method:
\[
x \leftarrow x-d x, d x=(q(x)-u) / p(x)
\]

An iteration counter forces the loop to terminate, but this is not often required although it is an important detail.
```

use rand gen int

```
```

    use show_int
    use Numerical_Libraries
    IMPLICIT NONE
    ! This is Example 4 for RAND_GEN.
integer i, i_map, k
integer, parameter :: n bins=36
integer, parameter :: Offset=18
integer, parameter :: n_samples=10000
integer, parameter :: n_samples_30=30
integer, parameter :: COUNT=15
real(kind(le0)) probabilities(n_bins)
real(kind(1e0)), dimension(n_bins) :: counts=0.0
real(kind(1e0)), dimension(n_samples) :: rn, x, f, fprime, dx
real(kind(1e0)), dimension(n_samples_30) :: rn_30, \&
x_30, f_30, fprime_30, dx_30
real(kind(\overline{le0)), parameter \: one=1e0, zero=0e0, half=0.5e0}
real(kind(1e0)), parameter :: tolerance=0.01
real(kind(1e0)) two_pi, omega
! Initialize values of 'two_pi' and 'omega'.
two_pi=2.0*const((/'pi'/))
omeg}\mp@subsup{\mp@code{g}}{}{\prime}=two_pi/n_bin
! Compute the probabilities for each bin according to
! the probability density (cos(x)+1)/(2*pi), -pi<x<pi.
do i=1,n_bins
probabilities(i)=(sin(omega*(i-offset)) \&
-sin(omega*(i-offset-1))+omega)/two_pi
end do
! Obtain uniform random numbers in (0,1).
call rand_gen(rn)
! Use Newton's method to solve the nonlinear equation:
! accumulated_distribution_function - random_number = 0.
x=zero; -k=0
solve_equation: do
f=(sin(x)+x)/two pi+half-rn
fprime=(one+cos(\overline{x}))/two_pi
dx=f/fprime
x=x-dx; k=k+1
if (maxval(abs(dx)) <= sqrt(epsilon(one)) \&
.or. k > COUNT) exit solve_equation
end do solve_equation
! Map the random numbers 'x' array into the 'counts' array.
do i=1,n_samples
i_map=int(x(i)/omega+offset)+1
counts(i_map)=counts(i_map)+one
end do

```
```

! Normalize the counts array.
counts=counts/n_samples
! Check that the generated random numbers are indeed
! based on the original distribution.
if (maxval(abs(counts(1:)-probabilities(1:))) \&
<= tolerance) then
write (*,'(a/)') 'Example 4 for RAND_GEN is correct.'
end if
! Generate 30 random numbers in (-pi,pi) according to
! the probability density (cos(x)+1)/(2*pi), -pi<x<pi.
call rand_gen(rn_30)
x_30=0.0; k=0
solve_equation_30: do
f_3\overline{0}=(sin(x_\overline{30})+x_30)/two_pi+half-rn_30
f\overline{prime_30=(\overline{one+cos}(x_30))/two_pi}
dx_30=\overline{f}_30/fprime_30
x_}\overline{3}0=x_30-dx_3
i\overline{f}(ma\overline{x}val(a\overline{bs}(dx_30)) <= sqrt(epsilon(one))\&
.or. k > COUNT) exit solve_equation_30
end do solve_equation_30
write(*,'(A)') 'Thirty random numbers generated ', \&
'according to the probability density ',\&
'pdf(x)=(cos(x)+1)/(2*pi), -pi<x<pi:'
call show(x_30)
end

```

\section*{Output}

Example 4 for RAND_GEN is correct.

\section*{RNGET}

Retrieves the current value of the seed used in the IMSL random number generators.

\section*{Required Arguments}

ISEED - The seed of the random number generator. (Output)
ISEED is in the range \((1,2147483646)\).

\section*{FORTRAN 90 Interface}

Generic: CALL RNGET (ISEED)
Specific: The specific interface name is RNGET.

\section*{FORTRAN 77 Interface}

Single: CALL RNGET (ISEED)

\section*{Description}

Routine RNGET retrieves the current value of the "seed" used in the IMSL random number generators. A reason for doing this would be to restart a simulation, using RNSET to reset the seed.

\section*{Example}

The following FORTRAN statements illustrate the use of RNGET:
```

    INTEGER ISEED
    ! Call RNSET to initialize the seed.
CALL RNSET(123457)
! Do some simulations.
...
...
CALL RNGET(ISEED)
Save ISEED. If the simulation is to be continued
in a different program, ISEED should be output,
possibly to a file.
When the simulations begun above are to be
restarted, restore ISEED to the value obtained
above and use as input to RNSET.
CALL RNSET(ISEED)
...
...

```

\section*{RNSET}

Initializes a random seed for use in the IMSL random number generators.

\section*{Required Arguments}

ISEED - The seed of the random number generator. (Input)
ISEED must be in the range \((0,2147483646)\). If ISEED is zero, a value is computed using the system clock; and, hence, the results of programs using the IMSL random number generators will be different at different times.

\section*{FORTRAN 90 Interface}

Generic: CALL RNSET (ISEED)

Specific: The specific interface name is RNSET .

\section*{FORTRAN 77 Interface}

Single: CALL RNSET (ISEED)

\section*{Description}

Routine RNSET is used to initialize the seed used in the IMSL random number generators. If the seed is not initialized prior to invocation of any of the routines for random number generation by calling RNSET, the seed is initialized via the system clock. The seed can be reinitialized to a clock-dependent value by calling RNSET with ISEED set to 0 .

The effect of RNSET is to set some values in a FORTRAN COMMON block that is used by the random number generators.

A common use of RNSET is in conjunction with RNGET to restart a simulation.

\section*{Example}

The following FORTRAN statements illustrate the use of RNSET:
INTEGER ISEED
\(!\)
!
CALL RNSET (0)
    -••
    ...
    CALL RNGET (ISEED)
    If the simulation is to be continued in a
    different program, ISEED should be output,
possibly to a file.
    .. .
                                    When the simulations begun above are to be
                                    restarted, restore ISEED to the value
                                    obtained above, and use as input to RNSET.
CALL RNSET (ISEED)
    Now continue the simulations.

\section*{RNOPT}

Selects the uniform \((0,1)\) multiplicative congruential pseudorandom number generator.

\section*{Required Arguments}

IOPT - Indicator of the generator. (Input)
The random number generator is either a multiplicative congruential generator with modulus \(2^{31}-1\) or a GFSR generator. IOPT is used to choose the multiplier and
whether or not shuffling is done, or is used to choose the GFSR method, or is used to choose the Mersenne Twister generator.

\section*{IOPT Generator}

1 The multiplier 16807 is used.
2 The multiplier 16807 is used with shuffling.
3 The multiplier 397204094 is used.
4 The multiplier 397204094 is used with shuffling.
5 The multiplier 950706376 is used.
6 The multiplier 950706376 is used with shuffling.
7 GFSR, with the recursion \(X_{t}=X_{t-1563} \oplus X_{t-96}\) is used.
8 A 32-bit Mersenne Twister generator is used. The real and double random numbers are generated from 32-bit integers.

9 A 64-bit Mersenne Twister generator is used. The real and double random numbers are generated from 64-bit integers. This ensures that all bits of both float and double are random.

\section*{FORTRAN 90 Interface}

Generic: CALL RNOPT (IOPT)
Specific: The specific interface name is RNOPT.

\section*{FORTRAN 77 Interface}

Single: CALL RNOPT (IOPT)

\section*{Description}

The uniform pseudorandom number generators use a multiplicative congruential method, with or without shuffling or a GFSR method, or the Mersenne Twister method. Routine RNOPT determines which method is used; and in the case of a multiplicative congruential method, it determines the value of the multiplier and whether or not to use shuffling. The description of RNUN may provide some guidance in the choice of the form of the generator. If no selection is made explicitly, the generators use the multiplier 16807 without shuffling. This form of the generator has been in use for some time (see Lewis, Goodman, and Miller, 1969). This is the generator formerly known as GGUBS in the IMSL Library. It is the "minimal standard generator" discussed by Park and Miller (1988).

Both of the Mersenne Twister generators have a period of \(2^{19937}-1\) and a 624 -dimensional equidistribution property. See Matsumoto et al. 1998 for details.
The IMSL Mersenne Twister generators are derived from code copyright (C) 1997-2002, Makoto Matsumoto and Takuji Nishimura, All rights reserved. It is subject to the following notice:

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The IMSL 32-bit Mersenne Twister generator is based on the Matsumoto and Nishimura code 'mt19937ar' and the 64-bit code is based on 'mt19937-64'.

\section*{Example}

The FORTRAN statement
CALL RNOPT (1)
would select the simple multiplicative congruential generator with multiplier 16807 . Since this is the same as the default, this statement would have no effect unless RNOPT had previously been called in the same program to select a different generator.

\section*{RNIN32}

Initializes the 32-bit Mersenne Twister generator using an array.

\section*{Required Arguments}
\(\boldsymbol{K E Y}\) - Integer array of length LEN used to initialize the 32-bit Mersenne Twister generator. (Input)

\section*{Optional Arguments}

LEN - Length of the array key. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL RNIN32 (KEY [, ...])
Specific: The specific interface name is \(S_{-}\)RNIN32.

\section*{FORTRAN 77 Interface}

Single: CALL RNIN32 (KEY, LEN)

\section*{Description}

By default, the Mersenne Twister random number generator is initialized using the current seed value (see RNGET). The seed is limited to one integer for initialization. This function allows an arbitrary length array to be used for initialization. This subroutine completely replaces the use of the seed for initialization of the 32-bit Mersenne Twister generator.

\section*{Example}

See routine RNGE32.

\section*{RNGE32}

Retrieves the current table used in the 32-bit Mersenne Twister generator.

\section*{Required Arguments}

MTABLE - Integer array of length 625 containing the table used in the 32-bit Mersenne Twister generator. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL RNGE32 (MTABLE)

Specific: The specific interface name is RNGE32

\section*{FORTRAN 77 Interface}

Single: CALL RNGE32 (MTABLE)

\section*{Description}

The values in the table contain the state of the 32-bit Mersenne Twister random number generator. The table can be used by RNSE 32 to set the generator back to this state.

\section*{Example}

In this example, four simulation streams are generated. The first series is generated with the seed used for initialization. The second series is generated using an array for initialization. The third series is obtained by resetting the generator back to the state it had at the beginning of the second stream. Therefore, the second and third streams are identical. The fourth stream is obtained by resetting the generator back to its original, uninitialized state, and having it reinitialize using the seed. The first and fourth streams are therefore the same.
```

    USE RNIN32 INT
    USE RNGE32_INT
    USE RNSET INT
    USE UMACH }\mp@subsup{}{}{-}\mathrm{ INT
    USE RNUN_INT
    IMPLICIT NONE
    INTEGER I, ISEED, NOUT
    INTEGER INIT(4)
    DATA INIT/291,564,837,1110/
    DATA ISEED/123457/
    INTEGER NR
    REAL R(5)
    INTEGER MTABLE (625)
    CHARACTER CLABEL(5)*5, FMT*8, RLABEL(3)*5
    RLABEL (1) = 'NONE'
    CLABEL (1) = 'NONE'
    DATA FMT/'(W10.4)'/
    NR=5
    CALL UMACH (2, NOUT)
    ISEED = 123457
    CALL RNOPT(8)
    CALL RNSET(ISEED)
    CALL RNUN(R)
    CALL WRRRL('FIRST STREAM OUTPUT',1,5,R,1,0, &
                            FMT, RLABEL, CLABEL)
    ! REINITIALIZE MERSENNE TWISTER SERIES WITH AN ARRAY
CALL RNIN32(INIT)
! SAVE THE STATE OF THE SERIES
CALL RNGE32 (MTABLE)
CALL RNUN(R)
CALL WRRRL('SECOND STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
! RESTORE THE STATE OF THE TABLE
CALL RNSE32 (MTABLE)
CALL RNUN (R)
CALL WRRRL('THIRD STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
! RESET THE SERIES - IT WILL REINITIALIZE FROM THE SEED
MTABLE (1)=1000
CALL RNSE32 (MTABLE)
CALL RNUN(R)
CALL WRRRL('FOURTH STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{First stream output} \\
\hline 0.4347 & 0.3522 & 0.0139 & 0.2091 & 0.4956 \\
\hline \multicolumn{5}{|c|}{Second stream output} \\
\hline 0.2486 & 0.2226 & 0.1111 & 0.9563 & 0.9846 \\
\hline \multicolumn{5}{|c|}{Third stream output} \\
\hline 0.2486 & 0.2226 & 0.1111 & 0.9563 & 0.9846 \\
\hline \multicolumn{5}{|c|}{Fourth stream output} \\
\hline 0.4347 & 0.3522 & 0.0139 & 0.2091 & 0.4956 \\
\hline
\end{tabular}

\section*{RNSE32}

Sets the current table used in the 32-bit Mersenne Twister generator.

\section*{Required Arguments}

MTABLE - Integer array of length 625 containing the table used in the 32-bit Mersenne Twister generator. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL RNSE32 (MTABLE)

Specific: The specific interface name is RNSE32

\section*{FORTRAN 77 Interface}

Single: CALL RNSE32 (MTABLE)

\section*{Description}

The values in MTABLE are the state of the 32-bit Mersenne Twister random number generator obtained by a call to RNGE 32. The values in the table can be used to restore the state of the generator.

Alternatively, if MTABLE [1] > 625 then the generator is set to its original, uninitialized, state.

\section*{Example}

See routine RNGE32.

\section*{RNIN64}

Initializes the 64-bit Mersenne Twister generator using an array.

\section*{Required Arguments}
\(\boldsymbol{K E Y}\) - Integer(kind=8) array of length LEN used to initialize the 64-bit Mersenne Twister generator. (Input)

\section*{Optional Arguments}

LEN - Length of the array key. (Input)

FORTRAN 90 Interface
Generic: CALL RNIN64 (KEY [,...])

Specific: The specific interface name is \(S_{-}\)RNIN64.

\section*{FORTRAN 77 Interface}
```

Single: CALL RNIN64 (KEY, LEN)

```

\section*{Description}

By default, the Mersenne Twister random number generator is initialized using the current seed value (see RNGET). The seed is limited to one integer for initialization. This function allows an arbitrary length array to be used for initialization. This subroutine completely replaces the use of the seed for initialization of the 64-bit Mersenne Twister generator.

\section*{RNGE64}

Retrieves the current table used in the 64-bit Mersenne Twister generator.

\section*{Required Arguments}

MTABLE - Integer(kind=8) array of length 313 containing the table used in the 64-bit Mersenne Twister generator. (Output)

\section*{FORTRAN 90 Interface}

Generic: CALL RNGE64 (MTABLE)

Specific: The specific interface name is RNGE64

\section*{FORTRAN 77 Interface}

\section*{Single: CALL RNGE64 (MTABLE)}

\section*{Description}

The values in the table contain the state of the 64-bit Mersenne Twister random number generator. The table can be used by RNSE 64 to set the generator back to this state.

\section*{Example}

In this example, four simulation streams are generated. The first series is generated with the seed used for initialization. The second series is generated using an array for initialization. The third series is obtained by resetting the generator back to the state it had at the beginning of the second stream. Therefore, the second and third streams are identical. The fourth stream is obtained by resetting the generator back to its original, uninitialized state, and having it reinitialize using the seed. The first and fourth streams are therefore the same.
```

USE RNIN64_INT
USE RNGE64_INT

```
```

    USE RNSET INT
    USE UMACH_INT
    USE RNUN_INT
    IMPLICIT }\mp@subsup{}{}{-}\mathrm{ NONE
    INTEGER I, ISEED, NOUT
    INTEGER(KIND=8) INIT(4)
    DATA INIT/291,564,837,1110/
    DATA ISEED/123457/
    INTEGER NR
    REAL R(5)
    INTEGER(KIND=8) MTABLE (313)
    CHARACTER CLABEL(5)*5, FMT*8, RLABEL(3)*5
    RLABEL(1) = 'NONE'
    CLABEL (1) = 'NONE'
    DATA FMT/'(W10.4)'/
    NR=5
    CALL UMACH (2, NOUT)
    ISEED = 123457
    CALL RNOPT(9)
    CALL RNSET(ISEED)
    CALL RNUN(R)
    CALL WRRRL('FIRST STREAM OUTPUT',1,5,R,1,0, &
        FMT, RLABEL, CLABEL)
    ! REINITIALIZE MERSENNE TWISTER SERIES WITH AN ARRAY
CALL RNIN64(INIT)
! SAVE THE STATE OF THE SERIES
CALL RNGE64 (MTABLE)
CALL RNUN(R)
CALL WRRRL('SECOND STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
! RESTORE THE STATE OF THE TABLE
CALL RNSE64 (MTABLE)
CALL RNUN(R)
CALL WRRRL('THIRD STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
! RESET THE SERIES - IT WILL REINITIALIZE FROM THE SEED
MTABLE (1)=1000
CALL RNSE64 (MTABLE)
CALL RNUN(R)
CALL WRRRL('FOURTH STREAM OUTPUT',1,5,R,1,0, \&
FMT, RLABEL, CLABEL)
END

```

\section*{Output}
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|c|}{First stream output} \\
\hline 0.5799 & 0.9401 & 0.7102 & 0.1640 & 0.5457 \\
\hline \multicolumn{5}{|c|}{Second stream output} \\
\hline 0.4894 & 0.7397 & 0.5725 & 0.0863 & 0.7588 \\
\hline \multicolumn{5}{|c|}{Third stream output} \\
\hline 0.4894 & 0.7397 & 0.5725 & 0.0863 & 0.7588 \\
\hline \multicolumn{5}{|c|}{Fourth stream output} \\
\hline 0.5799 & 0.9401 & 0.7102 & 0.1640 & 0.5457 \\
\hline
\end{tabular}

\section*{RNSE64}

Sets the current table used in the 64-bit Mersenne Twister generator.

\section*{Required Arguments}

MTABLE - Integer (kind=8) array of length 313 containing the table used in the 64-bit Mersenne Twister generator. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL RNSE64 (MTABLE)

Specific: The specific interface name is RNSE64

\section*{FORTRAN 77 Interface}

Single: CALL RNSE64 (MTABLE)

\section*{Description}

The values in MTABLE are the state of the 64-bit Mersenne Twister random number generator obtained by a call to RNGE 64. The values in the table can be used to restore the state of the generator. Alternatively, if MTABLE [1] > 313 then the generator is set to its original, uninitialized, state.

\section*{Example}

See function RNGE64.

\section*{RNUNF}

This function generates a pseudorandom number from a uniform \((0,1)\) distribution.

\section*{Function Return Value}

RNUNF - Function value, a random uniform \((0,1)\) deviate. (Output) See Comment 1.

\section*{Required Arguments}

None

\section*{FORTRAN 90 Interface}

Generic: RNUNF ()

Specific: The specific interface names are S_RNUNF and D_RNUNF.

\section*{FORTRAN 77 Interface}

Single: RNUNF ()
Double: The double precision name is DRNUNF.

\section*{Description}

Routine RNUNF is the function form of RNUN. The routine RNUNF generates pseudorandom numbers from a uniform \((0,1)\) distribution. The algorithm used is determined by RNOPt. The values returned by RNUNF are positive and less than 1.0.

If several uniform deviates are needed, it may be more efficient to obtain them all at once by a call to RNUN rather than by several references to RNUNF.

\section*{Comments}
1. If the generic version of this function is used, the immediate result must be stored in a variable before use in an expression. For example:
```

X = RNUNF (6)
Y = SQRT(X)

```
must be used rather than
\(\mathrm{Y}=\mathrm{SQRT}(\operatorname{RNUNF}(6))\)
If this is too much of a restriction on the programmer, then the specific name can be used without this restriction.
2. Routine RNSET can be used to initialize the seed of the random number generator. The routine RNOPT can be used to select the form of the generator.
3. This function has a side effect: it changes the value of the seed, which is passed through a common block.

\section*{Example}

In this example, RNUNF is used to generate five pseudorandom uniform numbers. Since RNOPT is not called, the generator used is a simple multiplicative congruential one with a multiplier of 16807.
```

USE RNUNF INT
USE RNSET INT
USE UMACH_INT
IMPLICIT NONE
INTEGER I, ISEED, NOUT

```
```

    REAL R(5)
    !
CALL UMACH (2, NOUT)
ISEED = 123457
CALL RNSET (ISEED)
DO 10 I=1, 5
R(I) = RNUNF()
10 CONTINUE
WRITE (NOUT,99999) R
99999 FORMAT (' Uniform random deviates: ', 5F8.4)
END

```

\section*{Output}

Uniform random deviates: \(\begin{array}{llllll}0.9662 & 0.2607 & 0.7663 & 0.5693 & 0.8448\end{array}\)

\section*{RNUN}

Generates pseudorandom numbers from a uniform \((0,1)\) distribution.

\section*{Required Arguments}
\(\boldsymbol{R}\) - Vector of length NR containing the random uniform \((0,1)\) deviates. (Output)

\section*{Optional Arguments}

NR - Number of random numbers to generate. (Input)
Default: NR = SIZE (R,1).

\section*{FORTRAN 90 Interface}

Generic: CALL RNUN (R [,...])

Specific: The specific interface names are S_RNUN and D_RNUN.

\section*{FORTRAN 77 Interface}

Single: CALL RNUN (NR, R)

Double: The double precision name is DRNUN.

\section*{Description}

Routine RNUN generates pseudorandom numbers from a uniform \((0,1)\) distribution using either a multiplicative congruential method or a generalized feedback shift register (GFSR) method, or the Mersenne Twister generator. The form of the multiplicative congruential generator is
\[
x_{i} \equiv c x_{i-1} \bmod \left(2^{31}-1\right)
\]

Each \(x_{i}\) is then scaled into the unit interval \((0,1)\). The possible values for \(c\) in the IMSL generators are 16807,397204094 , and 950706376 . The selection is made by the routine RNOPT. The choice of 16807 will result in the fastest execution time. If no selection is made explicitly, the routines use the multiplier 16807.

The user can also select a shuffled version of the multiplicative congruential generators. In this scheme, a table is filled with the first 128 uniform \((0,1)\) numbers resulting from the simple multiplicative congruential generator. Then, for each \(x_{i}\) from the simple generator, the low-order bits of \(x_{i}\) are used to select a random integer, \(j\), from 1 to 128 . The \(j\)-th entry in the table is then delivered as the random number; and \(x_{i}\), after being scaled into the unit interval, is inserted into the \(j\)-th position in the table.

The GFSR method is based on the recursion \(X_{t}=X_{t-1563} \oplus X_{t-96}\). This generator, which is different from earlier GFSR generators, was proposed by Fushimi (1990), who discusses the theory behind the generator and reports on several empirical tests of it.
Mersenne Twister(MT) is a pseudorandom number generating algorithm developed by Makoto Matsumoto and Takuji Nishimura in 1996-1997. MT has far longer period and far higher order of equidistribution than any other implemented generators. The values returned in \(R\) by RNUN are positive and less than 1.0. Values in R may be smaller than the smallest relative spacing, however. Hence, it may be the case that some value \(\mathrm{R}(i)\) is such that \(1.0-\mathrm{R}(i)=1.0\).

Deviates from the distribution with uniform density over the interval (A, B) can be obtained by scaling the output from RNUN. The following statements (in single precision) would yield random deviates from a uniform (A, B) distribution:
```

CALL RNUN (NR,R)
CALL SSCAL (NR,B-A,R,1)
CALL SADD (NR, A, R, 1)

```

\section*{Comments}

The routine RNSET can be used to initialize the seed of the random number generator. The routine RNOPT can be used to select the form of the generator.

\section*{Example}

In this example, RNUN is used to generate five pseudorandom uniform numbers. Since RNOPT is not called, the generator used is a simple multiplicative congruential one with a multiplier of 16807.
```

USE RNUN_INT
USE RNSET_INT
USE UMACH_INT
IMPLICIT NONE
INTEGER ISEED, NOUT, NR
REAL R(5)
CALL UMACH (2, NOUT)
NR= = 5
ISEED = 123457

```
!
```

    CALL RNSET (ISEED)
    CALL RNUN (R)
    WRITE (NOUT,99999) R
    9 9 9 9 9 ~ F O R M A T ~ ( ' ~ U n i f o r m ~ r a n d o m ~ d e v i a t e s : ~ ' , ~ 5 F 8 . 4 )
END

```

\section*{Output}
Uniform random deviates: \begin{tabular}{llllll}
0.9662 & 0.2607 & 0.7663 & 0.5693 & 0.8448
\end{tabular}

\section*{FAURE_INIT}

Shuffled Faure sequence initialization.

\section*{Required Arguments}

NDIM - The dimension of the hyper-rectangle. (Input)

STATE - An IMSL_FAURE pointer for the derived type created by the call to FAURE_INIT. The output contains information about the sequence. Use ?_IMSL_FAURE as the type, where ? \(\qquad\) is S or D depending on precision. (Output)

\section*{Optional Arguments}

NBASE - The base of the Faure sequence. (Input)
Default: The smallest prime number greater than or equal to NDIM.
NSKIP - The number of points to be skipped at the beginning of the Faure sequence.
(Input)
Default: \(\left\lfloor\right.\) base \(\left.^{m / 2-1}\right\rfloor\), where \(m=\lfloor\log B / \log\) base \(\rfloor\) and \(B\) is the largest machine representable integer.

\section*{FORTRAN 90 Interface}

Generic: CALL FAURE INIT (NDIM, STATE [,...])

Specific: The specific interface names are S_FAURE_INIT and D_FAURE INIT.

\section*{FAURE FREE}

Frees the structure containing information about the Faure sequence.

\section*{Required Arguments}

STATE - An IMSL_FAURE pointer containing the structure created by the call to FAURE_INIT. (Input/Output)

\section*{FORTRAN 90 Interface}

Generic: CALL FAURE_FREE (STATE)
Specific: The specific interface names are S_FAURE_FREE and D_FAURE_FREE.

\section*{FAURE_NEXT}

Computes a shuffled Faure sequence.

\section*{Required Arguments}

STATE - An IMSL_FAURE pointer containing the structure created by the call to FAURE_INIT. The structure contains information about the sequence. The structure should be freed using FAURE_FREE after it is no longer needed. (Input/Output)
\(\boldsymbol{N E X T}\) _PT - Vector of length NDIM containing the next point in the shuffled Faure sequence, where NDIM is the dimension of the hyper-rectangle specified in FAURE_INIT. (Output)

\section*{Optional Arguments}

IMSL_RETURN_SKIP — Returns the current point in the sequence. The sequence can be restarted by calling FAURE_INIT using this value for NSKIP, and using the same value for NDIM. (Input)

\section*{FORTRAN 90 Interface}

Generic: CALL FAURE_NEXT (STATE, NEXT_PT [,...])
Specific: The specific interface names are S_FAURE_NEXT and D_FAURE_NEXT.

\section*{Description}

The routines FAURE_INIT and FAURE_NEXT are used to generate shuffled Faure sequence of low discrepancy \(n\)-dimensional points. Low discrepency series fill an \(n\)-dimensional cube more uniformly than psuedo-random sequences, and are used in multivariate quadrature, simulation, and global optimization. Because of this uniformity, use of low discrepency series is generally more effiicient than psuedo-random series for multivariate Monte Carlo methods. See the IMSL routine QMC (Chapter 4, Integration and Differentiation) for a discussion of quasi-Monte Carlo quadrature based on low discrepancy series.

Discrepancy measures the deviation from uniformity of a point set.
The discrepancy of the point set \(x_{1}, \ldots, x_{n} \in[0,1]^{d}, d \geq 1\), is defined
\[
D_{n}^{(d)}=\sup _{E}\left|\frac{A(E ; n)}{n}-\lambda(E)\right|
\]
where the supremum is over all subsets of \([0,1]^{d}\) of the form
\(E=\left[0, t_{1}\right) \times \ldots \times\left[0, t_{d}\right), 0 \leq t_{j} \leq 1,1 \leq j \leq d\),
\(\lambda\) is the Lebesque measure, and \(A(E ; n)\) is the number of the \(x_{j}\) contained in \(E\).
The sequence \(x_{1}, x_{2}, \ldots\) of points \([0,1]^{d}\) is a low-discrepancy sequence if there exists a constant \(c(d)\), depending only on \(d\), such that
\[
D_{n}^{(d)} \leq c(d) \frac{(\log n)^{d}}{n}
\]
for all \(n>1\).
Generalized Faure sequences can be defined for any prime base \(b \geq d\). The lowest bound for the discrepancy is obtained for the smallest prime \(b \geq d\), so the optional argument NBASE defaults to the smallest prime greater than or equal to the dimension.
The generalized Faure sequence \(x_{1}, x_{2}, \ldots\), is computed as follows:
Write the positive integer \(n\) in its \(b\)-ary expansion,
\[
n=\sum_{i=0}^{\infty} a_{i}(n) b^{i}
\]
where \(a_{i}(n)\) are integers, \(0 \leq a_{i}(n)<b\).
The \(j\)-th coordinate of \(x_{n}\) is
\[
x_{n}^{(j)}=\sum_{k=0}^{\infty} \sum_{d=0}^{\infty} c_{k d}^{(j)} a_{d}(n) b^{-k-1}, \quad 1 \leq j \leq d
\]

The generator matrix for the series, \(c_{k d}^{(j)}\), is defined to be
\[
c_{k d}^{(j)}=j^{d-k} c_{k d}
\]
and \(c_{k d}\) is an element of the Pascal matrix,
\[
c_{k d}=\left\{\begin{array}{cc}
\frac{d!}{c!(d-c)!} & k \leq d \\
0 & k>d
\end{array}\right.
\]

It is faster to compute a shuffled Faure sequence than to compute the Faure sequence itself. It can be shown that this shuffling preserves the low-discrepancy property.

The shuffling used is the \(b\)-ary Gray code. The function \(G(n)\) maps the positive integer \(n\) into the integer given by its \(b\)-ary expansion.
The sequence computed by this function is \(x(G(n))\), where \(x\) is the generalized Faure sequence.

\section*{Example}

In this example, five points in the Faure sequence are computed. The points are in the threedimensional unit cube.

Note that FAURE_INIT is used to create a structure that holds the state of the sequence. Each call to FAURE_NEXT returns the next point in the sequence and updates the IMSL_FAURE structure. The final call to FAURE_FREE frees data items, stored in the structure, that were allocated by
FAURE_INIT.
```

use faure_int
implicit none
type (s_imsl_faure), pointer :: state
real(kind(1e0)) :: x(3)
integer,parameter :: ndim=3
integer :: k
CREATE THE STRUCTURE THAT HOLDS
THE STATE OF THE SEQUENCE.
call faure_init(ndim, state)
GET THE NEXT POINT IN THE SEQUENCE
do k=1,5
call faure_next(state, x)
write(*,'(\overline{3F15.3)') x(1), x(2) , x(3)}
enddo
FREE DATA ITEMS STORED IN
state STRUCTURE
call faure_free(state)
end

```

\section*{Output}
\begin{tabular}{lll}
0.334 & 0.493 & 0.064 \\
0.667 & 0.826 & 0.397 \\
0.778 & 0.270 & 0.175 \\
0.111 & 0.604 & 0.509 \\
0.445 & 0.937 & 0.842
\end{tabular}

\section*{IUMAG}

This routine handles MATH/LIBRARY and STAT/LIBRARY type INTEGER options.

\section*{Required Arguments}

PRODNM — Product name. Use either "МАТН" or "STAT." (Input)
ICHP - Chapter number of the routine that uses the options. (Input)
IACT - 1 if user desires to "get" or read options, or 2 if user desires to "put" or write options. (Input)

NUMOPT - Size of IOPTS. (Input)
IOPTS - Integer array of size NUMOPT containing the option numbers to "get" or "put." (Input)

IVALS - Integer array containing the option values. These values are arrays corresponding to the individual options in IOPTS in sequential order. The size of IVALS is the sum of the sizes of the individual options. (Input/Output)

\section*{FORTRAN 90 Interface}

Generic: CALL IUMAG (PRODNM, ICHP, IACT, NUMOPT, IOPTS, IVALS)

Specific: The specific interface name is IUMAG.

\section*{FORTRAN 77 Interface}

Single: CALL IUMAG (PRODNM, ICHP, IACT, NUMOPT, IOPTS, IVALS)

\section*{Description}

The Options Manager routine IUMAG reads or writes INTEGER data for some MATH/LIBRARY and STAT/LIBRARY codes. See Atchison and Hanson (1991) for more complete details.

There are MATH/LIBRARY routines in Chapters 1,2 , and 5 that now use IUMAG to communicate optional data from the user.

\section*{Comments}
1. Users can normally avoid reading about options when first using a routine that calls IUMAG.
2. Let I be any value between 1 and NUMOPT. A negative value of IOPTS(I) refers to option number-IOPTS(I) but with a different effect: For a "get" operation, the default values are returned in IVALS. For a "put" operation, the default values replace the current values. In the case of a "put," entries of IVALS are not allocated by the user and are not used by IUMAG.
3. Both positive and negative values of IOPTS can be used.

1 If the value is positive, print the next activity for any library routine that uses the Options Manager codes IUMAG, SUMAG, or DUMAG. Each printing step decrements the value if it is positive.
Default value is 0 .
2 If the value is 2, perform error checking in IUMAG, SUMAG, and DUMAG such as the verifying of valid option numbers and the validity of input data. If the value is 1 , do not perform error checking. Default value is 2 .

3 This value is used for testing the installation of IUMAG by other IMSL software. Default value is 3 .

\section*{Example}

The number of iterations allowed for the constrained least squares solver LCLSQ that calls L2LSQ is changed from the default value of \(\max (n r a, n c a)\) to the value 6 . The default value is restored after the call to LCLSQ. This change has no effect on the solution. It is used only for illustration. The first two arguments required for the call to IUMAG are defined by the product name, "MATH," and chapter number, 1 , where LCLSQ is documented. The argument IACT denotes a write or "put" operation. There is one option to change so NUMOPT has the value 1 . The arguments for the option number, 14 , and the new value, 6 , are defined by reading the documentation for LCLSQ.
```

USE IUMAG_INT
USE LCLSQ INT
USE UMACH-INT
USE SNRM2_INT
IMPLICIT NONE
Solve the following in the least squares sense:
3x1 + 2x2 + x3 = 3.3
4x1 + 2x2 + x3 = 2.3
2x1 + 2x2 + x3 = 1.3
x1 + x2 + x3 = 1.0
Subject to: x1 + x2 + x3 <= 1
0 <= x1 <= .5
0<= x2<= . 5
0<= x3 <= . 5
Declaration of variables
INTEGER ICHP, IPUT, LDA, LDC, MCON, NCA, NEWMAX, NRA, NUMOPT
PARAMETER (ICHP=1, IPUT=2, MCON=1, NCA=3, NEWMAX=14, NRA=4, \&
NUMOPT=1, LDA=NRA, LDC=MCON)
INTEGER IOPT(1), IRTYPE(MCON), IVAL(1), NOUT
REAL A(LDA,NCA), B(NRA), BC (MCON), C(LDC,NCA), RES(NRA), \&
RESNRM, XLB (NCA), XSOL (NCA), XUB (NCA)

```
```

! Data initialization
DATA A/3.0E0, 4.0E0, 2.0E0, 1.0E0, 2.0E0, 2.0E0, 2.0E0, 1.0E0, \&
1.0E0, 1.0EO, 1.0E0, 1.0E0/, B/3.3E0, 2.3EO, 1.3E0, 1.0E0/, \&
C/3*1.0E0/, BC/1.0E0/, IRTYPE/1/, XLB/3*0.0E0/, XUB/3*.5E0/
|-----------------------------------------------------------------------------
! Reset the maximum number of
iterations to use in the solver.
The value 14 is the option number.
The value 6 is the new maximum.
IOPT(1) = NEWMAX
IVAL(1) = 6
CALL IUMAG ('math', ICHP, IPUT, NUMOPT, IOPT, IVAL)
---------------------------------------
----------------------------------
Solve the bounded, constrained
least squares problem.
CALL LCLSQ (A, B, C, BC, B, IRTYPE, XLB, XUB, XSOL, RES=RES)
Compute the 2-norm of the residuals.
RESNRM = SNRM2 (NRA,RES,1)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) XSOL, RES, RESNRM
----------
------------------------------------
Reset the maximum number of
iterations to its default value.
This is not required but is
recommended programming practice.
IOPT(1) = -IOPT(1)
CALL IUMAG ('math', ICHP, IPUT, NUMOPT, IOPT, IVAL)
---------------------------------------
------------------------------------
!
99999 FORMAT (' The solution is ', 3F9.4, //, ' The residuals ', \&
'evaluated at the solution are ', /, 18X, 4F9.4, //, \&
' The norm of the residual vector is ', F8.4)
!
END

```

\section*{Output}

The solution is \(0.5000 \quad 0.3000 \quad 0.2000\)
The residuals evaluated at the solution are
\[
\begin{array}{cccc}
-1.0000 & 0.5000 & 0.5000 & 0.0000
\end{array}
\]

The norm of the residual vector is 1.2247

\section*{UMAG}

This routine handles MATH/LIBRARY and STAT/LIBRARY type REAL and double precision options.

\section*{Required Arguments}

PRODNM — Product name. Use either "МАТН" or "STAT." (Input)
ICHP - Chapter number of the routine that uses the options. (Input)
IACT - 1 if user desires to "get" or read options, or 2 if user desires to "put" or write options. (Input)

IOPTS - Integer array of size NUMOPT containing the option numbers to "get" or "put." (Input)

SVALS - Array containing the option values. These values are arrays corresponding to the individual options in IOPTS in sequential order. The size of SVALS is the sum of the sizes of the individual options. (Input/Output)

\section*{Optional Arguments}

NUMOPT - Size of IOPTS. (Input)
Default: NUMOPT \(=\) SIZE \((\) IOPTS, 1\()\).

\section*{FORTRAN 90 Interface}

Generic: CALL UMAG (PRODNM, ICHP, IACT, IOPTS, SVALS [,...])
Specific: The specific interface names are S_UMAG and D_UMAG.

\section*{FORTRAN 77 Interface}

Single: CALL SUMAG (PRODNM, ICHP, IACT, NUMOPT, IOPTS, SVALS)
Double: The double precision name is DUMAG.

\section*{Description}

The Options Manager routine SUMAG reads or writes REAL data for some MATH/LIBRARY and STAT/LIBRARY codes. See Atchison and Hanson (1991) for more complete details. There are MATH/LIBRARY routines in Chapters 1 and 5 that now use SUMAG to communicate optional data from the user.

\section*{Comments}
1. Users can normally avoid reading about options when first using a routine that calls Sumag.
2. Let I be any value between 1 and NUMOPT. A negative value of IOPTS(I) refers to option number-IOPTS(I) but with a different effect: For a "get" operation, the default values are returned in SVALS. For a "put" operation, the default values replace the current values. In the case of a "put," entries of SVALS are not allocated by the user and are not used by Sumag.
3. Both positive and negative values of IOPTS can be used.
4. Floating Point Options

1 This value is used for testing the installation of SUMAG by other IMSL software. Default value is 3.0 E 0 .

\section*{Example}

The rank determination tolerance for the constrained least squares solver LCLSQ that calls L2LSQ is changed from the default value of \(\operatorname{SQRT(AMACh(4))}\) to the value 0.01 . The default value is restored after the call to LCLSQ. This change has no effect on the solution. It is used only for illustration. The first two arguments required for the call to SUMAG are defined by the product name, "MATH," and chapter number, 1 , where LCLSQ is documented. The argument IACT denotes a write or "put" operation. There is one option to change so nUMOPT has the value 1. The arguments for the option number, 2 , and the new value, \(0.01 \mathrm{E}+0\), are defined by reading the documentation for LCLSQ.
```

USE UMAG_INT
USE LCLSQ_INT
USE UMACH_INT
USE SNRM2_INT
IMPLICIT NONE
Solve the following in the least squares sense:
3x1 + 2x2 + x3 = 3.3
4x1 + 2x2 + x3 = 2.3
2x1 + 2x2 + x3 = 1.3
x1 + x2 + x3 = 1.0
Subject to: x1 + x2 + x3 <= 1
0<= x1 <= . 5
0<= x2 <= . 5
0<= x3 <= . 5
Declaration of variables
INTEGER ICHP, IPUT, LDA, LDC, MCON, NCA, NEWTOL, NRA, NUMOPT

```
```

    PARAMETER (ICHP=1, IPUT=2, MCON=1, NCA=3, NEWTOL=2, NRA=4, &
                        NUMOPT=1, LDA=NRA, LDC=MCON)
    !
INTEGER IOPT(1), IRTYPE (MCON), NOUT
REAL A(LDA,NCA), B (NRA), BC (MCON), C(LDC,NCA), RES (NRA), \&
RESNRM, SVAL(1), XLB(NCA), XSOL (NCA), XUB (NCA)
Data initialization
DATA A/3.0E0, 4.0E0, 2.0E0, 1.0E0, 2.0E0, 2.0E0, 2.0E0, 1.0E0, \&
1.0E0, 1.0E0, 1.0E0, 1.0E0/, B/3.3E0, 2.3E0, 1.3E0, 1.0E0/, \&
C/3*1.0E0/, BC/1.0E0/, IRTYPE/1/, XLB/3*0.0E0/, XUB/3*.5E0/
Reset the rank determination
tolerance used in the solver.
The value 2 is the option number.
The value 0.01 is the new tolerance.
IOPT(1) = NEWTOL
SVAL (1) = 0.01E+0
CALL UMAG ('math', ICHP, IPUT, IOPT, SVAL)
--------------------------------------
----------------------------------
Solve the bounded, constrained
least squares problem.
CALL LCLSQ (A, B, C, BC, BC, IRTYPE, XLB, XUB, XSOL, RES=RES)
Compute the 2-norm of the residuals.
RESNRM = SNRM2 (NRA,RES,1)
Print results
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) XSOL, RES, RESNRM
-----------------------------------------
-----------------------------------
Reset the rank determination
tolerance to its default value.
This is not required but is
recommended programming practice.
IOPT(1) = -IOPT(1)
CALL UMAG ('math', ICHP, IPUT, IOPT, SVAL)
--------------------------------------
-------------------------------------
99999 FORMAT (' The solution is ', 3F9.4, //, ' The residuals ', \&
'evaluated at the solution are ', /, 18X, 4F9.4, //, \&
' The norm of the residual vector is ', F8.4)
!
END

```

\section*{Output}
The solution is \(0.5000 \quad 0.3000 \quad 0.2000\)
```

The residuals evaluated at the solution are
-1.0000 0.5000 0.5000 0.0000
The norm of the residual vector is 1.2247

```

\section*{SUMAG/DUMAG}

See umag.

\section*{PLOTP}

Prints a plot of up to 10 sets of points.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Vector of length NDATA containing the values of the independent variable. (Input)
A - Matrix of dimension NDATA by NFUN containing the NFUN sets of dependent variable values. (Input)

SYMBOL - CHARACTER string of length NFUN. (Input) SYMBOL(I : I) is the symbol used to plot function I.

XTITLE - CHARACTER string used to label the \(x\)-axis. (Input)
YTITLE - CHARACTER string used to label the \(y\)-axis. (Input)

TITLE - CHARACTER string used to label the plot. (Input)

\section*{Optional Arguments}

NDATA - Number of independent variable data points. (Input)
Default: NDATA \(=\operatorname{SIZE}(\mathrm{X}, 1)\).

NFUN - Number of sets of points. (Input)
NFUN must be less than or equal to 10 .
Default: NFUN = SIZE (A,2).

LDA - Leading dimension of A exactly as specified in the dimension statement of the calling program. (Input)
Default: LDA \(=\operatorname{SIZE}(\mathrm{A}, 1)\).
INC - Increment between elements of the data to be used. (Input)
PLOTP plots \(X(1+(I-1)\) * INC) for \(I=1,2, \ldots\), NDATA.
Default: \(\operatorname{INC}=1\).
\(\boldsymbol{R A N G E}\) - Vector of length four specifying minimum \(x\), maximum \(x\), minimum \(y\) and maximum \(y\). (Input)
PLOTP will calculate the range of the axis if the minimum and maximum of that range are equal.
Default: RANGE = 1.e0.

\section*{FORTRAN 90 Interface}

Generic: CALL PLOTP (X, A, SYMBOL, XTITLE, YTITLE, TITLE [,...])
Specific: The specific interface names are S_PLOTP and D_PLOTP.

\section*{FORTRAN 77 Interface}

Single: CALL PLOTP (NDATA, NFUN, X, A, LDA, INC, RANGE, SYMBOL, XTITLE, YTITLE, TITLE)

Double: The double precision name is DPLOTP.

\section*{Description}

Routine PLOTP produces a line printer plot of up to ten sets of points superimposed upon the same plot. A character " \(M\) " is printed to indicate multiple points. The user may specify the \(x\) and \(y\)-axis plot ranges and plotting symbols. Plot width and length may be reset in advance by calling PGOPT.

\section*{Comments}
1. Informational errors
\begin{tabular}{ccl} 
Type & Code & \\
3 & 7 & NFUN is greater than 10. Only the first 10 functions are plotted. \\
3 & 8 & TITLE is too long. TITLE is truncated from the right side. \\
3 & 9 & YTITLE is too long. YTITLE is truncated from the right side. \\
3 & 10 & \begin{tabular}{l} 
XTITLE is too long. XTITLE is truncated from the right side. The \\
maximum number of characters allowed depends on the page width \\
and the page length. See Comment 5 below for more information.
\end{tabular}
\end{tabular}
2. YTITLE and TITLE are automatically centered.
3. For multiple plots, the character \(M\) is used if the same print position is shared by two or more data sets.
4. Output is written to the unit specified by UMACH (see Reference Material).
5. Default page width is 78 and default page length is 60 . They may be changed by calling PGOPT in advance.

\section*{Example}

This example plots the sine and cosine functions from -3.5 to +3.5 and sets page width and length to 78 and 40 , respectively, by calling PGOPT in advance.
USE PLOTP_INT
USE CONST_INT
USE PGOPT_INT

IMPLICIT NONE
INTEGER I, IPAGE

REAL A(200,2), DELX, PI, RANGE(4), X(200)
CHARACTER SYMBOL*2
INTRINSIC COS, SIN
!
DATA SYMBOL/'SC'/
DATA RANGE/-3.5, 3.5, -1.2, 1.2/
!
\(\mathrm{PI}=3.14159\)
DELX \(=2 . * P I / 199\).
DO 10 I= 1, 200
\(\mathrm{X}(\mathrm{I})=-\mathrm{PI}+\mathrm{FLOAT}(\mathrm{I}-1)\) * DELX
\(\mathrm{A}(\mathrm{I}, 1)=\operatorname{SIN}(X(\mathrm{I}))\)
\(A(I, 2)=\operatorname{COS}(X(I))\)
10 CONTINUE
!
IPAGE = 78
CALL PGOPT (-1, IPAGE)
IPAGE = 40
CALL PGOPT ( -2, IPAGE)
CALL PLOTP (X, A, SYMBOL, 'X AXIS', 'Y AXIS', ' C = COS, S = SIN', \& RANGE=RANGE)
!
END

\section*{Output}



\section*{PRIME}

Decomposes an integer into its prime factors.

\section*{Required Arguments}
\(N\) - Integer to be decomposed. (Input)
\(\boldsymbol{N P F}\) - Number of different prime factors of \(\operatorname{ABS}(\mathrm{N})\). (Output) If N is equal to \(-1,0\), or \(1, \mathrm{NPF}\) is set to 0 .

IPF - Integer vector of length 13. (Output)
\(\operatorname{IPF}(I)\) contains the prime factors of the absolute value of \(N\), for \(I=1, \ldots, N P F\). The remaining \(13-\) NPF locations are not used.

IEXP - Integer vector of length 13. (Output)
\(\operatorname{IEXP}(\mathrm{I})\) is the exponent of \(\operatorname{IPF}(\mathrm{I})\), for \(I=1, \ldots, \mathrm{NPF}\). The remaining \(13-\mathrm{NPF}\) locations are not used.

IPW - Integer vector of length 13. (Output)
\(\operatorname{IPW}(I)\) contains the quantity \(\operatorname{IPF}(I) * * \operatorname{IEXP}(I)\), for \(I=1, \ldots, N P F\). The remaining 13 - NPF locations are not used.

\section*{FORTRAN 90 Interface}

Generic: CALL PRIME (N, NPF, IPF, IPW)
Specific: The specific interface name is PRIME.

\section*{FORTRAN 77 Interface}

Single: CALL PRIME (N, NPF, IPF, IEXP, IPW)

\section*{Description}

Routine PRIME decomposes an integer into its prime factors. The number to be factored, \(N\), may not have more than 13 distinct factors. The smallest number with more than 13 factors is about \(1.3 \times 10^{16}\). Most computers do not allow integers of this size.
The routine PRIME is based on a routine by Brenner (1973).

\section*{Comments}

The output from PRIME should be interpreted in the following way:
```

ABS(N)= IPF(1)**IEXP(1) * ... * * IPF(NPF)**IEXP(NPF).

```

\section*{Example}

This example factors the integer \(144=2^{4} 3^{2}\).
```

USE PRIME INT
USE UMACH_INT
IMPLICIT NONE
INTEGER N
PARAMETER (N=144)
INTEGER IEXP(13), IPF(13), IPW(13), NOUT, NPF
CALL PRIME (N, NPF, IPF, IEXP, IPW)
CALL UMACH (2, NOUT)
WRITE (NOUT,99999) N, IPF(1), IPF(2), IEXP(1), IEXP(2), IPW(1), \&
IPW(2), NPF
99999 FORMAT (' The prime factors for', I5, ' are: ', /, 10X, 2I6, // \&
IEXP =', 2I6, /, ' IPW =', 2I6, /, ' NPF =', I6, /)
END

```
!
!

\section*{Output}
```

The prime factors for 144 are:
2 3

```
\begin{tabular}{lrrr} 
IEXP & \(=\) & 4 & 2 \\
IPW & \(=\) & 16 & 9 \\
NPF & \(=\) & 2 &
\end{tabular}

\section*{CONST}

This function returns the value of various mathematical and physical constants.

\section*{Function Return Value}

CONST - Value of the constant. (Output)
See Comment 1.

\section*{Required Arguments}

NAME - Character string containing the name of the desired constant. (Input)
See Comment 3 for a list of valid constants.

\section*{FORTRAN 90 Interface}

Generic: CONST (NAME)

Specific: The specific interface names are S_CONST and D_CONST.

\section*{FORTRAN 77 Interface}

Single: CONST (NAME)
Double: The double precision name is DCONST.

\section*{Description}

Routine CONST returns the value of various mathematical and physical quantities. For all of the physical values, the Systeme International d'Unites (SI) are used.

The reference for constants are indicated by the code in [ ] Comment above.
[1] Cohen and Taylor (1986)
[2] Liepman (1964)
[3] Precomputed mathematical constants
The constants marked with an E before the [ ] are exact (to machine precision).
To change the units of the values returned by CONST, see CUNIT.

\section*{Comments}
1. If the generic version of this function is used, the immediate result must be stored in a variable before use in an expression. For example:
```

X = CONST('PI')
Y = CoS(x)

```
must be used rather than
```

Y = COS(CONST(`PI')).

```

If this is too much of a restriction on the programmer, then the specific name can be used without this restriction.
2. The case of the character string in NAME does not matter. The names "PI", "Pi", "Pi", and "pi" are equivalent.
3. The units of the physical constants are in SI units (meter kilogram-second).
4. The names allowed are as follows:
\begin{tabular}{llll} 
Name & Description & Value & Ref. \\
\hline AMU & Atomic mass unit & \(1.6605402 \mathrm{E}-27 \mathrm{~kg}\) & {\([1]\)} \\
ATM & Standard atm pressure & \(1.01325 \mathrm{E}+5 \mathrm{~N} / \mathrm{m}^{2} \mathrm{E}\) & {\([2]\)} \\
AU & Astronomical unit & \(1.496 \mathrm{E}+11 \mathrm{~m}\) & {[]} \\
Avogadro & Avogadro's number & \(6.0221367 \mathrm{E}+231 / \mathrm{mole}\) & {\([1]\)} \\
Boltzman & Boltzman's constant & \(1.380658 \mathrm{E}-23 \mathrm{~J} / \mathrm{K}\) & {\([1]\)} \\
C & Speed of light & \(2.997924580 \mathrm{E}+8 m / \mathrm{E}\) & {\([1]\)} \\
Catalan & Catalan's constant & \(0.915965 \ldots \mathrm{E}\) & {\([3]\)} \\
E & Base of natural logs & \(2.718 \ldots \mathrm{E}\) & {\([3]\)} \\
ElectronCharge & Electron change & \(1.60217733 \mathrm{E}-19 \mathrm{C}\) & {\([1]\)} \\
ElectronMass & Electron mass & \(9.1093897 \mathrm{E}-31 \mathrm{~kg}\) & {\([1]\)} \\
ElectronVolt & Electron volt & \(1.60217733 \mathrm{E}-19 \mathrm{~J}\) & {\([1]\)} \\
Euler & Euler's constant gamma & \(0.577 \ldots \mathrm{E}\) & {\([3]\)} \\
Faraday & Faraday constant & \(9.6485309 \mathrm{E}+4 \mathrm{C} / \mathrm{mole}\) & {\([1]\)} \\
FineStructure & fine structure & \(7.29735308 \mathrm{E}-3\) & {\([1]\)} \\
Gamma & Euler's constant & \(0.577 \ldots \mathrm{E}\) & {\([3]\)} \\
Gas & Gas constant & \(8.314510 \mathrm{~J} / \mathrm{mole} / \mathrm{k}\) & {\([1]\)} \\
Gravity & Gravitational constant & \(6.67259 \mathrm{E}-11 \mathrm{~N} * m^{2} / \mathrm{kg}{ }^{2}\) & {\([1]\)} \\
Hbar & Planck constant \(/ 2\) pi & \(1.05457266 \mathrm{E}-34 \mathrm{~J} * s\) & {\([1]\)} \\
PerfectGasVolume & Std vol ideal gas & \(2.241383 \mathrm{E}-2 m^{3} / \mathrm{mole}\) & {\([*]\)} \\
Pi & Pi & \(3.141 \ldots \mathrm{E}\) & {\([3]\)} \\
Planck & Planck's constant \(h\) & \(6.6260755 \mathrm{E}-34 \mathrm{~J} * s\) & {\([1]\)} \\
ProtonMass & Proton mass & \(1.6726231 \mathrm{E}-27 \mathrm{~kg}\) & {\([1]\)}
\end{tabular}
\begin{tabular}{lllc} 
Name & Description & Value & Ref. \\
\hline Rydberg & Rydberg's constant & \(1.0973731534 \mathrm{E}+7 / m\) & {\([1]\)} \\
SpeedLight & Speed of light & \(2.997924580 \mathrm{E}+8 m / s \mathrm{E}\) & {\([1]\)} \\
StandardGravity & Standard \(g\) & \(9.80665 m / s^{2} \mathrm{E}\) & {\([2]\)} \\
StandardPressure & Standard atm pressure & \(1.01325 \mathrm{E}+5 \mathrm{~N} / m^{2} \mathrm{E}\) & {\([2]\)} \\
StefanBoltzmann & Stefan-Boltzman & \(5.67051 \mathrm{E}-8 \mathrm{~W} / \mathrm{K}^{4} / m^{2}\) & {\([1]\)} \\
WaterTriple & Triple point of water & \(2.7316 \mathrm{E}+2 \mathrm{~K} \mathrm{E}\) & {\([2]\)}
\end{tabular}

\section*{Example}

In this example, Euler's constant \(\gamma\) is obtained and printed. Euler's constant is defined to be
\[
\gamma=\lim _{n \rightarrow \infty}\left[\sum_{k=1}^{n-1} \frac{1}{k}-\ln n\right]
\]
```

USE CONST INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
REAL GAMA
CALL UMACH (2, NOUT)
GAMA = CONST('GAMMA')
WRITE (NOUT,*) 'GAMMA = ', GAMA
END

```

\section*{Output}
```

GAMMA = 0.5772157

```

For another example, see CUNIT.

\section*{CUNIT}

Converts X in units XUNITS to Y in units YUNITS.

\section*{Required Arguments}
\(\boldsymbol{X}\) - Value to be converted. (Input)
XUNITS - Character string containing the name of the units for X . (Input) See Comments for a description of units allowed.
\(\boldsymbol{Y}\) - Value in YUNITS corresponding to X in XUNITS. (Output)

YUNITS - Character string containing the name of the units for Y . (Input) See Comments for a description of units allowed.

\section*{FORTRAN 90 Interface}

Generic: CALL CUNIT (X, XUNITS, Y, YUNITS [, ...])

Specific: The specific interface names are S_CUNIT and D_CUNIT.

\section*{FORTRAN 77 Interface}

\section*{Single: CALL CUNIT (X, XUNITS, Y, YUNITS)}

Double: The double precision name is DCUNIT.

\section*{Description}

Routine CUNIT converts a value expressed in one set of units to a value expressed in another set of units.

The input and output units are checked for consistency unless the input unit is "SI". SI means the Systeme International d'Unites. This is the meter-kilogram-second form of the metric system. If the input units are "SI", then the input is assumed to be expressed in the SI units consistent with the output units.

\section*{Comments}
1. Strings XUNITS and YUNITS have the form \(\mathrm{U}_{1} * \mathrm{U}_{2}{ }^{*} \ldots{ }^{*} \mathrm{U}_{m} / \mathrm{V}_{1} \ldots \mathrm{~V}_{n}\), where \(\mathrm{U}_{i}\) and \(\mathrm{V}_{i}\) are the names of basic units or are the names of basic units raised to a power. Examples are, "METER * KILOGRAM/SECOND", "M * KG/S", "METER", or "M/KG".
2. The case of the character string in XUNITS and YUNITS does not matter. The names "METER", "Meter" and "meter" are equivalent.
3. If XUNITS is "SI", then X is assumed to be in the standard international units corresponding to YUNITS. Similarly, if YUNITS is
"SI", then \(Y\) is assumed to be in the standard international units corresponding to XUNITS.
4. The basic unit names allowed are as follows:

Units of time
day, hour \(=\mathrm{hr}, \min =\) minute, \(\mathrm{s}=\mathrm{sec}=\) second, year
Units of frequency
Hertz \(=\mathrm{Hz}\)
```

Units of mass
AMU,g= gram, lb = pound, ounce = oz, slug
Units of distance
Angstrom, AU, feet = foot = ft, in = inch, m = meter = metre, micron, mile, mill,
parsec, yard
Units of area
acre
Units of volume
l= liter = litre
Units of force
dyne, N = Newton, poundal
Units of energy
BTU(thermochemical), Erg, J = Joule
Units of work
W = watt
Units of pressure
ATM = atomosphere, bar, Pascal
Units of temperature
degC = Celsius, degF = Fahrenheit, degK = Kelvin
Units of viscosity
poise, stoke
Units of charge
Abcoulomb, C = Coulomb, statcoulomb
Units of current
A = ampere, abampere, statampere,
Units of voltage
Abvolt, V = volt
Units of magnetic induction
T = Tesla, Wb = Weber
Other units
1, farad, mole, Gauss, Henry, Maxwell, Ohm
The following metric prefixes may be used with the above units. Note that the one or two letter prefixes may only be used with one letter unit abbreviations.

| A | Atto | 1.E -18 |
| :--- | :--- | :--- |
| F | Femto | $1 . \mathrm{E}-15$ |
| P | Pico | $1 . \mathrm{E}-12$ |
| N | Nano | $1 . \mathrm{E}-9$ |
| U | Micro | $1 . \mathrm{E}-6$ |
| M | Milli | $1 . \mathrm{E}-3$ |

```
\begin{tabular}{lll} 
C & Centi & \(1 . \mathrm{E}-2\) \\
D & Deci & \(1 . \mathrm{E}-1\) \\
DK & Deca & \(1 . \mathrm{E}+2\) \\
K & Kilo & \(1 . \mathrm{E}+3\) \\
& Myriad & \(1 . \mathrm{E}+4\) (no single letter prefix; M means milli \\
& Mega & \(1 . \mathrm{E}+6\) (no single letter prefix; M means milli \\
G & Giga & \(1 . \mathrm{E}+9\) \\
T & Tera & \(1 . \mathrm{E}+12\)
\end{tabular}
5. Informational error
\begin{tabular}{ccl} 
Type & Code & \\
3 & 8 & \begin{tabular}{l} 
A conversion of units of mass to units of force was required for \\
consistency.
\end{tabular}
\end{tabular}

\section*{Example}

The routine CONST is used to obtain the speed on light, \(c\), in SI units. CUNIT is then used to convert \(c\) to mile/second and to parsec/year. An example involving substitution of force for mass is required in conversion of Newtons/Meter \({ }^{2}\) to Pound/Inch \({ }^{2}\).
```

USE CONST INT
USE CUNIT INT
USE UMACH_INT
IMPLICIT NONE
INTEGER NOUT
REAL CMH, CMS, CPY, CPSI
Get output unit number
CALL UMACH (2, NOUT)
CMS = CONST('SpeedLight')
WRITE (NOUT,*) 'Speed of Light = ', CMS, ' meter/second'
Get speed of light in mile/second
CALL CUNIT (CMS, 'SI', CMH, 'Mile/Second')
WRITE (NOUT,*) 'Speed of Light = ', CMH, ' mile/second'
Get speed of light in parsec/year
CALL CUNIT (CMS, 'SI', CPY, 'Parsec/Year')
WRITE (NOUT,*) 'Speed of Light = ', CPY, ' Parsec/Year'
Convert Newton/Meter**2 to
Pound/Inch**2.
CALL CUNIT(1.E0, 'Newton/Meter**2', CPSI, \&
'Pound/Inch**2')
WRITE (NOUT,*)' Atmospheres, in Pound/Inch**2 = ',CPSI
END

```

\section*{Output}
```

Speed of Light = 299792440.0 meter/second
Speed of Light = 186282.39 mile/second
Speed of Light = 0.3063872 Parsec/Year

```
```

*** WARNING ERROR 8 from CUNIT. A conversion of units of mass to units of
*** force was required for consistency.
Atmospheres, in Pound/Inch**2 = 1.4503773E-4

```

\section*{HYPOT}

This functions computes \(\operatorname{SQRT}\left(A^{* *} 2+B * * 2\right)\) without underflow or overflow.

\section*{Function Return Value}

HYPOT \(-\operatorname{SQRT}\left(A * * 2+B^{* *} 2\right)\). (Output)

\section*{Required Arguments}

A - First parameter. (Input)
B - Second parameter. (Input)

\section*{FORTRAN 90 Interface}

Generic: \(\quad\) HYPOT (A, B)
Specific: The specific interface names are S_HYPOT and D_HYPOT.

\section*{FORTRAN 77 Interface}

Single: \(\quad \operatorname{HYPOT}(A, B)\)

Double: The double precision name is DНYPOT.

\section*{Description}

Routine нYРОт is based on the routine PYTHAG, used in EISPACK 3. This is an update of the work documented in Garbow et al. (1972).

\section*{Example}

Computes
\[
c=\sqrt{a^{2}+b^{2}}
\]
where \(a=10^{20}\) and \(b=2 \times 10^{20}\) without overflow.
```

USE HYPOT INT
USE UMACH_INT

```
IMPLICIT NONE
\(!\)
```

    INTEGER NOUT
    REAL A, B, C
    !
A = 1.0E+20
B = 2.0E+20
C = HYPOT (A,B)
! Get output unit number
CALL UMACH (2, NOUT)
! Print the results
WRITE (NOUT,'(A,1PE10.4)') ' C = ', C
END

```

\section*{Output}
\(C=2.2361 \mathrm{E}+20\)

\section*{MP_SETUP \\ REQUIRED}

Initializes or finalizes MPI.

\section*{Function Return Value}

Number of nodes, MP NPROCS, in the communicator, MP LIBRARY WORLD. (Output)
Returned when MP_SETUP is called with no arguments:
MP_NPROCS \(=\) MP_SETUP ( ) .

\section*{Required Argument}

None.

\section*{Optional Arguments}

NOTE - Character string 'Final'. (Input)
With 'Final' all pending error messages are sent from the nodes to the root and printed. If any node should STOP after printing messages, then MPI_Finalize() and a STOP are executed. Otherwise, only MPI_Finalize () is called. The character string 'Final' is the only valid string for this argument.
\(N\) - Size of array to be allocated for timing. (Input)
When this argument is supplied, the array MPI_NODE_PRIORITY is allocated with MP_PROCS components. The matrix products A .x. B are timed individually at each node of the machine. The elapsed time is noted and sorted to determine the node priority order. A and B are allocated to size N by N , and initialized with random data. The priority order is finally broadcast to the other nodes.

\section*{FORTRAN 90 Interface}

\author{
MP_SETUP ( [, ...])
}

\section*{Description}

Following a call to the function MP_SETUP (), the module MPI_node_int will contain information about the number of processors, the rank of a processor, the communicator for IMSL Fortran Numerical Library, and the usage priority order of the node machines:
```

MODULE MPI_NODE_INT
INTEGER, ALLOCATABLE :: MPI_NODE_PRIORITY(:)
INTEGER, SAVE :: MP_LIBRARY_WORLD = huge(1)
LOGICAL, SAVE :: MPI_ROOT_WORKS = .TRUE.
INTEGER, SAVE : : MP_RANK = 0, MP_NPROCS = 1
END MODULE

```

When the function MP_SETUP () is called with no arguments, the following events occur:
- If MPI has not been initialized, it is first initialized. This step uses the routines MPI_Initialized() and possibly MPI_Init(). Users who choose not to call MP_SETUP () must make the required initialization call before using any IMSL Fortran Numerical Library code that relies on MPI for its execution. If the user's code calls an IMSL Fortran Numerical Library function utilizing the box data type and MPI has not been initialized, then the computations are performed on the root node. The only MPI routine always called in this context is MPI_Initialized (). The name MP_SETUP is pushed onto the subprogram or call stack.
- If MP_LIBRARY_WORLD equals its initial value (=huge (1)) then MPI_COMM_WORLD, the default MPI communicator, is duplicated and becomes its handle. This uses the routine MPI_Comm_dup (). Users can change the handle of MP_LIBRARY_WORLD as required by their application code. Often this issue can be ignored.
- The integers MP_RANK and MP_NPROCS are respectively the node's rank and the number of nodes in the communicator, MP_LIBRARY_WORLD. Their values require the routines MPI_Comm_size() and MPI_Comm_rank(). The default values are important when MPI is not initialized and a box data type is computed. In this case the root node is the only node and it will do all the work. No calls to MPI communication routines are made when MP_NPROCS = 1 when computing the box data type functions. A program can temporarily assign this value to force box data type computation entirely at the root node. This is desirable for problems where using many nodes would be less efficient than using the root node exclusively.
- The array MPI_NODE_PRIORITY (: ) is not allocated unless the user allocates it. The IMSL Fortran Numerical Library codes use this array for assigning tasks to processors, if it is allocated. If it is not allocated, the default priority of the nodes is
\(\left(0,1, \ldots, M P \_N P R O C S-1\right)\). Use of the function call MP_SETUP (N) allocates the array, as explained below. Once the array is allocated its size is MP_NPROCS. The contents of the array is a permutation of the integers \(0, \ldots, \mathrm{MP}\) _NPROCS-1. Nodes appearing at the start of the list are used first for parallel computing. A node other than the root can avoid any computing,
except receiving the schedule, by setting the value MPI_NODE_PRIORITY(I) < 0. This means that node |MPI_NODE_PRIORITY(I)| will be sent the task schedule but will not perform any significant work as part of box data type function evaluations.
- The LOGICAL flag MPI_ROOT_WORKS designates whether or not the root node participates in the major computation of the tasks. The root node communicates with the other nodes to complete the tasks but can be designated to do no other work. Since there may be only one processor, this flag has the default value . TRUE., assuring that one node exists to do work. When more than one processor is available users can consider assigning MPI_ROOT_WORKS=.FALSE. This is desirable when the alternate nodes have equal or greater computational resources compared with the root node. Parallel Example 4 illustrates this usage. A single problem is given a box data type, with one rack. The computing is done at the node, other than the root, with highest priority. This example requires more than one processor since the root does no work.

When the generic function MP_SETUP ( N ) is called, where N is a positive integer, a call to MP_SETUP () is first made, using no argument. Use just one of these calls to MP_SETUP (). This initializes the MPI system and the other parameters described above. The array
MPI_NODE_PRIORITY (: ) is allocated with size MP_NPROCS. Then DOUBLE PRECISION matrix products \(\bar{C}=A B\), where \(A\) and \(B\) are \(N\) by \(N\) matrices, are computed at each node and the elapsed time is recorded. These elapsed times are sorted and the contents of MPI_NODE_PRIORITY (:) are permuted in accordance with the shortest times yielding the highest priority. All the nodes in the communicator MP_LIBRARY_WORLD are timed. The array MPI_NODE_PRIORITY(:) is then broadcast from the root to the remaining nodes of MP_LIBRARY_WORLD using the routine MPI_Bcast (). Timing matrix products to define the node priority is relevant because the effort to compute \(C\) is comparable to that of many linear algebra computations of similar size. Users are free to define their own node priority and broadcast the array MPI_NODE_PRIORITY (: ) to the alternate nodes in the communicator.
To print any IMSL Fortran Numerical Library error messages that have occurred at any node, and to finalize MPI, use the function call MP_SETUP ('Final'). The case of the string 'Final' is not important. Any error messages pending will be discarded after printing on the root node. This is triggered by popping the name 'MP_SETUP' from the subprogram stack or returning to Level 1 in the stack. Users can obtain error messages by popping the stack to Level 1 and still continuing with MPI calls. This requires executing call e1pop ('MP_SETUP'). To continue on after summarizing errors execute call e1psh ('MP_SETUP'). More details about the error processor are found in Reference Material chapter of this manual.
Messages are printed by nodes from largest rank to smallest, which is the root node. Use of the routine MPI_Finalize() is made within MP_SETUP ('Final'), which shuts down MPI. After MPI_Finalize() is called, the value of MP_NPROCS \(=0\). This flags that MPI has been initialized and terminated. It cannot be initialized again in the same program unit execution. No MPI routine is defined when MP_NPROCS has this value.

\section*{Examples}

\section*{Parallel Example (parallel_ex01.f90)}
```

use linear_operators
use mpi_setup_int

```
```

    implicit none
    ! This is the equivalent of Parallel Example 1 for .ix., with box data types
! and functions.
integer, parameter : : n=32, nr=4
real(kind(1e0)) :: one=1e0
real(kind(1e0)), dimension(n,n,nr) :: A, b, x, err(nr)
! Setup for MPI.
MP_NPROCS=MP_SETUP()
! Generate random matrices for A and b:
A = rand(A); b=rand(b)
! Compute the box solution matrix of Ax = b.
x = A.ix. b
! Check the results.
err = norm(b - (A .x. x))/(norm(A)*norm(x) +norm(b))
if (ALL(err <= sqrt(epsilon(one))) .and. MP_RANK == 0) \&
write (*,*) 'Parallel Example 1 is correct..'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{Parallel Example (parallel_ex04.f90)}

Here an alternate node is used to compute the majority of a single application, and the user does not need to make any explicit calls to MPI routines. The time-consuming parts are the evaluation of the eigenvalue-eigenvector expansion, the solving step, and the residuals. To do this, the rank-2 arrays are changed to a box data type with a unit third dimension. This uses parallel computing. The node priority order is established by the initial function call, MP_SETUP ( \(n\) ) . The root is restricted from working on the box data type by assigning MPI_ROOT_WORKS=.false. This example anticipates that the most efficient node, other than the root, will perform the heavy computing. Two nodes are required to execute.
```

    use linear operators
    use mpi_setup_int
    implicit none
    ! This is the equivalent of Parallel Example 4 for matrix exponential.
! The box dimension has a single rack.
integer, parameter :: n=32, k=128, nr=1
integer i
real(kind(le0)), parameter :: one=1e0, t_max=one, delta_t=t_max/(k-1)
real(kind(1e0)) err(nr), sizes(nr), A(n,n,nr)
real(kind(le0)) t(k), y(n,k,nr), y prime(n,k,nr)
complex(kind(le0)), dimension(n,nr) : : x (n,n,nr), z_0, \&
Z_1(n,nr,nr), y_0, d

```
```

! Setup for MPI. Establish a node priority order.
! Restrict the root from significant computing.
! Illustrates using the 'best' performing node that
! is not the root for a single task.
MP_NPROCS=MP_SETUP (n)
MPI_ROOT_WORKS=.false.
! Generate a random coefficient matrix.
A = rand(A)
! Compute the eigenvalue-eigenvector decomposition
! of the system coefficient matrix on an alternate node.
D = EIG(A, W=X)
! Generate a random initial value for the ODE system.
y_0 = rand(y_0)
! Solve complex data system that transforms the initial
! values, X z_0=y_0.
z_1= X .ix. y_0 ; z_0(:,nr) = z_1(:,nr,nr)
! The grid of points where a solution is computed:
t = (/(i*delta_t,i=0,k-1)/)
! Compute y and y' at the values t(1:k).
! With the eigenvalue-eigenvector decomposition AX = XD, this
! is an evaluation of EXP(A t)y_0 = y(t).
y = X .x.exp (spread(d(:,n\overline{r}),2,k)*spread(t,1,n))*spread(z_0(:,nr),2,k)
! This is y', derived by differentiating y(t).
y_prime = X .x. \&
spread(d(:,nr),2,k)*exp(spread(d(:,nr),2,k)*spread(t,1,n))* \&
spread(z_0(:,nr),2,k)
! Check results. Is y' - Ay = 0?
err = norm(y_prime-(A .x. y))
sizes=norm(y_prime)+norm(A)*norm(y)
if (ALL(err <= sqrt(epsilon(one))*sizes) .and. MP RANK == 0) \&
write (*,*) 'Parallel Example 4 is correct.'
! See to any error messages and quit MPI.
MP_NPROCS=MP_SETUP('Final')
end

```

\section*{Reference Material}

\section*{Contents}
User Errors ..... 1763
Automatic Workspace Allocation ..... 1785
Machine-Dependent Constants ..... 1769
Matrix Storage Modes ..... 1775
Reserved Names ..... 1784
Deprecated and Renamed Routines ..... 1785

\section*{User Errors}

IMSL routines attempt to detect user errors and handle them in a way that provides as much information to the user as possible. To do this, we recognize various levels of severity of errors, and we also consider the extent of the error in the context of the purpose of the routine; a trivial error in one situation may be serious in another. IMSL routines attempt to report as many errors as they can reasonably detect. Multiple errors present a difficult problem in error detection because input is interpreted in an uncertain context after the first error is detected.

\section*{What Determines Error Severity}

In some cases, the user's input may be mathematically correct, but because of limitations of the computer arithmetic and of the algorithm used, it is not possible to compute an answer accurately. In this case, the assessed degree of accuracy determines the severity of the error. In cases where the routine computes several output quantities, if some are not computable but most are, an error condition exists. The severity depends on an assessment of the overall impact of the error.

\section*{Terminal errors}

If the user's input is regarded as meaningless, such as \(N=-1\) when " \(N\) " is the number of equations, the routine prints a message giving the value of the erroneous input argument(s) and the reason for the erroneous input. The routine will then cause the user's program to stop. An error in which the user's input is meaningless is the most severe error and is called a terminal error. Multiple terminal error messages may be printed from a single routine.

\section*{Informational errors}

In many cases, the best way to respond to an error condition is simply to correct the input and rerun the program. In other cases, the user may want to take actions in the program itself based on errors that occur. An error that may be used as the basis for corrective action within the program is called an informational error. If an informational error occurs, a user-retrievable code is set. A routine can return at most one informational error for a single reference to the routine. The codes for the informational error codes are printed in the error messages.

\section*{Other errors}

In addition to informational errors, IMSL routines issue error messages for which no userretrievable code is set. Multiple error messages for this kind of error may be printed. These errors, which generally are not described in the documentation, include terminal errors as well as less serious errors. Corrective action within the calling program is not possible for these errors.

\section*{Kinds of Errors and Default Actions}

Five levels of severity of errors are defined in the MATH/LIBRARY. Each level has an associated PRINT attribute and a STOP attribute. These attributes have default settings (YES or NO), but they may also be set by the user. The purpose of having multiple error severity levels is to provide independent control of actions to be taken for errors of different severity. Upon return from an IMSL routine, exactly one error state exists. (A code 0 "error" is no informational error.) Even if more than one informational error occurs, only one message is printed (if the PRINT attribute is YES). Multiple errors for which no corrective action within the calling program is reasonable or necessary result in the printing of multiple messages (if the PRINT attribute for their severity level is YES). Errors of any of the severity levels except level 5 may be informational errors.

Level 1: \(\quad\) Note. A note is issued to indicate the possibility of a trivial error or simply to provide information about the computations. Default attributes: PRINT=NO, STOP=NO

Level 2: Alert. An alert indicates that the user should be advised about events occurring in the software. Default attributes: PRINT=NO, STOP=NO

Level 3: Warning. A warning indicates the existence of a condition that may require corrective action by the user or calling routine. A warning error may be issued because the results are accurate to only a few decimal places, because some of the output may be erroneous but most of the output is correct, or because some assumptions underlying the analysis technique are violated. Often no corrective action is necessary and the condition can be ignored. Default attributes: \(\mathrm{PRINT}=\mathrm{YES}, \mathrm{STOP}=\mathrm{NO}\)

Level 4: Fatal.A fatal error indicates the existence of a condition that may be serious. In most cases, the user or calling routine must take corrective action to recover. Default attributes: \(\mathrm{PRINT}=\mathrm{YES}, \mathrm{STOP}=\mathrm{YES}\)

Level 5: Terminal.A terminal error is serious. It usually is the result of an incorrect specification, such as specifying a negative number as the number of equations. These errors may also be caused by various programming errors impossible to diagnose correctly in FORTRAN. The resulting error message may be perplexing to the user. In
such cases, the user is advised to compare carefully the actual arguments passed to the routine with the dummy argument descriptions given in the documentation. Special attention should be given to checking argument order and data types.

A terminal error is not an informational error because corrective action within the program is generally not reasonable. In normal usage, execution is terminated immediately when a terminal error occurs. Messages relating to more than one terminal error are printed if they occur. Default attributes: PRINT=YES, STOP=YES

The user can set PRINT and STOP attributes by calling ERSET as described in "Routines for Error Handling."

\section*{Errors in Lower-Level Routines}

It is possible that a user's program may call an IMSL routine that in turn calls a nested sequence of lower-level IMSL routines. If an error occurs at a lower level in such a nest of routines and if the lower-level routine cannot pass the information up to the original user-called routine, then a traceback of the routines is produced. The only common situation in which this can occur is when an IMSL routine calls a user-supplied routine that in turn calls another IMSL routine.

\section*{Routines for Error Handling}

There are three ways in which the user may interact with the IMSL error handling system: (1) to change the default actions, (2) to retrieve the integer code of an informational error so as to take corrective action, and (3) to determine the severity level of an error. The routines to use are ERSET, IERCD, and N1RTY, respectively.

\section*{ERSET}

Change the default printing or stopping actions when errors of a particular error severity level occur.

\section*{Required Arguments}

IERSVR — Error severity level indicator. (Input)
If IERSVR \(=0\), actions are set for levels 1 to 5 . If IERSVR is 1 to 5 , actions are set for errors of the specified severity level.

IPACT - Printing action. (Input)

\section*{IPACT Action}
-1 Do not change current setting(s).
\(0 \quad\) Do not print.
1 Print.
2 Restore the default setting(s).

ISACT - Stopping action. (Input)
ISACT Action
-1 Do not change current setting(s).
\(0 \quad\) Do not stop.
1 Stop.
2 Restore the default setting(s).

\section*{FORTRAN 90 Interface}

Generic: CALL ERSET (IERSVR, IPACT, ISACT)
Specific: The specific interface name is ERSET.

\section*{FORTRAN 77 Interface}

Single: CALL ERSET (IERSVR, IPACT, ISACT)

\section*{IERCD and N1RTY}

The last two routines for interacting with the error handling system, IERCD and N1RTY, are INTEGER functions and are described in the following material.

IERCD retrieves the integer code for an informational error. Since it has no arguments, it may be used in the following way:
```

ICODE = IERCD()

```

The function retrieves the code set by the most recently called IMSL routine.
N1RTY retrieves the error type set by the most recently called IMSL routine. It is used in the following way:

ITYPE = N1RTY(1)
ITYPE \(=1,2,4\), and 5 correspond to error severity levels \(1,2,4\), and 5 , respectively. ITYPE \(=3\) and ITYPE \(=6\) are both warning errors, error severity level 3 . While ITYPE \(=3\) errors are informational errors \((\operatorname{IERCD}() \neq 0), \operatorname{ITYPE}=6\) errors are not informational errors \((\operatorname{IERCD}()=0)\).

For software developers requiring additional interaction with the IMSL error handling system, see Aird and Howell (1991).

\section*{Examples}

\section*{Changes to default actions}

Some possible changes to the default actions are illustrated below. The default actions remain in effect for the kinds of errors not included in the call to ERSET.

To turn off printing of warning error messages:
```

CALL ERSET (3,0,-1)

```

To stop if warning errors occur:
CALL ERSET ( \(3,-1,1\) )
To print all error messages:
CALL ERSET (0, 1, -1)
To restore all default settings:
CALL ERSET (0,2,2)

\section*{Use of informational error to determine program action}

In the program segment below, the Cholesky factorization of a matrix is to be performed. If it is determined that the matrix is not nonnegative definite (and often this is not immediately obvious), the program is to take a different branch.
CALL LFTDS (A, FACT)

IF (IERCD() .EQ. 2) THEN
! Handle matrix that is not nonnegative definite

END IF

\section*{Examples of errors}

The program below illustrates each of the different types of errors detected by the MATH/LIBRARY routines.

The error messages refer to the argument names that are used in the documentation for the routine, rather than the user's name of the variable used for the argument. In the message generated by IMSL routine LINRG in this example, reference is made to \(N\), whereas in the program a literal was used for this argument.
```

USE_IMSL_LIBRARIES
INTEGER N
PARAMETER (N=2)
REAL A(N,N), AINV (N,N), B(N), X(N)
DATA A/2.0, -3.0, 2.0, -3.0/
DATA B/1.0, 2.0/
Turn on printing and turn off
stopping for all error types.
CALL ERSET (0, 1, 0)
CALL LSARG (A, B, X)
CALL LINRG (A, AINV, N = -1)
END

```

\section*{Output}
```

*** FATAL ERROR 2 from LSARG. The input matrix is singular. Some of
*** the diagonal elements of the upper triangular matrix U of the
*** LU factorization are close to zero.
*** TERMINAL ERROR 1 from LINRG. The order of the matrix must be positive
*** while N = -1 is given.

```

\section*{Example of traceback}

The next program illustrates a situation in which a traceback is produced. The program uses the IMSL quadrature routines QDAG and QDAGS to evaluate the double integral
\[
\int_{0}^{1} \int_{0}^{1}(x+y) d x d y=\int_{0}^{1} g(y) d y
\]
where
\[
g(y)=\int_{0}^{1}(x+y) d x=\int_{0}^{1} f(x) d x, \text { with } f(x)=x+y
\]

Since both QDAG and QDAGS need 2500 numeric storage units of workspace, and since the workspace allocator uses some space to keep track of the allocations, 6000 numeric storage units of space are explicitly allocated for workspace. Although the traceback shows an error code associated with a terminal error, this code has no meaning to the user; the printed message contains all relevant information. It is not assumed that the user would take corrective action based on knowledge of the code.
```

USE QDAGS_INT
REAL A, B, ERRABS, ERREST, ERRREL, G, RESULT
EXTERNAL G
A = 0.0
B = 1.0
ERRABS = 0.0
ERRREL = 0.001
CALL QDAGS (G, A, B, RESULT, ERRABS, ERRREL, ERREST)
WRITE (*,*) RESULT, ERREST
END
REAL FUNCTION G (ARGY)
USE QDAG_INT
REAL ARGY
INTEGER IRULE
REAL C, D, ERRABS, ERREST, ERRREL, F, Y
COMMON /COMY/ Y
EXTERNAL F

```
```

Y = ARGY

```
Y = ARGY
C}=0.
```

C}=0.

```
!
!
!
```

        D = 1.0
        ERRABS = 0.0
        ERRREL = -0.001
        IRULE = 1
    !
CALL QDAG (F, C, D, G, ERRABS, ERRREL, IRULE, ERREST)
RETURN
END
!
REAL FUNCTION F (X)
REAL X
REAL Y
COMMON /COMY/ Y
!
F = X + Y
RETURN
END

```

\section*{Output}
```

*** TERMINAL ERROR 4 from Q2AG. The relative error desired ERRREL =
*** -1.000000E-03. It must be at least zero.
Here is a traceback of subprogram calls in reverse order:
Routine name Error type Error code
------------ ------------------------
Q2AG
QDAG
0 0
QDAG 0
Q2AGS 0
QDAGS
0
USER 0
0
Q2AGS 0
0 0

```

\section*{Machine-Dependent Constants}

The function subprograms in this section return machine-dependent information and can be used to enhance portability of programs between different computers. The routines IMACH, and AMACH describe the computer's arithmetic. The routine UMACH describes the input, ouput, and error output unit numbers.

\section*{IMACH}

This function retrieves machine integer constants that define the arithmetic used by the computer.

\section*{Function Return Value}
\(\operatorname{IMACH}(1)=\) Number of bits per integer storage unit.
\(\operatorname{IMACH}(2)=\) Number of characters per integer storage unit:
Integers are represented in \(M\)-digit, base \(A\) form as
\[
\sigma \sum_{k=0}^{M} x_{k} A^{k}
\]
where \(\sigma\) is the sign and \(0 \leq x_{k}<A, k=0, \ldots, M\).
Then,
\(\operatorname{IMACH}(3)=A\), the base.
\(\operatorname{IMACH}(4)=M\), the number of base- \(A\) digits.
\(\operatorname{IMACH}(5)=A^{M}-1\), the largest integer.
The machine model assumes that floating-point numbers are represented in normalized \(N\)-digit, base \(B\) form as
\[
\sigma B^{E} \sum_{k=1}^{N} x_{k} B^{-k}
\]
where \(\sigma\) is the sign, \(0<x_{1}<B, 0 \leq x_{k}<B, \mathrm{k}=2, \ldots, N\) and \(E_{\min } \leq E \leq E_{\max }\). Then,
\(\operatorname{IMACH}(6)=B\), the base.
\(\operatorname{IMACH}(7)=N_{s}\), the number of base- \(B\) digits in single precision.
\(\operatorname{IMACH}(8)=E_{\text {min }_{s}}\), the smallest single precision exponent.
\(\operatorname{IMACH}(9)=E_{\text {max }_{s}}\), the largest single precision exponent.
\(\operatorname{IMACH}(10)=N_{d}\), the number of base- \(B\) digits in double precision.
\(\operatorname{IMACH}(11)=E_{\min _{d}}\), the smallest double precision exponent.
\(\operatorname{IMACH}(12)=E_{\max _{d}}\), the number of base- \(B\) digits in double precision

\section*{Required Arguments}
\(\boldsymbol{I}\) - Index of the desired constant. (Input)

\section*{FORTRAN 90 Interface}

Generic: IMACH (I)
Specific: The specific interface name is IMACH.

\section*{FORTRAN 77 Interface}

Single: IMACH (I)

\section*{AMACH}

The function subprogram AMACH retrieves machine constants that define the computer's singleprecision or double precision arithmetic. Such floating-point numbers are represented in normalized \(N\)-digit, base \(B\) form as
\[
\sigma B^{E} \sum_{k=1}^{N} x_{k} B^{-k}
\]
where \(\sigma\) is the sign, \(0<x_{1}<B, 0 \leq x_{k}<B, k=2, \ldots, N\) and
\[
E_{\min } \leq E \leq E_{\max }
\]

\section*{Function Return Value}
\(\operatorname{AMACH}(1)=B^{E_{\text {min }}-1}\), the smallest normalized positive number.
\(\operatorname{AMACH}(2)=B^{E_{\max }}\left(1-B^{-N}\right)\), the largest number.
\(\operatorname{AMACH}(3)=B^{-N}\), the smallest relative spacing.
\(\operatorname{AMACH}(4)=B^{1-N}\), the largest relative spacing.
\(\operatorname{AMACH}(5)=\log _{10}(B)\).
\(\operatorname{AMACH}(6)=\mathrm{NaN}\) (quiet not a number).
AMACH(7)=positive machine infinity.
\(\operatorname{AMACH}(8)=\) negative machine infinity.

See Comment 1 for a description of the use of the generic version of this function.
See Comment 2 for a description of \(\min , \max\), and \(N\).

\section*{Required Arguments}
\(\boldsymbol{I}\) - Index of the desired constant. (Input)

\section*{FORTRAN 90 Interface}

Generic: AMACH (I)

Specific: The specific interface names are S_AMACH and D_AMACH.

\section*{FORTRAN 77 Interface}

Single: AMACH (I)
Double: The double precision name is DMACH.

\section*{Comments}
1. If the generic version of this function is used, the immediate result must be stored in a variable before use in an expression. For example:
```

X = AMACH(I)
Y = SQRT(X)

```
must be used rather than
\(\mathrm{Y}=\operatorname{SQRT}(\mathrm{AMACH}(\mathrm{I}))\).

If this is too much of a restriction on the programmer, then the specific name can be used without this restriction.
2. Note that for single precision \(B=\operatorname{IMACH}(6), N=\operatorname{IMACH}(7)\).
\(E_{\text {min }}=\operatorname{IMACH}(8)\), and \(E_{\text {max }},=\operatorname{IMACH}(9)\).
For double precision \(B=\operatorname{IMACH}(6), N=\operatorname{IMACH}(10)\).
\(E_{\text {min }}=\operatorname{IMACH}(11)\), and \(E_{\text {max }},=\operatorname{IMACH}(12)\).
3. The IEEE standard for binary arithmetic (see IEEE 1985) specifies quiet NaN (not a number) as the result of various invalid or ambiguous operations, such as \(0 / 0\). The intent is that \(\operatorname{AMACH}(6)\) return a quiet NaN . On IEEE format computers that do not support a quiet NaN , a special value near \(\operatorname{AMACH}(2)\) is returned for \(\operatorname{AMACH}(6)\). On computers that do not have a special representation for infinity, \(\operatorname{AMACH}(7)\) returns the same value as AMACH (2).

\section*{DMACH}

See Amach.

\section*{IFNAN(X)}

This logical function checks if the argument X is NaN (not a number).

\section*{Function Return Value}

IFNAN - Logical function value. True is returned if the input argument is a NAN. Otherwise, False is returned. (Output)

\section*{Required Arguments}
\(\boldsymbol{X}\) - Argument for which the test for NAN is desired. (Input)

\section*{FORTRAN 90 Interface}

Generic: IFNAN (X)
Specific: The specific interface names are S_IFNAN and D_IFNAN.

\section*{FORTRAN 77 Interface}

Single: IFNAN (X)
Double: The double precision name is DIFNAN.

\section*{Example}
```

    USE IFNAN_INT
    USE AMACH ' INT
    USE UMACH_INT
    INTEGER NOUT
    REAL X
    CALL UMACH (2, NOUT)
    X = AMACH (6)
    IF (IFNAN(X)) THEN
    WRITE (NOUT,*) ' X is NaN (not a number).'
    ELSE
        WRITE (NOUT,*) ' X = ', X
    END IF
    END

```
!
!
!

\section*{Output}

X is NaN (not a number).

\section*{Description}

The logical function IFNAN checks if the single or double precision argument X is NaN (not a number). The function IFNAN is provided to facilitate the transfer of programs across computer systems. This is because the check for NaN can be tricky and not portable across computer systems that do not adhere to the IEEE standard. For example, on computers that support the IEEE standard for binary arithmetic (see IEEE 1985), NaN is specified as a bit format not equal to itself. Thus, the check is performed as
```

IFNAN = X .NE. X

```

On other computers that do not use IEEE floating-point format, the check can be performed as:
```

IFNAN = X .EQ. AMACH(6)

```

The function IFNAN is equivalent to the specification of the function Isnan listed in the Appendix, (IEEE 1985). The above following example illustrates the use of IFNAN. If \(x\) is NaN , a message is
printed instead of X. (Routine UMACH, which is described in the following section, is used to retrieve the output unit number for printing the message.)

\section*{UMACH}

Routine UMACH sets or retrieves the input, output, or error output device unit numbers.

\section*{Required Arguments}
\(N\) - Integer value indicating the action desired. If the value of \(N\) is negative, the input, output, or error output unit number is reset to NUNIT. If the value of N is positive, the input, output, or error output unit number is returned in NUNIT. See the table in argument NUNIT for legal values of N . (Input)

NUNIT - The unit number that is either retrieved or set, depending on the value of input argument N. (Input/Output)
The arguments are summarized by the following table:
\begin{tabular}{|r|l|}
\hline N & Effect \\
\hline 1 & Retrieves input unit number in NUNIT. \\
\hline 2 & Retrieves output unit number in NUNIT. \\
\hline 3 & Retrieves error output unit number in NUNIT. \\
\hline-1 & Sets the input unit number to NUNIT. \\
\hline-2 & Sets the output unit number to NUNIT. \\
\hline-3 & Sets the error output unit number to NUNIT. \\
\hline
\end{tabular}

\section*{FORTRAN 90 Interface}

Generic: CALL UMACH (N, NUNIT)
Specific: The specific interface name is UMACH.

\section*{FORTRAN 77 Interface}

Single: CALL UMACH (N, NUNIT)

\section*{Description}

Routine UMACH sets or retrieves the input, output, or error output device unit numbers. UMACH is set automatically so that the default FORTRAN unit numbers for standard input, standard output, and standard error are used. These unit numbers can be changed by inserting a call to UMACH at the beginning of the main program that calls MATH/LIBRARY routines. If these unit numbers are changed from the standard values, the user should insert an appropriate OPEN statement in the calling program.

\section*{Example}

In the following example, a terminal error is issued from the MATH/LIBRARY AMACH function since the argument is invalid. With a call to UMACH, the error message will be written to a local file named "CHECKERR".
```

USE AMACH_INT
USE UMACH_INT
INTEGER - N, NUNIT
REAL X
N}=
NUNIT = 9
CALL UMACH (-3, NUNIT)
OPEN (UNIT=NUNIT, FILE=' CHECKERR')
X = AMACH (N)
END

```
!

\section*{Output}
```

The output from this example, written to "CHECKERR" is:
*** TERMINAL ERROR 5 from AMACH. The argument must be between 1 and 8
*** inclusive. N = 0

```

\section*{Matrix Storage Modes}

In this section, the word matrix will be used to refer to a mathematical object, and the word array will be used to refer to its representation as a FORTRAN data structure.

\section*{General Mode}

A general matrix is an \(N \times N\) matrix \(A\). It is stored in a FORTRAN array that is declared by the following statement:

DIMENSION A(LDA,N)
The parameter LDA is called the leading dimension of A. It must be at least as large as N. IMSL general matrix subprograms only refer to values \(A_{i j}\) for \(i=1, \ldots, N\) and \(j=1, \ldots, N\). The data type of a general array can be one of REAL, DOUBLE PRECISION, or COMPLEX. If your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX can also be declared.

\section*{Rectangular Mode}

A rectangular matrix is an \(M \times N\) matrix \(A\). It is stored in a FORTRAN array that is declared by the following statement:
```

DIMENSION A(LDA,N)

```

The parameter LDA is called the leading dimension of A. It must be at least as large as M. IMSL rectangular matrix subprograms only refer to values \(A_{i j}\) for \(i=1, \ldots, M\) and \(j=1, \ldots, N\). The data
type of a rectangular array can be REAL, DOUBLE PRECISION, or COMPLEX. If your FORTRAN compiler allows, you can declare the nonstandard data type DOUBLE COMPLEX.

\section*{Symmetric Mode}

A symmetric matrix is a square \(N \times N\) matrix \(A\), such that \(A^{T}=A\). \(\left(A^{T}\right.\) is the transpose of \(A\).) It is stored in a FORTRAN array that is declared by the following statement:
```

DIMENSION A(LDA,N)

```

The parameter LDA is called the leading dimension of A. It must be at least as large as N. IMSL symmetric matrix subprograms only refer to the upper or to the lower half of A (i.e., to values \(A_{i j}\) for \(i=1, \ldots, N\) and \(j=i, \ldots, N\), or \(A_{i j}\) for \(j=1, \ldots, N\) and \(\left.i=j, \ldots, N\right)\). The data type of a symmetric array can be one of REAL or DOUBLE PRECISION. Use of the upper half of the array is denoted in the BLAS that compute with symmetric matrices, see Chapter 9, Basic Matrix/Vector Operations, using the CHARACTER*1 flag UPLO \(={ }^{\prime} \mathrm{U}^{\prime}\). Otherwise, UPLO \(=' \mathrm{~L}\) ' denotes that the lower half of the array is used.

\section*{Hermitian Mode}

A Hermitian matrix is a square \(N \times N\) matrix \(A\), such that
\[
\bar{A}^{T}=A
\]

The matrix
\[
\bar{A}
\]
is the complex conjugate of \(A\) and
\[
A^{H} \equiv \bar{A}^{T}
\]
is the conjugate transpose of \(A\). For Hermitian matrices, \(A^{H}=A\). The matrix is stored in a FORTRAN array that is declared by the following statement:
```

DIMENSION A(LDA,N)

```

The parameter LDA is called the leading dimension of A. It must be at least as large as \(N\). IMSL Hermitian matrix subprograms only refer to the upper or to the lower half of A (i.e., to values \(A_{i j}\) for \(i=1, \ldots, N\) and \(j=i, \ldots, N\). , or \(A_{i j}\) for \(j=1, \ldots, N\) and \(\left.i=j, \ldots, N\right)\). Use of the upper half of the array is denoted in the BLAS that compute with Hermitian matrices, see Chapter 9, Basic Matrix/Vector Operations, using the CHARACTER* 1 flag UPLO \(={ }^{\prime} \mathrm{U}^{\prime}\). Otherwise, UPLO \(=\prime^{\prime} \mathrm{L}^{\prime}\) denotes that the lower half of the array is used. The data type of a Hermitian array can be COMPLEX or, if your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX.

\section*{Triangular Mode}

A triangular matrix is a square \(N \times N\) matrix \(A\) such that values \(A_{i j}=0\) for \(i<j\) or \(A_{i j}=0\) for \(i>j\). The first condition defines a lower triangular matrix while the second condition defines an upper triangular matrix. A lower triangular matrix \(A\) is stored in the lower triangular part of a FORTRAN array A. An upper triangular matrix is stored in the upper triangular part of a FORTRAN array. Triangular matrices are called unit triangular whenever \(A_{j j}=1, j=1, \ldots, N\). For unit triangular matrices, only the strictly lower or upper parts of the array are referenced. This is
denoted in the BLAS that compute with triangular matrices, see Chapter 9, Basic Matrix/Vector Operations, using the CHARACTER* 1 flag DIAGNL \(={ }^{\prime} U^{\prime}\). Otherwise, DIAGNL \(={ }^{\prime} \mathrm{N}^{\prime}\) denotes that the diagonal array terms should be used. For unit triangular matrices, the diagonal terms are each used with the mathematical value 1 . The array diagonal term does not need to be 1.0 in this usage. Use of the upper half of the array is denoted in the BLAS that compute with triangular matrices, see Chapter 9, Basic Matrix/Vector Operations, using the CHARACTER* 1 flag UPLO \(=\) 'U'. Otherwise, UPLO \(=\) ' L' denotes that the lower half of the array is used. The data type of an array that contains a triangular matrix can be one of REAL, DOUBLE PRECISION, or COMPLEX. If your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX can also be declared.

\section*{Band Storage Mode}

A band matrix is an \(M \times N\) matrix \(A\) with all of its nonzero elements "close" to the main diagonal. Specifically, values \(A_{i j}=0\) if \(i-j>\) NLCA or \(j-i>\) NUCA. The integers NLCA and NUCA are the lower and upper band widths. The integer \(m=\) NLCA + NUCA +1 is the total band width. The diagonals, other than the main diagonal, are called codiagonals. While any \(M \times N\) matrix is a band matrix, the band matrix mode is most useful only when the number of nonzero codiagonals is much less than \(m\).

In the band storage mode, the NLCA lower codiagonals and NUCA upper codiagonals are stored in the rows of a FORTRAN array of dimension \(m \times N\). The elements are stored in the same column of the array as they are in the matrix. The values \(A_{i j}\) inside the band width are stored in array positions \((i-j+\) NUCA \(+1, j)\). This array is declared by the following statement:

DIMENSION A(LDA,N)
The parameter LDA is called the leading dimension of \(A\). It must be at least as large as \(m\). The data type of a band matrix array can be one of REAL, DOUBLE PRECISION, COMPLEX or, if your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX . Use of the CHARACTER* 1 flag TRANS \(={ }^{\prime} N^{\prime}\) in the BLAS, see Chapter 9, Basic Matrix/Vector Operations, specifies that the matrix \(A\) is used. The flag value
\[
\text { TRANS }=\text { 'T' uses } A^{T}
\]
while
\[
\text { TRANS }=\text { 'C' uses } \bar{A}^{T}
\]

For example, consider a real \(5 \times 5\) band matrix with 1 lower and 2 upper codiagonals, stored in the FORTRAN array declared by the following statements:
```

PARAMETER (N=5, NLCA=1, NUCA=2 )
REAL A (NLCA+NUCA+1,N)

```

The matrix \(A\) has the form
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{12} & A_{13} & 0 & 0 \\
A_{21} & A_{22} & A_{23} & A_{24} & 0 \\
0 & A_{32} & A_{33} & A_{34} & A_{35} \\
0 & 0 & A_{43} & A_{44} & A_{45} \\
0 & 0 & 0 & A_{54} & A_{55}
\end{array}\right]
\]

As a FORTRAN array, it is
\[
A=\left[\begin{array}{ccccc}
\times & \times & A_{13} & A_{24} & A_{35} \\
\times & A_{12} & A_{23} & A_{34} & A_{45} \\
A_{11} & A_{22} & A_{33} & A_{44} & A_{55} \\
A_{21} & A_{32} & A_{43} & A_{54} & \times
\end{array}\right]
\]

The entries marked with an x in the above array are not referenced by the IMSL band subprograms.

\section*{Band Symmetric Storage Mode}

A band symmetric matrix is a band matrix that is also symmetric. The band symmetric storage mode is similar to the band mode except only the lower or upper codiagonals are stored.

In the band symmetric storage mode, the NCODA upper codiagonals are stored in the rows of a FORTRAN array of dimension \((\mathrm{NCODA}+1) \times N\). The elements are stored in the same column of the array as they are in the matrix. Specifically, values \(A_{i j}, j \leq i\) inside the band are stored in array positions ( \(i-j+\mathrm{NCODA}+1, j\) ). This is the storage mode designated by using the CHARACTER* 1 flag UPLO \(=\) ' U ' in Level 2 BLAS that compute with band symmetric matrices, see Chapter 9, Basic Matrix/Vector Operations. Alternatively, \(A_{i j}, j \leq i\), inside the band, are stored in array positions ( \(i-j+1, j\) ). This is the storage mode designated by using the CHARACTER*1 flag UPLO = ' L ' in these Level 2 BLAS, see Chapter 9, Basic Matrix/Vector Operations. The array is declared by the following statement:
DIMENSION A (LDA,N)
The parameter LDA is called the leading dimension of \(A\). It must be at least as large as NCODA +1 . The data type of a band symmetric array can be REAL or DOUBLE PRECISION.
For example, consider a real \(5 \times 5\) band matrix with 2 codiagonals. Its FORTRAN declaration is
```

PARAMETER (N=5, NCODA=2)
REAL A (NCODA+1,N)

```

The matrix \(A\) has the form
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{12} & A_{13} & 0 & 0 \\
A_{12} & A_{22} & A_{23} & A_{24} & 0 \\
A_{13} & A_{23} & A_{33} & A_{34} & A_{35} \\
0 & A_{24} & A_{34} & A_{44} & A_{45} \\
0 & 0 & A_{35} & A_{45} & A_{55}
\end{array}\right]
\]

Since \(A\) is symmetric, the values \(A_{i j}=A_{j i}\). In the FORTRAN array, it is
\[
A=\left[\begin{array}{ccccc}
\times & \times & A_{13} & A_{24} & A_{35} \\
\times & A_{12} & A_{23} & A_{34} & A_{45} \\
A_{11} & A_{22} & A_{33} & A_{44} & A_{55}
\end{array}\right]
\]

The entries marked with an \(\times\) in the above array are not referenced by the IMSL band symmetric subprograms.

An alternate storage mode for band symmetric matrices is designated using the CHARACTER* 1 flag UPLO \(=\) ' \(L^{\prime}\) in Level 2 BLAS that compute with band symmetric matrices, see Chapter 9, Basic Matrix/Vector Operations. In that case, the example matrix is represented as
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{22} & A_{33} & A_{44} & A_{55} \\
A_{12} & A_{23} & A_{34} & A_{45} & \times \\
A_{13} & A_{24} & A_{35} & \times & \times
\end{array}\right]
\]

\section*{Band Hermitian Storage Mode}

A band Hermitian matrix is a band matrix that is also Hermitian. The band Hermitian mode is a complex analogue of the band symmetric mode.
In the band Hermitian storage mode, the NCODA upper codiagonals are stored in the rows of a FORTRAN array of dimension (NCODA +1 ) \(\times N\). The elements are stored in the same column of the array as they are in the matrix. In the Level 2 BLAS, see Chapter 9, Basic Matrix/Vector Operations, this is denoted by using the CHARACTER* 1 flag UPLO \(=^{\prime} \mathrm{U}^{\prime}\). The array is declared by the following statement:
```

DIMENSION A(LDA,N)

```

The parameter LDA is called the leading dimension of \(A\). It must be at least as large as (NCODA +1 ). The data type of a band Hermitian array can be COMPLEX or, if your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX.
For example, consider a complex \(5 \times 5\) band matrix with 2 codiagonals. Its FORTRAN declaration is
```

PARAMETER (N=5,NCODA = 2)
COMPLEX A(NCODA + 1,N)

```

The matrix \(A\) has the form
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{12} & A_{13} & 0 & 0 \\
\bar{A}_{12} & A_{22} & A_{23} & A_{24} & 0 \\
\bar{A}_{13} & \bar{A}_{23} & A_{33} & A_{34} & A_{35} \\
0 & \bar{A}_{24} & \bar{A}_{34} & A_{44} & A_{45} \\
0 & 0 & \bar{A}_{35} & \bar{A}_{45} & A_{55}
\end{array}\right]
\]
where the value
\[
\bar{A}_{i j}
\]
is the complex conjugate of \(A_{i j}\). This matrix represented as a FORTRAN array is
\[
A=\left[\begin{array}{ccccc}
\times & \times & A_{13} & A_{24} & A_{35} \\
\times & A_{12} & A_{23} & A_{34} & A_{45} \\
A_{11} & A_{22} & A_{33} & A_{44} & A_{55}
\end{array}\right]
\]

The entries marked with an \(\times\) in the above array are not referenced by the IMSL band Hermitian subprograms.

An alternate storage mode for band Hermitian matrices is designated using the CHARACTER* 1 flag UPLO \(=\) 'L' in Level 2 BLAS that compute with band Hermitian matrices, see Chapter 9, Basic Matrix/Vector Operations. In that case, the example matrix is represented as
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{22} & A_{33} & A_{44} & A_{55} \\
\bar{A}_{12} & \bar{A}_{23} & \bar{A}_{34} & \bar{A}_{45} & \times \\
\bar{A}_{13} & \bar{A}_{24} & \bar{A}_{35} & \times & \times
\end{array}\right]
\]

\section*{Band Triangular Storage Mode}

A band triangular matrix is a band matrix that is also triangular. In the band triangular storage mode, the NCODA codiagonals are stored in the rows of a FORTRAN array of dimension \((\mathrm{NCODA}+1) \times N\). The elements are stored in the same column of the array as they are in the matrix. For usage in the Level 2 BLAS, see Chapter 9, Programming Notes for BLAS, the CHARACTER* 1 flag DIAGNL has the same meaning as used in section "Triangular Storage Mode". The flag UPLO has the meaning analogous with its usage in the section "Banded Symmetric Storage Mode". This array is declared by the following statement:

DIMENSION A (LDA,N)
The parameter LDA is called the leading dimension of \(A\). It must be at least as large as ( \(\mathrm{NCODA}+1\) ).

For example, consider a \(5 \times 5\) band upper triangular matrix with 2 codiagonals. Its FORTRAN declaration is
```

PARAMETER (N = 5,NCODA = 2)
COMPLEX A (NCODA + 1,N)

```

The matrix \(A\) has the form
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{12} & A_{13} & 0 & 0 \\
0 & A_{22} & A_{23} & A_{24} & 0 \\
0 & 0 & A_{33} & A_{34} & A_{35} \\
0 & 0 & 0 & A_{44} & A_{45} \\
0 & 0 & 0 & 0 & A_{55}
\end{array}\right]
\]

This matrix represented as a FORTRAN array is
\[
A=\left[\begin{array}{ccccc}
\times & \times & A_{13} & A_{24} & A_{35} \\
\times & A_{12} & A_{23} & A_{34} & A_{45} \\
A_{11} & A_{22} & A_{33} & A_{44} & A_{55}
\end{array}\right]
\]

This corresponds to the CHARACTER* 1 flags DIAGNL \(={ }^{\prime} \mathrm{N}^{\prime}\) and UPLO \(={ }^{\prime} \mathrm{U}^{\prime}\). The matrix \(A^{T}\) is represented as the FORTRAN array
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{22} & A_{33} & A_{44} & A_{55} \\
A_{12} & A_{23} & A_{34} & A_{45} & \times \\
A_{13} & A_{24} & A_{35} & \times & \times
\end{array}\right]
\]

This corresponds to the CHARACTER \(\star 1\) flags DIAGNL \(={ }^{\prime} \mathrm{N}^{\prime}\) and UPLO \(=' \mathrm{~L}^{\prime}\). In both examples, the entries indicated with an \(\times\) are not referenced by IMSL subprograms.

\section*{Codiagonal Band Symmetric Storage Mode}

This is an alternate storage mode for band symmetric matrices. It is not used by any of the BLAS, see Chapter 9, Basic Matrix/Vector Operations. Storing data in a form transposed from the Band Symmetric Storage Mode maintains unit spacing between consecutive referenced array elements. This data structure is used to get good performance in the Cholesky decomposition algorithm that solves positive definite symmetric systems of linear equations \(A x=b\). The data type can be REAL or DOUBLE PRECISION. In the codiagonal band symmetric storage mode, the NCODA upper codiagonals and right-hand-side are stored in columns of this FORTRAN array. This array is declared by the following statement:
```

DIMENSION A(LDA, NCODA + 2)

```

The parameter LDA is the leading positive dimension of \(A\). It must be at least as large as \(\mathrm{n}+\mathrm{nCODA}\).

Consider a real symmetric \(5 \times 5\) matrix with 2 codiagonals
\[
A=\left[\begin{array}{ccccc}
A_{11} & A_{12} & A_{13} & 0 & 0 \\
A_{12} & A_{22} & A_{23} & A_{24} & 0 \\
A_{13} & A_{23} & A_{33} & A_{34} & A_{35} \\
0 & A_{24} & A_{34} & A_{44} & A_{45} \\
0 & 0 & A_{35} & A_{45} & A_{55}
\end{array}\right]
\]
and a right-hand-side vector
\[
b=\left[\begin{array}{l}
b_{1} \\
b_{2} \\
b_{3} \\
b_{4} \\
b_{5}
\end{array}\right]
\]

A FORTRAN declaration for the array to hold this matrix and right-hand-side vector is
```

PARAMETER ( N = 5,NCODA = 2, LDA = N + NCODA)
REAL A(LDA,NCODA + 2)

```

The matrix and right-hand-side entries are placed in the FORTRAN array \(A\) as follows:
\[
A=\left[\begin{array}{cccc}
\times & \times & \times & \times \\
\times & \times & \times & \times \\
A_{11} & \times & \times & b_{1} \\
A_{22} & A_{12} & \times & b_{2} \\
A_{33} & A_{23} & A_{13} & b_{3} \\
A_{44} & A_{34} & A_{24} & b_{4} \\
A_{55} & A_{45} & A_{35} & b_{5}
\end{array}\right]
\]

Entries marked with an \(\times\) do not need to be defined. Certain of the IMSL band symmetric subprograms will initialize and use these values during the solution process. When a solution is computed, the \(b_{i}, i=1, \ldots, 5\), are replaced by \(x_{i}, i=1, \ldots, 5\).

The nonzero \(A_{i j}, j \geq i\), are stored in array locations \(A(j+\) NCODA, \((j-i)+1)\). The right-hand-side entries \(b_{j}\) are stored in locations \(A(j+\mathrm{NCODA}, \mathrm{NCODA}+2)\). The solution entries \(x_{j}\) are returned in \(A(j+\mathrm{NCODA}, \mathrm{NCODA}+2)\).

\section*{Codiagonal Band Hermitian Storage Mode}

This is an alternate storage mode for band Hermitian matrices. It is not used by any of the BLAS, see Chapter 9, Basic Matrix/Vector Operations. In the codiagonal band Hermitian storage mode, the real and imaginary parts of the \(2 *\) NCODA +1 upper codiagonals and right-hand-side are stored in columns of a FORTRAN array. Note that there is no explicit use of the COMPLEX or the nonstandard data type DOUBLE COMPLEX data type in this storage mode.
For Hermitian complex matrices,
\[
A=U+\sqrt{-1} V
\]
where \(U\) and \(V\) are real matrices. They satisfy the conditions \(U=U^{T}\) and \(V=-V^{T}\). The right-hand-side
\[
b=c+\sqrt{-1} d
\]
where \(c\) and \(d\) are real vectors. The solution vector is denoted as
\[
x=u+\sqrt{-1} v
\]
where \(u\) and \(v\) are real. The storage is declared with the following statement
```

DIMENSION A(LDA, 2*NCODA + 3)

```

The parameter LDA is the leading positive dimension of \(A\). It must be at least as large as \(\mathrm{N}+\mathrm{NCODA}\).

The diagonal terms \(U_{j j}\) are stored in array locations \(A(j+\mathrm{NCODA}, 1)\). The diagonal \(V_{i j}\) are zero and are not stored. The nonzero \(U_{i j}, j>i\), are stored in locations \(A(j+\) NCODA, \(2 *(j-i))\).

The nonzero \(V_{i j}\) are stored in locations \(A(j+\mathrm{NCODA}, 2 \star(j-i)+1)\). The right side vector \(b\) is stored with \(c_{j}\) and \(d_{j}\) in locations \(A\left(j+\mathrm{NCODA}, 2{ }^{*} \mathrm{NCODA}+2\right)\) and \(A\left(j+\mathrm{NCODA}, 2{ }^{*} \mathrm{NCODA}+3\right)\)
respectively. The real and imaginary parts of the solution, \(u_{j}\) and \(v_{j}\), respectively overwrite \(c_{j}\) and \(d_{j}\).

Consider a complex hermitian \(5 \times 5\) matrix with 2 codiagonals
\[
A=\left[\begin{array}{ccccc}
U_{11} & U_{12} & U_{13} & 0 & 0 \\
U_{12} & U_{22} & U_{23} & U_{24} & 0 \\
U_{13} & U_{23} & U_{33} & U_{34} & U_{35} \\
0 & U_{24} & U_{34} & U_{44} & U_{45} \\
0 & 0 & U_{35} & U_{45} & U_{55}
\end{array}\right]+\sqrt{-1}\left[\begin{array}{ccccc}
0 & V_{12} & V_{13} & 0 & 0 \\
-V_{12} & 0 & V_{23} & V_{24} & 0 \\
-V_{13} & -V_{23} & 0 & V_{34} & V_{35} \\
0 & -V_{24} & -V_{34} & 0 & V_{45} \\
0 & 0 & -V_{35} & -V_{45} & 0
\end{array}\right]
\]
and a right-hand-side vector
\[
b=\left[\begin{array}{l}
c_{1} \\
c_{2} \\
c_{3} \\
c_{4} \\
c_{5}
\end{array}\right]+\sqrt{-1}\left[\begin{array}{l}
d_{1} \\
d_{2} \\
d_{3} \\
d_{4} \\
d_{5}
\end{array}\right]
\]

A FORTRAN declaration for the array to hold this matrix and right-hand-side vector is
```

PARAMETER ( }\textrm{N}=5,NCODA = 2,LDA = N + NCODA)
REAL A(LDA, 2*NCODA + 3)

```

The matrix and right-hand-side entries are placed in the FORTRAN array \(A\) as follows:
\[
A=\left[\begin{array}{ccccccc}
\times & \times & \times & \times & \times & \times & \times \\
\times & \times & \times & \times & \times & \times & \times \\
U_{11} & \times & \times & \times & \times & c_{1} & d_{1} \\
U_{22} & U_{12} & V_{12} & \times & \times & c_{2} & d_{2} \\
U_{33} & U_{23} & V_{23} & U_{13} & V_{13} & c_{3} & d_{3} \\
U_{44} & U_{34} & V_{34} & U_{24} & V_{24} & c_{4} & d_{4} \\
U_{55} & U_{45} & V_{45} & U_{35} & V_{35} & c_{5} & d_{5}
\end{array}\right]
\]

Entries marked with an \(\times\) do not need to be defined.

\section*{Sparse Matrix Storage Mode}

The sparse linear algebraic equation solvers in Chapter 1 accept the input matrix in sparse storage mode. This structure consists of INTEGER values \(N\) and NZ, the matrix dimension and the total number of nonzero entries in the matrix. In addition, there are two INTEGER arrays IROW(*) and \(\operatorname{JCOL}(*)\) that contain unique matrix row and column coordinates where values are given. There is also an array \(A(*)\) of values. All other entries of the matrix are zero. Each of the arrays IROW \((*)\), \(\operatorname{JCOL}(*), A(*)\) must be of size NZ. The correspondence between matrix and array entries is given by
\[
A_{\operatorname{IROW}(i), \operatorname{JCOL}(i)}=A(i), i=1, \ldots, \mathrm{NZ}
\]

The data type for \(A(*)\) can be one of REAL, DOUBLE PRECISION, or COMPLEX. If your FORTRAN compiler allows, the nonstandard data type DOUBLE COMPLEX can also be declared.

For example, consider a real \(5 \times 5\) sparse matrix with 11 nonzero entries. The matrix \(A\) has the form
\[
A=\left[\begin{array}{ccccc}
A_{11} & 0 & A_{13} & A_{14} & 0 \\
A_{21} & A_{22} & 0 & 0 & 0 \\
0 & A_{32} & A_{33} & A_{34} & 0 \\
0 & 0 & A_{43} & 0 & 0 \\
0 & 0 & 0 & A_{54} & A_{55}
\end{array}\right]
\]

Declarations of arrays and definitions of the values for this sparse matrix are
```

PARAMETER (NZ = 11, N = 5)
DIMENSION IROW(NZ), JCOL (NZ), A (NZ)
DATA IROW / 1, 1,1,2,2,3,3,3,4,5,5/
DATA JCOL /1, 3, 4,1,2,2,3,4,3,4,5/
DATA A / A11, A13, A14, A21, A22, A 32, A33, A34, \& A A43, A 54, A55/

```

\section*{Reserved Names}

When writing programs accessing the MATH/LIBRARY, the user should choose FORTRAN names that do not conflict with names of IMSL subroutines, functions, or named common blocks, such as the workspace common block worksp (see Automatic Workspace Allocation). The user needs to be aware of two types of name conflicts that can arise. The first type of name conflict occurs when a name (technically a symbolic name) is not uniquely defined within a program unit (either a main program or a subprogram). For example, such a name conflict exists when the name RCURV is used to refer both to a type REAL variable and to the IMSL subroutine RCURV in a single program unit. Such errors are detected during compilation and are easy to correct. The second type of name conflict, which can be more serious, occurs when names of program units and named common blocks are not unique. For example, such a name conflict would be caused by the user defining a subroutine named WORKSP and also referencing an MATH/LIBRARY subroutine that uses the named common block wORKSP. Likewise, the user must not define a subprogram with the same name as a subprogram in the MATH/LIBRARY, that is referenced directly by the user's program or is referenced indirectly by other MATH/LIBRARY subprograms.

The MATH/LIBRARY consists of many routines, some that are described in the User's Manual and others that are not intended to be called by the user and, hence, that are not documented. If the choice of names were completely random over the set of valid FORTRAN names, and if a program uses only a small subset of the MATH/LIBRARY, the probability of name conflicts is very small. Since names are usually chosen to be mnemonic, however, the user may wish to take some precautions in choosing FORTRAN names.

Many IMSL names consist of a root name that may have a prefix to indicate the type of the routine. For example, the IMSL single precision subroutine for fitting a polynomial by least squares has the name RCURV, which is the root name, and the corresponding IMSL double
precision routine has the name DRCURV. Associated with these two routines are R2URV and DR2URV. RCURV is listed in the Alphabetical Index of Routines, but DRCURV, R2URV, and DR2URV are not. The user of RCURV must consider both names RCURV and R2URV to be reserved; likewise, the user of DRCURV must consider both names DRCURV and DR2URV to be reserved. The root names of all routines and named common blocks that are used by the MATH/LIBRARY and that do not have a numeral in the second position of the root name are listed in the Alphabetical Index of Routines. Some of the routines in this Index (such as the "Level 2 BLAS") are not intended to be called by the user and so are not documented.
The careful user can avoid any conflicts with IMSL names if the following rules are observed:
- Do not choose a name that appears in the Alphabetical Summary of Routines in the User's Manual, nor one of these names preceded by a \(\mathrm{D}_{\text {, }}, \mathrm{D}_{-}, \mathrm{C}_{-}\), or \(\mathrm{Z}_{-}\).
- Do not choose a name of three or more characters with a numeral in the second or third position.

These simplified rules include many combinations that are, in fact, allowable. However, if the user selects names that conform to these rules, no conflict will be encountered.

\section*{Deprecated Features and Renamed Routines}

\section*{Automatic Workspace Allocation}

FORTRAN subroutines that work with arrays as input and output often require extra arrays for use as workspace while doing computations or moving around data. IMSL routines generally do not require the user explicitly to allocate such arrays for use as workspace. On most systems the workspace allocation is handled transparently. The only limitation is the actual amount of memory available on the system.
On some systems the workspace is allocated out of a stack that is passed as a FORTRAN array in a named common block wORKSP. A very similar use of a workspace stack is described by Fox et al. (1978, pages 116-121). (For compatiblity with older versions of the IMSL Libraries, space is allocated from the Сомmon block, if possible.)

The arrays for workspace appear as arguments in lower-level routines. For example, the IMSL routine LSARG (in Chapter 1, "Linear Systems"), which solves systems of linear equations, needs arrays for workspace. LSARG allocates arrays from the common area, and passes them to the lower-level routine L2ARG which does the computations. In the "Comments" section of the documentation for LSARG, the amount of workspace is noted and the call to L2ARG is described. This scheme for using lower-level routines is followed throughout the IMSL Libraries. The names of these routines have a " 2 " in the second position (or in the third position in double precision routines having a " \(D\) " prefix). The user can provide workspace explicitly and call directly the "2-level" routine, which is documented along with the main routine. In a very few cases, the 2-level routine allows additional options that the main routine does not allow.
Prior to returning to the calling program, a routine that allocates workspace generally deallocates that space so that it becomes available for use in other routines.

\section*{Changing the Amount of Space Allocated}

This section is relevant only to those systems on which the transparent workspace allocator is not available.

By default, the total amount of space allocated in the common area for storage of numeric data is 5000 numeric storage units. (A numeric storage unit is the amount of space required to store an integer or a real number. By comparison, a double precision unit is twice this amount. Therefore the total amount of space allocated in the common area for storage of numeric data is 2500 double precision units.) This space is allocated as needed for INTEGER, REAL, or other numeric data. For larger problems in which the default amount of workspace is insufficient, the user can change the allocation by supplying the FORTRAN statements to define the array in the named common block and by informing the IMSL workspace allocation system of the new size of the common array. To request 7000 units, the statements are
```

COMMON /WORKSP/ RWKSP
REAL RWKSP (7000)
CALL IWKIN(7000)

```

If an IMSL routine attempts to allocate workspace in excess of the amount available in the common stack, the routine issues a fatal error message that indicates how much space is needed and prints statements like those above to guide the user in allocating the necessary amount. The program below uses IMSL routine PERMA to permute rows or columns of a matrix. This routine requires workspace equal to the number of columns, which in this example is too large. (Note that the work vector RWKSP must also provide extra space for bookkeeping.)
```

USE_PERMA_INT

```
```

Specifications for local variables
INTEGER NRA, NCA, LDA, IPERMU(6000), IPATH
REAL A(2,6000)
NRA = 2
NCA = 6000
LDA = 2
!
DO 10 I = 1, NCA
IPERMU(I) = NCA + 1 - I
1 0 ~ C O N T I N U E ~
IPATH = 2
CALL PERMA (A, IPERMU, A, IPATH=IPATH)
END

```

\section*{Output}
```

*** TERMINAL ERROR 10 from PERMA. Insufficient workspace for current

```
*** allocation (s). Correct by calling IWKIN from main program with
*** the three following statements: (REGARDLESS OF PRECISION)
*** COMMON /WORKSP/ RWKSP
*** REAL RWKSP (6018)
*** CALL IWKIN(6018)
*** TERMINAL ERROR 10 from PERMA. Workspace allocation was based on NCA =
*** 6000 .

In most cases, the amount of workspace is dependent on the parameters of the problem so the amount needed is known exactly. In a few cases, however, the amount of workspace is dependent on the data (for example, if it is necessary to count all of the unique values in a vector), so the IMSL routine cannot tell in advance exactly how much workspace is needed. In such cases the error message printed is an estimate of the amount of space required.

\section*{Character Workspace}

Since character arrays cannot be equivalenced with numeric arrays, a separate named common block WKSPCH is provided for character workspace. In most respects this stack is managed in the same way as the numeric stack. The default size of the character workspace is 2000 character units. (A character unit is the amount of space required to store one character.) The routine analogous to IWKIN used to change the default allocation is IWKCIN.
The routines in the following list are being deprecated in Version 2.0 of MATH/LIBRARY. A deprecated routine is one that is no longer used by anything in the library but is being included in the product for those users who may be currently referencing it in their application. However, any future versions of MATH/LIBRARY will not include these routines. If any of these routines are being called within an application, it is recommended that you change your code or retain the deprecated routine before replacing this library with the next version. Most of these routines were called by users only when they needed to set up their own workspace. Thus, the impact of these changes should be limited.
\begin{tabular}{|c|c|c|c|}
\hline CZADD & DE2LRH & DNCONF & E3CRG \\
\hline CZINI & DE2LSB & DNCONG & E4CRG \\
\hline CZMUL & DE3CRG & E2ASF & E4ESF \\
\hline CZSTO & DE3CRH & E2AHF & E5CRG \\
\hline DE2AHF & DE3LSE & E2BHF & E7CRG \\
\hline DE2ASF & DE4CRG & E2BSB & G2CCG \\
\hline DE2BHF & DE4ESF & E2BSF & G2CRG \\
\hline DE2BSB & DE5CRG & E2CCG & G2LCG \\
\hline DE2BSF & DE7CRG & E2CCH & G2LRG \\
\hline DE2CCG & DG2CCG & E2CHF & G3CCG \\
\hline DE2CCH & DG2CRG & E2CRG & G4CCG \\
\hline DE2CHF & DG2 DF & E2CRH & G5CCG \\
\hline DE2CRG & DG2IND & E2CSB & G7CRG \\
\hline DE2CRH & DG2LCG & E2EHF & NOONF \\
\hline DE2CSB & DG2LRG & E2ESB & NCONF \\
\hline DE2EHF & DG3CCG & E2FHF & NCONG \\
\hline DE2ESB & DG3DF & E2FSB & SDADD \\
\hline DE2FHF & DG4CCG & E2FSF & SDINI \\
\hline DE2FSB & DG5CCG & E2LCG & SDMUL \\
\hline DE2FSF & DG7CRG & E2LCH & SDSTO \\
\hline DE2LCG & DHOUAP & E2LHF & SHOUAP \\
\hline
\end{tabular}
\begin{tabular}{llll} 
DE2LCH & DHOUTR & E2LRG & SHOUTR \\
DE2LHF & DIVPBS & E2LRH & \\
DE2LRG & DN0ONF & E2LSB &
\end{tabular}

The following routines have been renamed due to naming conflicts with other software manufacturers.

CTIME - replaced with CPSEC
DTIME - replaced with TIMDY
PAGE - replaced with PGOPT

\section*{Appendix A: GAMS Index}

\section*{Description}

This index lists routines in MATH/LIBRARY by a tree-structured classification scheme known as GAMS Version 2.0 (Boisvert, Howe, Kahaner, and Springmann (1990). Only the GAMS classes that contain MATH/LIBRARY routines are included in the index. The page number for the documentation and the purpose of the routine appear alongside the routine name.
The first level of the full classification scheme contains the following major subject areas:
A. Arithmetic, Error Analysis
B. Number Theory
C. Elementary and Special Functions
D. Linear Algebra
E. Interpolation
F. Solution of Nonlinear Equations
G. Optimization
H. Differentiation and Integration
I. Differential and Integral Equations
J. Integral Transforms
K. Approximation
L. Statistics, Probability
M. Simulation, Stochastic Modeling
N. Data Handling
O. Symbolic Computation
P. Computational Geometry
Q. Graphics
R. Service Routines
S. Software Development Tools
Z. Other

There are seven levels in the classification scheme. Classes in the first level are identified by a capital letter as is given above. Classes in the remaining levels are identified by alternating letter-and-number combinations. A single letter ( \(\mathrm{a}-\mathrm{z}\) ) is used with the odd-numbered levels. A number (1-26) is used within the even-numbered levels.

\section*{IMSL MATH/LIBRARY}
A...........ARITHMETIC, ERROR ANALYSIS

A3.........Real
A3c.......Extended precision
DQADD Adds a double-precision scalar to the accumulator in extended precision.
DQINI Initializes an extended-precision accumulator with a double-precision scalar.
DQMUL Multiplies double-precision scalars in extended precision.
DQSTO Stores a double-precision approximation to an extendedprecision scalar.
A4.........Complex
A4c.......Extended precision
ZQADD Adds a double complex scalar to the accumulator in extended precision.
ZQINI Initializes an extended-precision complex accumulator to a double complex scalar.
ZQMUL Multiplies double complex scalars using extended precision.
ZQSTO Stores a double complex approximation to an extendedprecision complex scalar.

A6.........Change of representation
A6c.......Decomposition, construction
PRIME Decomposes an integer into its prime factors.
B ...........NUMBER THEORY
PRIME Decomposes an integer into its prime factors.
C...........ELEMENTARY AND SPECIAL FUNCTIONS

C2 .........Powers, roots, reciprocals
HYPOT Computes \(\sqrt{a^{2}+b^{2}}\) without underflow or overflow.
C19.......Other special functions
CONST Returns the value of various mathematical and physical constants.
CUNIT Converts \(X\) in units XUNITS to \(Y\) in units YUNITS.
D...........LINEAR ALGEBRA

D1.........Elementary vector and matrix operations

D1a....... Elementary vector operations
D1a1..... Set to constant
CSET Sets the components of a vector to a scalar, all complex.
ISET Sets the components of a vector to a scalar, all integer.
SSET Sets the components of a vector to a scalar, all single precision.
D1a2..... Minimum and maximum components
ICAMAX Finds the smallest index of the component of a complex vector having maximum magnitude.
ICAMIN Finds the smallest index of the component of a complex vector having minimum magnitude.
IIMAX Finds the smallest index of the maximum component of a integer vector.
IIMIN Finds the smallest index of the minimum of an integer vector.
ISAMAX Finds the smallest index of the component of a singleprecision vector having maximum absolute value.
ISAMIN Finds the smallest index of the component of a singleprecision vector having minimum absolute value.
ISMAX Finds the smallest index of the component of a singleprecision vector having maximum value.
ISMIN Finds the smallest index of the component of a singleprecision vector having minimum value.
D1a3.....Norm
D1a3a... \(L_{1}\) (sum of magnitudes)
DISL1 Computes the 1-norm distance between two points.
SASUM Sums the absolute values of the components of a singleprecision vector.
SCASUM Sums the absolute values of the real part together with the absolute values of the imaginary part of the components of a complex vector.

D1a3b... \(L_{2}\) (Euclidean norm)
DISL2 Computes the Euclidean (2-norm) distance between two points.
NORM2, CNORM2 Computes the Euclidean length of a vector or matrix, avoiding out-of-scale intermediate subexpressions.
MNORM2, CMNORM2 Computes the Euclidean length of a vector or matrix, avoiding out-of-scale intermediate subexpressions
NRM2, CNRM2 Computes the Euclidean length of a vector or matrix, avoiding out-of-scale intermediate subexpressions.
SCNRM2 Computes the Euclidean norm of a complex vector.
SNRM2 Computes the Euclidean length or \(L_{2}\) norm of a singleprecision vector.

D1a3c ... \(L_{\infty}\) (maximum magnitude)
DISLI Computes the infinity norm distance between two points.

ICAMAX Finds the smallest index of the component of a complex vector having maximum magnitude.
ISAMAX Finds the smallest index of the component of a singleprecision vector having maximum absolute value.

D1a4.....Dot product (inner product)
CDOTC Computes the complex conjugate dot product, \(\bar{x}^{T} y\).
CDOTU Computes the complex dot product \(x^{T} y\).
CZCDOT Computes the sum of a complex scalar plus a complex conjugate dot product, \(a+\bar{x}^{T} y\), using a double-precision accumulator.
CZDOTA Computes the sum of a complex scalar, a complex dot product and the double-complex accumulator, which is set to the result \(\mathrm{ACC} \leftarrow \mathrm{ACC}+a+x^{T} y\).
CZDOTC Computes the complex conjugate dot product, \(\bar{x}^{T} y\), using a double-precision accumulator.
CZDOTI Computes the sum of a complex scalar plus a complex dot product using a double-complex accumulator, which is set to the result ACC \(\leftarrow a+x^{T} y\).
CZDOTU Computes the complex dot product \(x^{T} y\) using a doubleprecision accumulator.
CZUDOT Computes the sum of a complex scalar plus a complex dot product, \(a+x^{T} y\), using a double-precision accumulator.
DSDOT Computes the single-precision dot product \(x^{T} y\) using a double precision accumulator.
SDDOTA Computes the sum of a single-precision scalar, a singleprecision dot product and the double-precision accumulator, which is set to the result \(\mathrm{ACC} \leftarrow \mathrm{ACC}+a+x^{T} y\).
SDDOTI Computes the sum of a single-precision scalar plus a singleprecision dot product using a double-precision accumulator, which is set to the result ACC \(\leftarrow a+x^{T} y\).
SDOT Computes the single-precision dot product \(x^{T} y\).
SDSDOT Computes the sum of a single-precision scalar and a single precision dot product, \(a+x^{T} y\), using a double-precision accumulator.

D1a5.....Copy or exchange (swap)
CCOPY Copies a vector \(x\) to a vector \(y\), both complex.
CSWAP Interchanges vectors \(x\) and \(y\), both complex.
ICOPY Copies a vector \(x\) to a vector \(y\), both integer.
ISWAP Interchanges vectors \(x\) and \(y\), both integer.
SCOPY Copies a vector \(x\) to a vector \(y\), both single precision.
SSWAP Interchanges vectors \(x\) and \(y\), both single precision.
D1a6.....Multiplication by scalar
CSCAL Multiplies a vector by a scalar, \(y \leftarrow a y\), both complex.

CSSCAL Multiplies a complex vector by a single-precision scalar, \(y \leftarrow a y\).
CSVCAL Multiplies a complex vector by a single-precision scalar and store the result in another complex vector, \(y \leftarrow a x\).
CVCAL Multiplies a vector by a scalar and store the result in another vector, \(y \leftarrow a x\), all complex.
SSCAL Multiplies a vector by a scalar, \(y \leftarrow a y\), both single precision.
SVCAL Multiplies a vector by a scalar and store the result in another vector, \(y \leftarrow a x\), all single precision.
D1a7..... Triad \((a x+y\) for vectors \(x, y\) and scalar \(a)\)
CAXPY Computes the scalar times a vector plus a vector, \(y \leftarrow a x+y\), all complex.
SAXPY Computes the scalar times a vector plus a vector, \(y \leftarrow a x+y\), all single precision.
D1a8..... Elementary rotation (Givens transformation) (search also class D1b10)
CSROT Applies a complex Givens plane rotation.
CSROTM Applies a complex modified Givens plane rotation.
SROT Applies a Givens plane rotation in single precision.
SROTM Applies a modified Givens plane rotation in single precision.
D1a10... Convolutions
RCONV Computes the convolution of two real vectors.
VCONC Computes the convolution of two complex vectors.
VCONR Computes the convolution of two real vectors.
D1a11...Other vector operations
CADD Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all complex.
CSUB Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all complex.
DISL1 Computes the 1-norm distance between two points.
DISL2 Computes the Euclidean (2-norm) distance between two points.
DISLI Computes the infinity norm distance between two points.
IADD Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all integer.
ISUB Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all integer.
ISUM Sums the values of an integer vector.
SADD Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all single precision.
SHPROD Computes the Hadamard product of two single-precision vectors.
SPRDCT Multiplies the components of a single-precision vector.
SSUB Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all single precision.

SSUM Sums the values of a single-precision vector.
SXYZ Computes a single-precision \(x y z\) product.
D1b. \(\qquad\) Elementary matrix operations

CGERC Computes the rank-one update of a complex general matrix:
\(A \leftarrow A+\alpha x \bar{y}^{T}\).
CGERU Computes the rank-one update of a complex general matrix:
\(A \leftarrow A+\alpha x y^{T}\).
CHER Computes the rank-one update of an Hermitian matrix:
\(A \leftarrow A+\alpha x \bar{x}^{T}\) with \(x\) complex and \(\alpha\) real.
CHER2 Computes a rank-two update of an Hermitian matrix:
\(A \leftarrow A+\alpha x \bar{y}^{T}+\bar{\alpha} y \bar{x}^{T}\).
CHER2K Computes one of the Hermitian rank \(2 k\) operations:
\(C \leftarrow \alpha A \bar{B}^{T}+\bar{\alpha} B \bar{A}^{T}+\beta C\) or \(C \leftarrow \alpha \bar{A}^{T} B+\bar{\alpha} \bar{B}^{T} A+\beta C\), where \(C\) is an \(n\) by \(n\) Hermitian matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.
CHERK Computes one of the Hermitian rank \(k\) operations:
\(C \leftarrow \alpha A \bar{A}^{T}+\beta C\) or \(C \leftarrow \alpha \bar{A}^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) Hermitian matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
CSYR2k Computes one of the symmetric rank \(2 k\) operations:
\(C \leftarrow \alpha A B^{T}+\alpha B A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} B+\alpha B^{T} A+\beta C\), where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.
CSYRK Computes one of the symmetric rank \(k\) operations:
\(C \leftarrow \alpha A A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
CTBSV Solves one of the complex triangular systems:
\(x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x\), or \(x \leftarrow\left(\bar{A}^{T}\right)^{-1} x\),
where \(A\) is a triangular matrix in band storage mode.
CTRSM Solves one of the complex matrix equations:
\(B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B, B \leftarrow \alpha B\left(A^{-1}\right)^{T}\),
\(B \leftarrow \alpha\left(\bar{A}^{T}\right)^{-1} B\), or \(B \leftarrow \alpha B\left(\bar{A}^{T}\right)^{-1}\)
where \(A\) is a triangular matrix.

CTRSV Solves one of the complex triangular systems:
\(x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x\), or \(x \leftarrow\left(\bar{A}^{T}\right)^{-1} x\),
where \(A\) is a triangular matrix.
HRRRR Computes the Hadamard product of two real rectangular matrices.
SGER Computes the rank-one update of a real general matrix:
\(A \leftarrow A+\alpha x y^{T}\).
SSYR Computes the rank-one update of a real symmetric matrix: \(A \leftarrow A+\alpha x x^{T}\).
SSYR2 Computes the rank-two update of a real symmetric matrix: \(A \leftarrow A+\alpha x y^{T}+\alpha y x^{T}\).
SSYR2K Computes one of the symmetric rank \(2 k\) operations:
\(C \leftarrow \alpha A B^{T}+\alpha B A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} B+\alpha B^{T} A+\beta C\), where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.
SSYRK Computes one of the symmetric rank \(k\) operations:
\(C \leftarrow \alpha A A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
STBSV Solves one of the triangular systems:
\(x \leftarrow A^{-1} x\) or \(x \leftarrow\left(A^{-1}\right)^{T} x\),
where \(A\) is a triangular matrix in band storage mode.
STRSM Solves one of the matrix equations:
\(B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B, \quad\) or \(B \leftarrow \alpha B\left(A^{-1}\right)^{T}\),
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
STRSV Solves one of the triangular linear systems:
\(x \leftarrow A^{-1} x\) or \(x \leftarrow\left(A^{-1}\right)^{T} x\),
where \(A\) is a triangular matrix.
D1b2.....Norm
NR1CB Computes the 1-norm of a complex band matrix in band storage mode.
NR1RB Computes the 1-norm of a real band matrix in band storage mode.
NR1RR Computes the 1-norm of a real matrix.
NR2RR Computes the Frobenius norm of a real rectangular matrix.
NRIRR Computes the infinity norm of a real matrix.
D1b3..... Transpose
TRNRR Transposes a rectangular matrix.
D1b4 Multiplication by vector
BLINF Computes the bilinear form \(x^{T} A y\).

CGBMV Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y\), or \(y \leftarrow \alpha \bar{A}^{T}+\beta y\), where \(A\) is a matrix stored in band storage mode.
CGEMV Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y\), or \(y \leftarrow \alpha \bar{A}^{T}+\beta y\),
Chbmv Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\),
where \(A\) is an Hermitian band matrix in band Hermitian storage.
CHEMV Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\),
where \(A\) is an Hermitian matrix.
Ctbmy Computes one of the matrix-vector operations:
\(x \leftarrow A x, x \leftarrow A^{T} x\), or \(x \leftarrow \bar{A}^{T} x\),
where \(A\) is a triangular matrix in band storage mode.
CTRMV Computes one of the matrix-vector operations:
\(x \leftarrow A x, x \leftarrow A^{T} x\), or \(x \leftarrow \bar{A}^{T} x\),
where \(A\) is a triangular matrix.
MUCBV Multiplies a complex band matrix in band storage mode by a complex vector.
MUCRV Multiplies a complex rectangular matrix by a complex vector.
MURBV Multiplies a real band matrix in band storage mode by a real vector.
MURRV Multiplies a real rectangular matrix by a vector.
SGBMV Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y\), or \(y \leftarrow \alpha A^{T} x+\beta y\), where \(A\) is a matrix stored in band storage mode.
SGEMV Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y\), or \(y \leftarrow \alpha A^{T} x+\beta y\),
SSBMV Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\),
where \(A\) is a symmetric matrix in band symmetric storage mode.
SSYMV Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\),
where \(A\) is a symmetric matrix.
STBMV Computes one of the matrix-vector operations:
\(x \leftarrow A x\) or \(x \leftarrow A^{T} x\),
where \(A\) is a triangular matrix in band storage mode.
STRMV Computes one of the matrix-vector operations:
\(x \leftarrow A x\) or \(x \leftarrow A^{T} x\),
where \(A\) is a triangular matrix.
D1b5..... Addition, subtraction
\(A C B C B\) Adds two complex band matrices, both in band storage mode.
ARBRB Adds two band matrices, both in band storage mode.

D1b6..... Multiplication
CGEMM Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C, C \leftarrow \alpha A^{T} B+\beta C, C \leftarrow \alpha A B^{T}\)
\(+\beta C, C \leftarrow \alpha A^{T} B^{T}+\beta C, C \leftarrow \alpha A \bar{B}^{T}+\beta C\),
or \(C \leftarrow \alpha \bar{A}^{T} B+\beta C, C \leftarrow \alpha A^{T} \bar{B}^{T}+\beta C\),
\(C \leftarrow \alpha \bar{A}^{T} B^{T}+\beta C\), or \(C \leftarrow \alpha \bar{A}^{T} \bar{B}^{T}+\beta C\)
CHEMM Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C\) or \(C \leftarrow \alpha B A+\beta C\),
where \(A\) is an Hermitian matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
CSYMM Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C\) or \(C \leftarrow \alpha B A+\beta C\),
where \(A\) is a symmetric matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
CTRMM Computes one of the matrix-matrix operations:
\(B \leftarrow \alpha A B, B \leftarrow \alpha A^{T} B, B \leftarrow \alpha B A, B \leftarrow \alpha B A^{T}\),
\(B \leftarrow \alpha \bar{A}^{T} B\), or \(B \leftarrow \alpha B \bar{A}^{T}\)
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
MCRCR Multiplies two complex rectangular matrices, \(A B\).
MRRRR Multiplies two real rectangular matrices, \(A B\).
MXTXF Computes the transpose product of a matrix, \(A^{T} A\).
MXTYF Multiplies the transpose of matrix \(A\) by matrix \(B, A^{T} B\).
MXYTF Multiplies a matrix \(A\) by the transpose of a matrix \(B, A B^{T}\).
SGEMM Compute one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C, C \leftarrow \alpha A^{T} B+\beta C, C \leftarrow \alpha A B^{T}\)
\(+\beta C\), or \(C \leftarrow \alpha A^{T} B^{T}+\beta C\)
SSYMM Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C\) or \(C \leftarrow \alpha B A+\beta C\),
where \(A\) is a symmetric matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
STRMM Computes one of the matrix-matrix operations:
\(B \leftarrow \alpha A B, B \leftarrow \alpha A^{T} B\), or \(B \leftarrow \alpha B A, B \leftarrow \alpha B A^{T}\),
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
D1b7..... Matrix polynomial
POLRG 1207 Evaluates a real general matrix polynomial.
D1b8..... Copy

CCBCB Copies a complex band matrix stored in complex band storage mode.
CCGCG Copies a complex general matrix.
CRBRB Copies a real band matrix stored in band storage mode.
CRGRG Copies a real general matrix.
D1b9.....Storage mode conversion
CCBCG Converts a complex matrix in band storage mode to a complex matrix in full storage mode.
CCGCB Converts a complex general matrix to a matrix in complex band storage mode.
CHBCB Copies a complex Hermitian band matrix stored in band Hermitian storage mode to a complex band matrix stored in band storage mode.
CHFCG Extends a complex Hermitian matrix defined in its upper triangle to its lower triangle.
CRBCB Converts a real matrix in band storage mode to a complex matrix in band storage mode.
CRBRG Converts a real matrix in band storage mode to a real general matrix.
CRGCG Copies a real general matrix to a complex general matrix.
CRGRB Converts a real general matrix to a matrix in band storage mode.
CRRCR Copies a real rectangular matrix to a complex rectangular matrix.
CSBRB Copies a real symmetric band matrix stored in band symmetric storage mode to a real band matrix stored in band storage mode.
CSFRG Extends a real symmetric matrix defined in its upper triangle to its lower triangle.

D1b10...Elementary rotation (Givens transformation) (search also class D1a8)
SROTG Constructs a Givens plane rotation in single precision.
SROTMG Constructs a modified Givens plane rotation in single precision.
D2.........Solution of systems of linear equations (including inversion, \(L U\) and related decompositions)

D2a.......Real nonsymmetric matrices
LSLTO Solves a real Toeplitz linear system.
D2a1 .....General
LFCRG Computes the \(L U\) factorization of a real general matrix and estimate its \(L_{1}\) condition number.
LFIRG Uses iterative refinement to improve the solution of a real general system of linear equations.
LFSRG Solves a real general system of linear equations given the \(L U\) factorization of the coefficient matrix.
LFTRG Computes the \(L U\) factorization of a real general matrix.
LINRG Computes the inverse of a real general matrix.
\begin{tabular}{|c|c|}
\hline LSARG & Solves a real general system of linear equations with iterative refinement. \\
\hline LSLRG & Solves a real general system of linear equations without iterative refinement. \\
\hline LIN_SOL_GEN & Solves a general system of linear equations \(A x=b\). Using optional arguments, any of several related computations can be performed. These extra tasks include computing the \(L U\) factorization of \(A\) using partial pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), and solving \(A^{T} x=b\) or \(A x=b\) given the \(L U\) factorization of \(A\). \\
\hline \multicolumn{2}{|l|}{D2a2..... Banded} \\
\hline LFCRB & Computes the \(L U\) factorization of a real matrix in band storage mode and estimate its \(L_{1}\) condition number. \\
\hline LFIRB & Uses iterative refinement to improve the solution of a real system of linear equations in band storage mode. \\
\hline LFSRB & Solves a real system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode. \\
\hline LFTRB & Computes the \(L U\) factorization of a real matrix in band storage mode. \\
\hline LSARB & Solves a real system of linear equations in band storage mode with iterative refinement. \\
\hline LSLRB & Solves a real system of linear equations in band storage mode without iterative refinement. \\
\hline STBSV & Solves one of the triangular systems: \\
\hline & \begin{tabular}{l}
\[
x \leftarrow A^{-1} x \text { or } x \leftarrow\left(A^{-1}\right)^{T} x,
\] \\
where \(A\) is a triangular matrix in band storage mode.
\end{tabular} \\
\hline \multicolumn{2}{|l|}{D2a2a... Tridiagonal} \\
\hline LSLCR & Computes the \(L D U\) factorization of a real tridiagonal matrix \(A\) using a cyclic reduction algorithm. \\
\hline LSLTR & Solves a real tridiagonal system of linear equations. \\
\hline LIN_SOL_TRI & Solves multiple systems of linear equations \(A_{j} x_{j}=y_{j}, j=1\), \(\ldots, k\). Each matrix \(A_{j}\) is tridiagonal with the same dimension, \(n\) : The default solution method is based on \(L U\) factorization computed using cyclic reduction. An option is used to select Gaussian elimination with partial pivoting. \\
\hline TRI_SOLVE & A real, tri-diagonal, multiple system solver. Uses both cyclic reduction and Gauss elimination. Similar in function to lin_sol_tri. \\
\hline \multicolumn{2}{|l|}{D2a3..... Triangular} \\
\hline LFCRT & Estimates the condition number of a real triangular matrix. \\
\hline LINRT & Computes the inverse of a real triangular matrix. \\
\hline LSLRT & Solves a real triangular system of linear equations. \\
\hline
\end{tabular}

STRSM Solves one of the matrix equations:
\(B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B\)
or \(B \leftarrow \alpha B\left(A^{-1}\right)^{T}\),
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
STRSV Solves one of the triangular linear systems:
\[
x \leftarrow A^{-1} x \text { or } x \leftarrow\left(A^{-1}\right)^{T} x
\]
where \(A\) is a triangular matrix.
D2a4.....Sparse
LESXG Solves a sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.
LFTXG Computes the \(L U\) factorization of a real general sparse matrix.
LSLXG Solves a sparse system of linear algebraic equations by Gaussian elimination.
gmRes Uses restarted GMRES with reverse communication to generate an approximate solution of \(A x=b\).
D2b.......Real symmetric matrices
D2b1.....General
D2b1a...Indefinite
LCHRG Computes the Cholesky decomposition of a symmetric positive semidefinite matrix with optional column pivoting.
LFCSE Computes the \(U D U^{T}\) factorization of a real symmetric matrix and estimate its \(L_{1}\) condition number.
LFISF Uses iterative refinement to improve the solution of a real symmetric system of linear equations.
LFSSF Solves a real symmetric system of linear equations given the \(U D U^{T}\) factorization of the coefficient matrix.
LFTSF Computes the \(U D U^{T}\) factorization of a real symmetric matrix.
LSASF Solves a real symmetric system of linear equations with iterative refinement.
LSLSF Solves a real symmetric system of linear equations without iterative refinement.
LIN_SOL_SELF Solves a system of linear equations \(A x=b\), where \(A\) is a self-adjoint matrix. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of \(A\) using symmetric pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), or computing the solution of \(A x=b\) given the factorization of \(A\). An optional argument is provided indicating that \(A\) is positive definite so that the Cholesky decomposition can be used.

D2b1b...Positive definite
LChrg Computes the Cholesky decomposition of a symmetric positive semidefinite matrix with optional column pivoting.
LFCDS Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix and estimate its \(L_{1}\) condition number.
LFIDS Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations.
LFSDS Solves a real symmetric positive definite system of linear equations given the \(R^{T} R\) Choleksy factorization of the coefficient matrix.
LFTDS Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix.
LINDS Computes the inverse of a real symmetric positive definite matrix.
LSADS Solves a real symmetric positive definite system of linear equations with iterative refinement.
LSLDS Solves a real symmetric positive definite system of linear equations without iterative refinement.
LIN_SOL_SELF Solves a system of linear equations \(A x=b\), where \(A\) is a self-adjoint matrix. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of \(A\) using symmetric pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), or computing the solution of \(A x=b\) given the factorization of \(A\). An optional argument is provided indicating that \(A\) is positive definite so that the Cholesky decomposition can be used.

D2b2..... Positive definite banded
LFCQS Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode and estimate its \(L_{1}\) condition number.
LFDQS Computes the determinant of a real symmetric positive definite matrix given the \(R^{T} R\) Cholesky factorization of the band symmetric storage mode.
LFIQS Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations in band symmetric storage mode.
LFSQS Solves a real symmetric positive definite system of linear equations given the factorization of the coefficient matrix in band symmetric storage mode.
LFTQS Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode.

LSAQS Solves a real symmetric positive definite system of linear equations in band symmetric storage mode with iterative refinement.
LSLPB Computes the \(R^{T} D R\) Cholesky factorization of a real symmetric positive definite matrix \(A\) in codiagonal band symmetric storage mode. Solve a system \(A x=b\).
LSLQS Solves a real symmetric positive definite system of linear equations in band symmetric storage mode without iterative refinement.

D2b4.....Sparse
JCGRC
Solves a real symmetric definite linear system using the Jacobi preconditioned conjugate gradient method with reverse communication.
LFSXD Solves a real sparse symmetric positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.
LNFXD Computes the numerical Cholesky factorization of a sparse symmetrical matrix \(A\).
LSCXD Performs the symbolic Cholesky factorization for a sparse symmetric matrix using a minimum degree ordering or a userspecified ordering, and set up the data structure for the numerical Cholesky factorization.
LSLXD Solves a sparse system of symmetric positive definite linear algebraic equations by Gaussian elimination.
PCGRC Solves a real symmetric definite linear system using a preconditioned conjugate gradient method with reverse communication.

D2c.......Complex non-Hermitian matrices
LSLCC Solves a complex circulant linear system.
LSLTC Solves a complex Toeplitz linear system.
D2c1.....General
LFCCG Computes the \(L U\) factorization of a complex general matrix and estimate its \(L_{1}\) condition number.
LFICG Uses iterative refinement to improve the solution of a complex general system of linear equations.
LFSCG Solves a complex general system of linear equations given the \(L U\) factorization of the coefficient matrix.
LFTCG Computes the \(L U\) factorization of a complex general matrix.
LINCG Computes the inverse of a complex general matrix.
LSACG Solves a complex general system of linear equations with iterative refinement.
LSLCG Solves a complex general system of linear equations without iterative refinement.
LIN_SOL_GEN Solves a general system of linear equations \(A x=b\). Using optional arguments, any of several related computations can be performed. These extra tasks include computing the
\(L U\) factorization of \(A\) using partial pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), and solving \(A^{T} x=b\) or \(A x=b\) given the \(L U\) factorization of \(A\).

D2c2.....Banded
CTBSV
Solves one of the complex triangular systems:
\(x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x\), or \(x \leftarrow\left(\bar{A}^{T}\right)^{-1} x\),
where \(A\) is a triangular matrix in band storage mode.
LFCCB

LFICB

LFSCB Solves a complex system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode.
LFTCB Computes the \(L U\) factorization of a complex matrix in band storage mode.
LSACB Solves a complex system of linear equations in band storage mode with iterative refinement.
LSLCB Solves a complex system of linear equations in band storage mode without iterative refinement.
D2c2a... Tridiagonal
LSLCQ
Computes the \(L D U\) factorization of a complex tridiagonal matrix \(A\) using a cyclic reduction algorithm.
LSLTQ Solves a complex tridiagonal system of linear equations.
LIN_SOL_TRI
Solves multiple systems of linear equations \(A_{j} x_{j}=y_{j}, j=1\), \(\ldots, k\). Each matrix \(A_{j}\) is tridiagonal with the same dimension, \(n\) : The default solution method is based on \(L U\) factorization computed using cyclic reduction. An option is used to select Gaussian elimination with partial pivoting.

D2c3..... Triangular
CTRSM
Solves one of the complex matrix equations:
\(B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B, B \leftarrow \alpha B\left(A^{-1}\right)^{T}\),
\(B \leftarrow \alpha\left(\bar{A}^{T}\right)^{-1} B\), or \(B \leftarrow \alpha B\left(\bar{A}^{T}\right)^{-1}\)
where \(A\) is a traiangular matrix.
CTRSV
Solves one of the complex triangular systems:
\(x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x\), or \(x \leftarrow\left(\bar{A}^{T}\right)^{-1} x\)
where \(A\) is a triangular matrix.
LFCCT Estimates the condition number of a complex triangular matrix.
LINCT Computes the inverse of a complex triangular matrix.
LSLCT Solves a complex triangular system of linear equations.

D2c4..... Sparse
LFSZG Solves a complex sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.
LFTZG Computes the \(L U\) factorization of a complex general sparse matrix.
LSLZG Solves a complex sparse system of linear equations by Gaussian elimination.

D2d.......Complex Hermitian matrices
D2d1.....General
D2d1a...Indefinite
LFCHF Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix and estimate its \(L_{1}\) condition number.
LfDhf Computes the determinant of a complex Hermitian matrix given the \(U D U^{H}\) factorization of the matrix.
LFIHF Uses iterative refinement to improve the solution of a complex Hermitian system of linear equations.
LFSHF Solves a complex Hermitian system of linear equations given the \(U D U^{H}\) factorization of the coefficient matrix.
LFTHF Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix.
LSAHF Solves a complex Hermitian system of linear equations with iterative refinement.
LSLHF Solves a complex Hermitian system of linear equations without iterative refinement.
LIN_SOL_SELF Solves a system of linear equations \(A x=b\), where \(A\) is a self-adjoint matrix. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of \(A\) using symmetric pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), or computing the solution of \(A x=b\) given the factorization of \(A\). An optional argument is provided indicating that \(A\) is positive definite so that the Cholesky decomposition can be used.

D2d1b...Positive definite
LFCDH Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix and estimate its \(L_{1}\) condition number.
LFIDH Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations.
LFSDH Solves a complex Hermitian positive definite system of linear equations given the \(R^{H} R\) factorization of the coefficient matrix.
LFTDH Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix.

LSADH Solves a Hermitian positive definite system of linear equations with iterative refinement.
LSLDH Solves a complex Hermitian positive definite system of linear equations without iterative refinement.
LIN_SOL_SELF Solves a system of linear equations \(A x=b\), where \(A\) is a self-adjoint matrix. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of \(A\) using symmetric pivoting, representing the determinant of \(A\), computing the inverse matrix \(A^{-1}\), or computing the solution of \(A x=b\) given the factorization of \(A\). An optional argument is provided indicating that \(A\) is positive definite so that the Cholesky decomposition can be used.

D2d2..... Positive definite banded
LFCQH Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode and estimate its \(L_{1}\) condition number.
LFIQH Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations in band Hermitian storage mode.
LFSQH Solves a complex Hermitian positive definite system of linear equations given the factorization of the coefficient matrix in band Hermitian storage mode.
LFTQH Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode.
LSAQH Solves a complex Hermitian positive definite system of linear equations in band Hermitian storage mode with iterative refinement.
LSLQB Computes the \(R^{H} D R\) Cholesky factorization of a complex hermitian positive-definite matrix \(A\) in codiagonal band hermitian storage mode. Solve a system \(A x=b\).
LSLQH Solves a complex Hermitian positive definite system of linearequations in band Hermitian storage mode without iterative refinement.
D2d4..... Sparse
LFSZD Solves a complex sparse Hermitian positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.
LNFZD Computes the numerical Cholesky factorization of a sparse Hermitian matrix \(A\).
LSLZD Solves a complex sparse Hermitian positive definite system of linear equations by Gaussian elimination.

\footnotetext{
D3.........Determinants
}

D3a.......Real nonsymmetric matrices
D3a1.....General

LFDRG Computes the determinant of a real general matrix given the \(L U\) factorization of the matrix.

D3a2.....Banded
LFDRB Computes the determinant of a real matrix in band storage mode given the \(L U\) factorization of the matrix.
D3a3..... Triangular
LFDRT Computes the determinant of a real triangular matrix.
D3b.......Real symmetric matrices
D3b1.....General
D3b1a...Indefinite
LFDSF Computes the determinant of a real symmetric matrix given the \(U D U^{T}\) factorization of the matrix.
D3b1b...Positive definite
LFDDS Computes the determinant of a real symmetric positive definite matrix given the \(R^{H} R\) Cholesky factorization of the matrix.

D3c.......Complex non-Hermitian matrices
D3c1 .....General
LFDCG Computes the determinant of a complex general matrix given the \(L U\) factorization of the matrix.
D3c2.....Banded
LFDCB Computes the determinant of a complex matrix given the \(L U\) factorization of the matrix in band storage mode.

D3c3.....Triangular
LFDCT Computes the determinant of a complex triangular matrix.
D3d....... Complex Hermitian matrices
D3d1.....General
D3d1b...Positive definite
LFDDH Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\) Cholesky factorization of the matrix.

D3d2.....Positive definite banded
LFDQH Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\) Cholesky factorization in band Hermitian storage mode.
D4.........Eigenvalues, eigenvectors
D4a.......Ordinary eigenvalue problems \((A x=\lambda x)\)
D4a1 .....Real symmetric
EVASF Computes the largest or smallest eigenvalues of a real symmetric matrix.

EVBSF Computes selected eigenvalues of a real symmetric matrix.
EVCSF Computes all of the eigenvalues and eigenvectors of a real symmetric matrix.
EVESF Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix.
EVFSF Computes selected eigenvalues and eigenvectors of a real symmetric matrix.
EVLSF Computes all of the eigenvalues of a real symmetric matrix.
LIN_EIG_SELF Computes the eigenvalues of a self-adjoint matrix, \(A\). Optionally, the eigenvectors can be computed. This gives the decomposition \(A=V D V^{T}\), where \(V\) is an \(n \times n\) orthogonal matrix and \(D\) is a real diagonal matrix.

D4a2..... Real nonsymmetric
EVCRG Computes all of the eigenvalues and eigenvectors of a real matrix.
EVLRG Computes all of the eigenvalues of a real matrix.
LIN_EIG_GEN Computes the eigenvalues of an \(n \times n\) matrix, \(A\). Optionally, the eigenvectors of \(A\) or \(A^{T}\) are computed. Using the eigenvectors of \(A\) gives the decomposition \(A V=V E\), where \(V\) is an \(n \times n\) complex matrix of eigenvectors, and \(E\) is the complex diagonal matrix of eigenvalues. Other options include the reduction of \(A\) to upper triangular or Schur form, reduction to block upper triangular form with \(2 \times 2\) or unit sized diagonal block matrices, and reduction to upper Hessenberg form.
D4a3..... Complex Hermitian
EVAHF Computes the largest or smallest eigenvalues of a complex Hermitian matrix.
EVBHF Computes the eigenvalues in a given range of a complex Hermitian matrix.
EVCHF Computes all of the eigenvalues and eigenvectors of a complex Hermitian matrix.
EVEHF Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a complex Hermitian matrix.
EVFHF Computes the eigenvalues in a given range and the corresponding eigenvectors of a complex Hermitian matrix.
EVLHF Computes all of the eigenvalues of a complex Hermitian matrix.
LIN_EIG_SELF Computes the eigenvalues of a self-adjoint matrix, \(A\). Optionally, the eigenvectors can be computed. This gives the decomposition \(A=V D V^{T}\), where \(V\) is an \(n \times n\) orthogonal matrix and \(D\) is a real diagonal matrix.

D4a4..... Complex non-Hermitian

EVCCG Computes all of the eigenvalues and eigenvectors of a complex matrix.
EVLCG Computes all of the eigenvalues of a complex matrix. LIN_EIG_GEN Computes the eigenvalues of an \(n \times n\) matrix, \(A\). Optionally, the eigenvectors of \(A\) or \(A^{T}\) are computed. Using the eigenvectors of \(A\) gives the decomposition \(A V=V E\), where \(V\) is an \(n \times n\) complex matrix of eigenvectors, and \(E\) is the complex diagonal matrix of eigenvalues. Other options include the reduction of \(A\) to upper triangular or Schur form, reduction to block upper triangular form with \(2 \times 2\) or unit sized diagonal block matrices, and reduction to upper Hessenberg form.

D4a6.....Banded
EVASB Computes the largest or smallest eigenvalues of a real symmetric matrix in band symmetric storage mode.
EVBSB Computes the eigenvalues in a given interval of a real symmetric matrix stored in band symmetric storage mode.
EVCSB Computes all of the eigenvalues and eigenvectors of a real symmetric matrix in band symmetric storage mode.
EVESB Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix in band symmetric storage mode.
EVFSB Computes the eigenvalues in a given interval and the corresponding eigenvectors of a real symmetric matrix stored in band symmetric storage mode.
EVLSB Computes all of the eigenvalues of a real symmetric matrix in band symmetric storage mode.
\(\mathrm{D} 4 \mathrm{~b} . . . . .\). Generalized eigenvalue problems (e.g., \(A x=\lambda B x\) )
D4b1.....Real symmetric
GVCSP Computes all of the eigenvalues and eigenvectors of the generalized real symmetric eigenvalue problem \(A z=\lambda B z\), with \(B\) symmetric positive definite.
GVLSP Computes all of the eigenvalues of the generalized real symmetric eigenvalue problem \(A z=\lambda B z\), with \(B\) symmetric positive definite.
LIN_GEIG_GEN Computes the generalized eigenvalues of an \(n \times n\) matrix pencil, \(A v \cong \lambda B v\). Optionally, the generalized eigenvectors are computed. If either of \(A\) or \(B\) is nonsingular, there are diagonal matrices \(\alpha\) and \(\beta\) and a complex matrix \(V\) computed such that \(A V \beta=B V \alpha\).

D4b2.....Real general
GVCRG Computes all of the eigenvalues and eigenvectors of a generalized real eigensystem \(A z=\lambda B z\).
GVLRG Computes all of the eigenvalues of a generalized real eigensystem \(A z=\lambda B z\).

LIN_GEIG_GEN Computes the generalized eigenvalues of an \(n \times n\) matrix pencil, \(A v \cong \lambda B v\). Optionally, the generalized eigenvectors are computed. If either of \(A\) or \(B\) is nonsingular, there are diagonal matrices \(\alpha\) and \(\beta\) and a complex matrix \(V\) computed such that \(A V \beta=B V \alpha\).

D4b4..... Complex general
GVCCG Computes all of the eigenvalues and eigenvectors of a generalized complex eigensystem \(A z=\lambda B z\).
GVLCG Computes all of the eigenvalues of a generalized complex eigensystem \(A z=\lambda B z\).
LIN_GEIG_GEN Computes the generalized eigenvalues of an \(n \times n\) matrix pencil, \(A v \cong \lambda B v\). Optionally, the generalized eigenvectors are computed. If either of \(A\) or \(B\) is nonsingular, there are diagonal matrices \(\alpha\) and \(\beta\) and a complex matrix \(V\) computed such that \(A V \beta=B V \alpha\).
D4c....... Associated operations
BALANC, CBSLANC Balances a general matrix before computing the eigenvalue-eigenvector decomposition.
EPICG Computes the performance index for a complex eigensystem.
EPIHF Computes the performance index for a complex Hermitian eigensystem.
EPIRG Computes the performance index for a real eigensystem.
EPISB Computes the performance index for a real symmetric eigensystem in band symmetric storage mode.
EPISF Computes the performance index for a real symmetric eigensystem.
GPICG Computes the performance index for a generalized complex eigensystem \(A z=\lambda B z\).
GPIRG Computes the performance index for a generalized real eigensystem \(A z=\lambda B z\).
GPISP Computes the performance index for a generalized real symmetric eigensystem problem.
PERFECT_SHIFT Computes eigenvectors using actual eigenvalue as an explicit shift. Called by lin_eig_self.
PWK A rational QR algorithm for computing eigenvalues of real, symmetric tri-diagonal matrices. Called by lin_svd and lin_eig_self.

D4c2..... Compute eigenvalues of matrix in compact form
D4c2b...Hessenberg
EVCCH Computes all of the eigenvalues and eigenvectors of a complex upper Hessenberg matrix.
EVCRH Computes all of the eigenvalues and eigenvectors of a real upper Hessenberg matrix.
EVLCH Computes all of the eigenvalues of a complex upper Hessenberg matrix.

EVLRH Computes all of the eigenvalues of a real upper Hessenberg matrix.

D5......... \(Q R\) decomposition, Gram-Schmidt orthogonalization
LQERR Accumulates the orthogonal matrix \(Q\) from its factored form given the \(Q R\) factorization of a rectangular matrix \(A\).
\(L Q R R R \quad\) Computes the \(Q R\) decomposition, \(A P=Q R\), using Householder transformations.
LQRSL Computes the coordinate transformation, projection, and complete the solution of the least-squares problem \(A x=b\).
LSBRR Solves a linear least-squares problem with iterative refinement.
LSQRR Solves a linear least-squares problem without iterative refinement.

D6.........Singular value decomposition
LSVCR Computes the singular value decomposition of a complex matrix.
LSVRR Computes the singular value decomposition of a real matrix.
LIN_SOL_SVD Solves a rectangular least-squares system of linear equations \(A x \cong b\) using singular value decomposition, \(A=U S V^{T}\). Using optional arguments, any of several related computations can be performed. These extra tasks include computing the rank of \(A\), the orthogonal \(m \times m\) and \(n \times n\) matrices \(U\) and \(V\), and the \(m \times n\) diagonal matrix of singular values, \(S\).
LIN_SVD Computes the singular value decomposition (SVD) of a rectangular matrix, \(A\). This gives the decomposition \(A=U S V^{T}\), where \(V\) is an \(n \times n\) orthogonal matrix, \(U\) is an \(m \times m\) orthogonal matrix, and \(S\) is a real, rectangular diagonal matrix.

D7.........Update matrix decompositions
D7b.......Cholesky
LDNCH Downdates the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix after a rank-one matrix is removed.
LUPCH Updates the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix after a rank-one matrix is added.

D7c....... \(Q R\)
LUPQR Computes an updated \(Q R\) factorization after the rank-one matrix \(\alpha x y^{T}\) is added.

D9.........Singular, overdetermined or underdetermined systems of linear equations, generalized inverses

D9a.......Unconstrained
D9a1 .....Least squares \(\left(L_{2}\right)\) solution
```

BAND
ACCUMALATION Accumulatez and solves banded least-squares problem
using Householder transformations.
BAND_SOLVE Accumulatez and solves banded least-squares problem
using Householder transformations.
HOUSE_HOLDER Accumulates and solves banded least-squares problem
using Householder transformations.
$L Q R R R$ Computes the $Q R$ decomposition, $A P=Q R$, using Householder transformations.
LQRRV Computes the least-squares solution using Householder transformations applied in blocked form.
$L Q R S L$ Computes the coordinate transformation, projection, and complete the solution of the least-squares problem $A x=b$.
LSBRR Solves a linear least-squares problem with iterative refinement.
LSQRR Solves a linear least-squares problem without iterative refinement.
LIN_SOL_LSQ Solves a rectangular system of linear equations $A x \cong b$, in a least-squares sense. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of $A$ using column and row pivoting, representing the determinant of $A$, computing the generalized inverse matrix $A^{\dagger}$, or computing the least-squares solution of $A x \cong b$ or $A^{T} y \cong d$ given the factorization of $A$. An optional argument is provided for computing the following unscaled covariance matrix: $C=\left(A^{T} A\right)^{-1}$.
LIN_SOL_SVD Solves a rectangular least-squares system of linear equations $A x \cong b$ using singular value decomposition, $A=U S V$. Using optional arguments, any of several related computations can be performed. These extra tasks include computing the rank of $A$, the orthogonal $m \times m$ and $n \times n$ matrices $U$ and $V$, and the $m \times n$ diagonal matrix of singular values, $S$.
D9b....... Constrained
D9b1..... Least squares $\left(L_{2}\right)$ solution
LCLSQ Solves a linear least-squares problem with linear constraints.
D9c....... Generalized inverses
LSGRR Computes the generalized inverse of a real matrix. LIN_SOL_LSQ Solves a rectangular system of linear equations $A x \cong b$, in a least-squares sense. Using optional arguments, any of several related computations can be performed. These extra tasks include computing and saving the factorization of $A$ using column and row pivoting, representing the determinant of $A$, computing the generalized inverse

```
matrix \(A^{\dagger}\), or computing the least-squares solution of \(A x \cong b\) or \(A^{T} y \cong d\) given the factorization of \(A\). An optional argument is provided for computing the following unscaled covariance matrix: \(C=(A A)^{-1}\).
E...........INTERPOLATION

E1 ......... Univariate data (curve fitting)
E1a .......Polynomial splines (piecewise polynomials)
BSINT Computes the spline interpolant, returning the B-spline coefficients.
CSAKM Computes the Akima cubic spline interpolant.
CSCON Computes a cubic spline interpolant that is consistent with the concavity of the data.
CSDEC Computes the cubic spline interpolant with specified derivative endpoint conditions.
CSHER Computes the Hermite cubic spline interpolant.
CSIEZ Computes the cubic spline interpolant with the 'not-a-knot' condition and return values of the interpolant at specified points.
CSINT Computes the cubic spline interpolant with the 'not-a-knot' condition.
CSPER Computes the cubic spline interpolant with periodic boundary conditions.
QDVAL Evaluates a function defined on a set of points using quadratic interpolation.
SPLEZ Computes the values of a spline that either interpolates or fits user-supplied data.
SPLINE_FITTING Solves constrained least-squares fitting of one-dimensional data by B-splines.
SPIINE_SUPPORT B-spline function and derivative evaluation package.
E2 ..........Multivariate data (surface fitting)
E2a ....... Gridded
BS2 IN Computes a two-dimensional tensor-product spline interpolant, returning the tensor-product B-spline coefficients.
BS3IN Computes a three-dimensional tensor-product spline interpolant, returning the tensor-product B-spline coefficients.
QD2DR Evaluates the derivative of a function defined on a rectangular grid using quadratic interpolation.
QD2VL Evaluates a function defined on a rectangular grid using quadratic interpolation.
QD3DR Evaluates the derivative of a function defined on a rectangular three-dimensional grid using quadratic interpolation.
QD3VL Evaluates a function defined on a rectangular threedimensional grid using quadratic interpolation.

SURFACE_FITTING Solves constrained least-squares fitting of two-dimensional data by tensor products of B-splines.

E2b ....... Scattered
SURF Computes a smooth bivariate interpolant to scattered data that is locally a quintic polynomial in two variables.
SURFACE_FAIRING Constrained weighted least-squares fitting of tensor product B -splines to discrete data, with covariance matrix and constraints at points.

E3 ......... Service routines for interpolation
E3a ....... Evaluation of fitted functions, including quadrature
E3a1 .....Function evaluation
BS1GD Evaluates the derivative of a spline on a grid, given its Bspline representation.
BS2DR Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation.
BS2GD Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid.
BS2VL Evaluates a two-dimensional tensor-product spline, given its tensor-product B-spline representation.
BS3GD Evaluates the derivative of a three-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid.
BS3VL Evaluates a three-dimensional tensor-product spline, given its tensor-product B -spline representation.
BSVAL Evaluates a spline, given its B-spline representation.
CSVAL Evaluates a cubic spline.
PPVAL Evaluates a piecewise polynomial.
QDDER Evaluates the derivative of a function defined on a set of points using quadratic interpolation.

E3a2 ..... Derivative evaluation
BS1GD Evaluates the derivative of a spline on a grid, given its Bspline representation.
BS2DR Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation.
BS2GD Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid.
BS3DR Evaluates the derivative of a three-dimensional tensorproduct spline, given its tensor-product B-spline representation.
BS3GD Evaluates the derivative of a three-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid.

BSDER Evaluates the derivative of a spline, given its B-spline representation.
CS1GD Evaluates the derivative of a cubic spline on a grid.
CSDER Evaluates the derivative of a cubic spline.
PP1GD Evaluates the derivative of a piecewise polynomial on a grid.
PPDER Evaluates the derivative of a piecewise polynomial.
QDDER Evaluates the derivative of a function defined on a set of points using quadratic interpolation.
E3a3 ..... Quadrature
BS2 IG Evaluates the integral of a tensor-product spline on a rectangular domain, given its tensor-product B-spline representation.
BS3IG Evaluates the integral of a tensor-product spline in three dimensions over a three-dimensional rectangle, given its tensorproduct B-spline representation.
BSITG Evaluates the integral of a spline, given its B-spline representation.
CSITG Evaluates the integral of a cubic spline.
E3b ....... Grid or knot generation
BSNAK Computes the 'not-a-knot' spline knot sequence.
BSOPK Computes the 'optimal' spline knot sequence.
E3c .......Manipulation of basis functions (e.g., evaluation, change of basis)
BSCPP Converts a spline in B-spline representation to piecewise polynomial representation.

F ........... SOLUTION OF NONLINEAR EQUATIONS
F1 ......... Single equation
F1a........Polynomial
F1a1......Real coefficients
ZPLRC Finds the zeros of a polynomial with real coefficients using Laguerre's method.
ZPORC Finds the zeros of a polynomial with real coefficients using the Jenkins-Traub three-stage algorithm.

F1a2......Complex coefficients
ZPOCC Finds the zeros of a polynomial with complex coefficients using the Jenkins-Traub three-stage algorithm.

F1b ....... Nonpolynomial
ZANLY Finds the zeros of a univariate complex function using Müller's method.
ZBREN Finds a zero of a real function that changes sign in a given interval.
ZREAL Finds the real zeros of a real function using Müller's method.

F2 ......... System of equations

NEQBF Solves a system of nonlinear equations using factored secant update with a finite-difference approximation to the Jacobian.
NEQBJ Solves a system of nonlinear equations using factored secant update with a user-supplied Jacobian.
NEQNF Solves a system of nonlinear equations using a modified Powell hybrid algorithm and a finite-difference approximation to the Jacobian.
NEQNJ Solves a system of nonlinear equations using a modified Powell hybrid algorithm with a user-supplied Jacobian.
G. \(\qquad\) OPTIMIZATION (search also classes K, L8)

G1......... Unconstrained
G1a.......Univariate
G1a1..... Smooth function
G1ala... User provides no derivatives
UVMIF Finds the minimum point of a smooth function of a single variable using only function evaluations.
G1alb... User provides first derivatives
UVMID Finds the minimum point of a smooth function of a single variable using both function evaluations and first derivative evaluations.

G1a2..... General function (no smoothness assumed)
UVMGS Finds the minimum point of a nonsmooth function of a single variable.
G1b.......Multivariate
G1b1..... Smooth function
Glbla... User provides no derivatives
UMCGF Minimizes a function of \(N\) variables using a conjugate gradient algorithm and a finite-difference gradient.
UMINF Minimizes a function of N variables using a quasi-New method and a finite-difference gradient.
UNLSF Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.

G1b1b... User provides first derivatives
UMCGG Minimizes a function of N variables using a conjugate gradient algorithm and a user-supplied gradient.
UMIDH Minimizes a function of N variables using a modified Newton method and a finite-difference Hessian.
UMING Minimizes a function of N variables using a quasi-New method and a user-supplied gradient.
UNLS J Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.

G1b1c...User provides first and second derivatives
UMIAH Minimizes a function of N variables using a modified Newton method and a user-supplied Hessian.

G1b2.....General function (no smoothness assumed)
UMPOL Minimizes a function of \(N\) variables using a direct search polytope algorithm.
G2.........Constrained
G2a.......Linear programming
G2a1 ..... Dense matrix of constraints
DLPRS Solves a linear programming problem via the revised simplex algorithm.

G2a2 .....Sparse matrix of constraints SLPRS Solves a sparse linear programming problem via the revised simplex algorithm.

G2e.......Quadratic programming
G2e1 .....Positive definite Hessian (i.e., convex problem)
QPROG Solves a quadratic programming problem subject to linear equality/inequality constraints.

G2h.......General nonlinear programming
G2h1.....Simple bounds
G2h1a...Smooth function
G2h1a1.User provides no derivatives
BCLSF Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.
BCONF Minimizes a function of \(N\) variables subject to bounds the variables using a quasi-Newton method and a finitedifference gradient.

G2h1a2.User provides first derivatives
BCLSJ Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.
BCODH Minimizes a function of N variables subject to bounds the variables using a modified Newton method and a finitedifference Hessian.
BCONG Minimizes a function of \(N\) variables subject to bounds the variables using a quasi-Newton method and a usersupplied gradient.

G2h1a3.User provides first and second derivatives
BCOAH Minimizes a function of N variables subject to bounds the variables using a modified Newton method and a usersupplied Hessian.

G2h1b... General function (no smoothness assumed)
BCPOL Minimizes a function of N variables subject to bounds the variables using a direct search complex algorithm.

G2h2.....Linear equality or inequality constraints
G2h2a... Smooth function
G2h2a1.User provides no derivatives
LCONF Minimizes a general objective function subject to linear equality/inequality constraints.

G2h2a2. User provides first derivatives
LCONG Minimizes a general objective function subject to linear equality/inequality constraints.
G2h3..... Nonlinear constraints
G2h3b... Equality and inequality constraints
NNLPG Uses a sequential equality constrained QP method.
NNLPF Uses a sequential equality constrained QP method.
G2h3b1. Smooth function and constraints
G2h3b1a. User provides no derivatives
G2h3b1b User provides first derivatives of function and constraints

G4......... Service routines
G4c....... Check user-supplied derivatives
CHGRD Checks a user-supplied gradient of a function.
CHHES Checks a user-supplied Hessian of an analytic function.
CHJAC Checks a user-supplied Jacobian of a system of equations with \(M\) functions in \(N\) unknowns.

G4d....... Find feasible point
GGUES Generates points in an N-dimensional space.
G4f ....... Other
CDGRD Approximates the gradient using central differences.
FDGRD Approximates the gradient using forward differences.
FDHES Approximates the Hessian using forward differences and function values.
FDJAC Approximates the Jacobian of M functions in N unknowns using forward differences.
gDHES Approximates the Hessian using forward differences and a user-supplied gradient.
H........... DIFFERENTIATION, INTEGRATION

H1..........Numerical differentiation
DERIV Computes the first, second or third derivative of a usersupplied function.

H2......... Quadrature (numerical evaluation of definite integrals)

H2a.......One-dimensional integrals
H2al .....Finite interval (general integrand)
H2a1a...Integrand available via user-defined procedure
H2ala1. Automatic (user need only specify required accuracy)
QDAG Integrates a function using a globally adaptive scheme based on Gauss-Kronrod rules.
QDAGS Integrates a function (which may have endpoint singularities).
QDNG Integrates a smooth function using a nonadaptive rule.
H2a2 .....Finite interval (specific or special type integrand including weight functions, oscillating and singular integrands, principal value integrals, splines, etc.)
\(\mathrm{H} 2 \mathrm{a} 2 \mathrm{a} . .\). Integrand available via user-defined procedure
H2a2al .Automatic (user need only specify required accuracy)
QDAGP Integrates a function with singularity points given.
QDAWC Integrates a function \(F(X) /(X-C)\) in the Cauchy principal value sense.
QDAWO Integrates a function containing a sine or a cosine.
QDAWS Integrates a function with algebraic-logarithmic singularities.

H 2 a 2 b ...Integrand available only on grid
H2a2b1.Automatic (user need only specify required accuracy)
BSITG Evaluates the integral of a spline, given its B-spline representation.

H 2 a 3 ..... Semi-infinite interval (including \(e-^{x}\) weight function)
H2a3a. ..Integrand available via user-defined procedure
H2a3a1. Automatic (user need only specify required accuracy)
QDAGI Integrates a function over an infinite or semi-infinite interval.
QDAWF Computes a Fourier integral.
H2b.......Multidimensional integrals
H2b1..... One or more hyper-rectangular regions (including iterated integrals) QMC Integrates a function over a hyperrectangle using a quasi-Monte Carlo method.

H2bla...Integrand available via user-defined procedure
H2bla1.Automatic (user need only specify required accuracy)
QAND Integrates a function on a hyper-rectangle.
TWODQ Computes a two-dimensional iterated integral.
H2b1b...Integrand available only on grid
H2b1b2.Nonautomatic

BS2IG Evaluates the integral of a tensor-product spline on a rectangular domain, given its tensor-product B-spline representation.
BS3IG Evaluates the integral of a tensor-product spline in three dimensions over a three-dimensional rectangle, given its tensorproduct B-spline representation.
\(\mathrm{H} 2 \mathrm{c} . . . . .\). Service routines (compute weight and nodes for quadrature formulas) FQRUL Computes a Fejér quadrature rule with various classical weight functions.
GQRCF Computes a Gauss, Gauss-Radau or Gauss-Lobatto quadrature rule given the recurrence coefficients for the monic polynomials orthogonal with respect to the weight function.
GQRUL Computes a Gauss, Gauss-Radau, or Gauss-Lobatto quadrature rule with various classical weight functions.
RECCF Computes recurrence coefficients for various monic polynomials.
RECQR Computes recurrence coefficients for monic polynomials given a quadrature rule.
I.............DIFFERENTIAL AND INTEGRAL EQUATIONS

I1 .......... Ordinary differential equations (ODE's)
I1a. ....... Initial value problems
I1a1 ...... General, nonstiff or mildly stiff
Ila1a..... One-step methods (e.g., Runge-Kutta)
IVMRK Solves an initial-value problem \(y^{\prime}=f(t, y)\) for ordinary differential equations using Runge-Kutta pairs of various orders.
IVPRK Solves an initial-value problem for ordinary differential equations using the Runge-Kutta-Verner fifth-order and sixth-order method.

Ilalb. ... Multistep methods (e.g., Adams predictor-corrector)
IVPAG Solves an initial-value problem for ordinary differential equations using either Adams-Moulton's or Gear's BDF method.

I1a2 ...... Stiff and mixed algebraic-differential equations
DASPG Solves a first order differential-algebraic system of equations, \(g\left(t, y, y^{\prime}\right)=0\), using Petzold-Gear BDF method.

I1b ........ Multipoint boundary value problems
I1b2 ...... Nonlinear
BVPFD Solves a (parameterized) system of differential equations with boundary conditions at two points, using a variable order, variable step size finite-difference method with deferred corrections.

BVPMS Solves a (parameterized) system of differential equations with boundary conditions at two points, using a multipleshooting method.

I1b3 ......Eigenvalue (e.g., Sturm-Liouville)
SLCNT Calculates the indices of eigenvalues of a Sturm-Liouville problem with boundary conditions (at regular points) in a specified subinterval of the real line, \([\alpha, \beta]\).
SLEIG Determines eigenvalues, eigenfunctions and/or spectral density functions for Sturm-Liouville problems in the form with boundary conditions (at regular points).

I2 ..........Partial differential equations
I2a. .......Initial boundary value problems
I2a1 ......Parabolic
PDE_1D_MG Integrates an initial-value PDE problem with one space variable.

I2a1a..... One spatial dimension
MOLCH Solves a system of partial differential equations of the form \(u_{t}=f\left(x, t, u, u_{x}, u_{x x}\right)\) using the method of lines. The solution is represented with cubic Hermite polynomials.

I2b ........Elliptic boundary value problems
I2b1 ......Linear
I2b1a. ...Second order
I2b1a1 ..Poisson (Laplace) or Helmholtz equation
I2b1ala.Rectangular domain (or topologically rectangular in the coordinate system)

FPS2H Solves Poisson's or Helmholtz's equation on a twodimensional rectangle using a fast Poisson solver based on the HODIE finite-difference scheme on a uni mesh.
fPS3H Solves Poisson's or Helmholtz's equation on a threedimensional box using a fast Poisson solver based on the HODIE finite-difference scheme on a uniform mesh.
J............INTEGRAL TRANSFORMS

J1 ..........Trigonometric transforms including fast Fourier transforms
J1a ........One-dimensional
J1a1 ......Real
FFTRB Computes the real periodic sequence from its Fourier coefficients.
FFTRF Computes the Fourier coefficients of a real periodic sequence.
FFTRI Computes parameters needed by FFTRF and FFTRB.
J1a2 ......Complex

FAST-DFT Computes the Discrete Fourier Transform (DFT) of a rank1 complex array, \(x\).
FFTCB Computes the complex periodic sequence from its Fourier coefficients.
FFTCF Computes the Fourier coefficients of a complex periodic sequence.
FFTCI Computes parameters needed by FFTCF and FFTCB.
J1a3 ...... Sine and cosine transforms
FCOSI Computes parameters needed by FCOST.
FCOST Computes the discrete Fourier cosine transformation of an even sequence.
FSINI Computes parameters needed by FSINT.
FSINT Computes the discrete Fourier sine transformation of an odd sequence.
QCOSB Computes a sequence from its cosine Fourier coefficients with only odd wave numbers.
QCOSF Computes the coefficients of the cosine Fourier transform with only odd wave numbers.
QCOSI Computes parameters needed by QCOSF and QCOSB.
QSINB Computes a sequence from its sine Fourier coefficients with only odd wave numbers.
QSINF Computes the coefficients of the sine Fourier transform with only odd wave numbers.
QSINI Computes parameters needed by QSINF and QSINB.
J1b........Multidimensional
FFT2B Computes the inverse Fourier transform of a complex periodic two-dimensional array.
FFT2D Computes Fourier coefficients of a complex periodic twodimensional array.
FFT3B Computes the inverse Fourier transform of a complex periodic three-dimensional array.
FFT3F Computes Fourier coefficients of a complex periodic threedimensional array.
FAST_2DFT Computes the Discrete Fourier Transform (DFT) of a rank2 complex array, \(x\).
FAST_3DFT Computes the Discrete Fourier Transform (DFT) of a rank3 complex array, \(x\).

J2.......... Convolutions
CCONV Computes the convolution of two complex vectors.
RCONV Computes the convolution of two real vectors.
J3 .......... Laplace transforms
INLAP Computes the inverse Laplace transform of a complex function.
SINLP Computes the inverse Laplace transform of a complex function.
K........... APPROXIMATION (search also class L8)
\(\qquad\) .Least squares \(\left(L_{2}\right)\) approximation

K1a.......Linear least squares (search also classes D5, D6, D9)
K1a1 .....Unconstrained
K1a1a...Univariate data (curve fitting)
K1ala1 .Polynomial splines (piecewise polynomials)
BSLSQ Computes the least-squares spline approximation, and return the B-spline coefficients.
BSVLS Computes the variable knot B-spline least squares approximation to given data.
CONFT Computes the least-squares constrained spline approximation, returning the B -spline coefficients. FRENCH_CURVE Constrained weighted least-squares fitting of B-splines to discrete data, with covariance matrix.and constraints at points.
K1ala2.Polynomials
RCURV Fits a polynomial curve using least squares.
K1a1a3 .Other functions (e.g., trigonometric, user-specified)
FNLSQ Compute a least-squares approximation with user-supplied basis functions.

K1alb...Multivariate data (surface fitting)
BSLS2 Computes a two-dimensional tensor-product spline approximant using least squares, returning the tensorproduct \(B\)-spline coefficients.
BSLS3 Computes a three-dimensional tensor-product spline approximant using least squares, returning the tensorproduct B-spline coefficients.
SURFACE_FAIRING Constrained weighted least-squares fitting of tensor product B -splines to discrete data, with covariance matrix and constraints at points.
K1a2.....Constrained
LIN_SOL_LSQ_CON Routine for constrained linear-least squares based on a least-distance, dual algorithm.
LIN_SOL_LSQ_INQ Routine for constrained linear-least squares based on a least-distance, dual algorithm.
LEAST_PROJ_
DISTANCE - Routine for constrained linear-least squares based on a least-distance, dual algorithm.

PARALLEL_\&
NONONEGATIVE_LSQ Solves multiple systems of linear equations
\(A_{j} x_{j}=y_{j}, j=1, \ldots, k\). Each matrix \(A_{j}\) is tridiagonal with the same dimension, \(n\) : The default solution method is based on \(L U\) factorization computed using cyclic reduction. An option is used to select Gaussian elimination with partial pivoting.
```

PARALLEL_\& BOUNDED_LSQ
Parallel routines for simple bounded constrained linear-
least squares based on a descent algorithm.

```
K1a2a... Linear constraints
            LCLSQ Solves a linear least-squares problem with linear
                        constraints.
PARALLEL_
NONNEGATIVE_LSQ Solves a large least-squares system with non-negative
                constraints, using parallel computing.
PARALLEL
BOUNDED_LSQ Solves a large least-squares system with simple bounds,
                        using parallel computing.

K1b.......Nonlinear least squares
K1b1..... Unconstrained
K1b1a... Smooth functions
K1b1a1. User provides no derivatives
UNLSF Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.

K1b1a2. User provides first derivatives
UNLSJ Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.

K1b2..... Constrained
K1b2a... Linear constraints
BCLSF Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.
BCLSJ Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.
BCNLS Solves a nonlinear least-squares problem subject to bounds on the variables and general linear constraints.
K2......... Minimax \(\left(L_{\infty}\right)\) approximation
RATCH Computes a rational weighted Chebyshev approximation to a continuous function on an interval.

K5......... Smoothing
CSSCV Computes a smooth cubic spline approximation to noisy data using cross-validation to estimate the smoothing parameter.
CSSED Smooths one-dimensional data by error detection.
CSSMH Computes a smooth cubic spline approximation to noisy data.
K6......... Service routines for approximation

K6a.......Evaluation of fitted functions, including quadrature
K6a1.....Function evaluation
BSVAL Evaluates a spline, given its B-spline representation.
CSVAL Evaluates a cubic spline.
PPVAL Evaluates a piecewise polynomial.
K6a2 ..... Derivative evaluation
BSDER Evaluates the derivative of a spline, given its B-spline representation.
CS1GD Evaluates the derivative of a cubic spline on a grid.
CSDER Evaluates the derivative of a cubic spline.
PP1GD Evaluates the derivative of a piecewise polynomial on a grid.
PPDER Evaluates the derivative of a piecewise polynomial.
K6a3.....Quadrature
CSITG Evaluates the integral of a cubic spline.
PPITG Evaluates the integral of a piecewise polynomial.
K6c.......Manipulation of basis functions (e.g., evaluation, change of basis)
BSCPP Converts a spline in B-spline representation to piecewise polynomial representation.

L ...........STATISTICS, PROBABILITY
L1 ......... Data summarization
L1c. ......Multi-dimensional data
L1c1 .....Raw data
L1c1b. ..Covariance, correlation
CCORL Computes the correlation of two complex vectors.
RCORL Computes the correlation of two real vectors.
L3 .........Elementary statistical graphics (search also class Q)
L3e. ......Multi-dimensional data
L3e3. ....Scatter diagrams
L3e3a....Superimposed \(Y\) vs. \(X\)
PLOTP Prints a plot of up to 10 sets of points.
L6 ......... Random number generation
L6a. ......Univariate
RAND_GEN Generates a rank-1 array of random numbers. The output array entries are positive and less than 1 in value.

L6a21 ...Uniform (continuous, discrete), uniform order statistics
RNUN \(\quad\) Generates pseudorandom numbers from a uniform \((0,1)\) distribution.
RNUNF Generates a pseudorandom number from a uniform \((0,1)\) distribution.

L6b \(\qquad\) Mulitivariate

L6b21 ... Linear L-1 (least absolute value) approximation random numbers FAURE_INIT Shuffles Faure sequence initialization.
FAURE_FREE Frees the structure containing information about the Faure sequence.
FAURE_NEXT Computes a shuffled Faure sequence.
L6c. ...... Service routines (e.g., seed)
RNGET Retrieves the current value of the seed used in the IMSL random number generators.
RNOPT Selects the uniform \((0,1)\) multiplicative congruential pseudorandom number generator.
RNSET Initializes a random seed for use in the IMSL random number generators.
RAND_GEN Generates a rank-1 array of random numbers. The output array entries are positive and less than 1 in value.

L8 ......... Regression (search also classes D5, D6, D9, G, K)
L8a. ...... Simple linear (e.g., \(y=\beta_{0}+\beta_{1} x+\varepsilon\) ) (search also class L8h)
L8a1. .... Ordinary least squares
FNLSQ Computes a least-squares approximation with usersupplied basis functions.

L8a1a ...Parameter estimation
L8a1a1. Unweighted data
RLINE Fits a line to a set of data points using least squares.
L8b. ...... Polynomial (e.g., \(\mathrm{y}=\beta_{0}+\beta_{1} x+\beta_{2} x 2+\varepsilon\) ) (search also class L8c)
L8b1 ..... Ordinary least squares
L8b1b ...Parameter estimation
L8b1b2. Using orthogonal polynomials
RCURV Fits a polynomial curve using least squares.
L8c ....... Multiple linear (e.g., \(y=\beta_{0}+\beta_{1} x_{1}+\ldots+\beta_{k} x_{k}+\varepsilon\) )
L8c1 ..... Ordinary least squares
L8c1b ...Parameter estimation (search also class L8c1a)
L8c1b1 . Using raw data
LSBRR Solves a linear least-squares problem with iterative refinement.
LSQRR Solves a linear least-squares problem without iterative refinement.
N........... DATA HANDLING

N1.........Input, output
PGOPT Sets or retrieves page width and length for printing.
\begin{tabular}{|c|c|c|}
\hline & WRCRL & Prints a complex rectangular matrix with a given format and labels. \\
\hline & WRCRN & Prints a complex rectangular matrix with integer row and column labels. \\
\hline & WRIRL & Prints an integer rectangular matrix with a given format and labels. \\
\hline & WRIRN & Prints an integer rectangular matrix with integer row and column labels. \\
\hline & WROPT & Sets or retrieves an option for printing a matrix. \\
\hline & WRRRL & Prints a real rectangular matrix with a given format and labels. \\
\hline & WRRRN & Prints a real rectangular matrix with integer row and column labels. \\
\hline SCALAPACK & READ & Reads matrix data from a file and place in a twodimensional block-cyclic form on a process grid. \\
\hline SCALAPACK_ち & WRITE & Writes matrix data to a file, starting with a twodimensional block-cyclic form on a process grid. \\
\hline & SHOW & Prints rank-1 and rank-2 arrays with indexing and text. \\
\hline N3........Cha & aracter m & anipulation \\
\hline & ACHAR & Returns a character given its ASCII value. \\
\hline & CVTSI & Converts a character string containing an integer number into the corresponding integer form. \\
\hline & IACHAR & Returns the integer ASCII value of a character argument. \\
\hline & ICASE & Returns the ASCII value of a character converted to uppercase. \\
\hline & IICSR & Compares two character strings using the ASCII collating sequence but without regard to case. \\
\hline & IIDEX & Determines the position in a string at which a given character sequence begins without regard to case. \\
\hline N4........ Stor & rage man & agement (e.g., stacks, heaps, trees) \\
\hline & IWKCIN & Initializes bookkeeping locations describing the character workspace stack. \\
\hline & IWKIN & Initializes bookkeeping locations describing the workspace stack. \\
\hline ScaLAPAC & K_READ & Moves data from a file to Block-Cyclic form, for use in ScaLAPACK. \\
\hline ScaLAPACK & _WRITE & Move data from Block-Cyclic form, following use in ScaLAPACK, to a file. \\
\hline N5.........Sea & rching & \\
\hline N5b.......Inse & ertion pos & ition \\
\hline & ISRCH & Searches a sorted integer vector for a given integer and return its index. \\
\hline & SRCH & Searches a sorted vector for a given scalar and return its index. \\
\hline & SSRCH & Searches a character vector, sorted in ascending ASCII order, for a given string and return its index. \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{N5c....... On a key} \\
\hline IIDEX & Determines the position in a string at which a given character sequence begins without regard to case. \\
\hline ISRCH & Searches a sorted integer vector for a given integer and return its index. \\
\hline SRCH & Searches a sorted vector for a given scalar and return its index. \\
\hline SSRCH & Searches a character vector, sorted in ascending ASCII order, for a given string and return its index. \\
\hline \multicolumn{2}{|l|}{N6........ Sorting} \\
\hline \multicolumn{2}{|l|}{N6a...... Internal} \\
\hline \multicolumn{2}{|l|}{N6a1 ..... Passive (i.e., construct pointer array, rank)} \\
\hline \multicolumn{2}{|l|}{N6a1a...Integer} \\
\hline SVIBP & Sorts an integer array by nondecreasing absolute value and return the permutation that rearranges the array. \\
\hline SVIGP & Sorts an integer array by algebraically increasing value and return the permutation that rearranges the array. \\
\hline \multicolumn{2}{|l|}{N6a1b...Real} \\
\hline SVRBP & Sorts a real array by nondecreasing absolute value and return the permutation that rearranges the array. \\
\hline SVRGP & Sorts a real array by algebraically increasing value and return the permutation that rearranges the array. \\
\hline LIN_SOL_TRI & Sorts a rank-1 array of real numbers \(x\) so the \(y\) results are algebraically nondecreasing, \(y_{1} \leq y_{2} \leq \ldots y_{n}\). \\
\hline \multicolumn{2}{|l|}{N6a2..... Active} \\
\hline \multicolumn{2}{|l|}{N6a2a...Integer} \\
\hline SVIBN & Sorts an integer array by nondecreasing absolute value. \\
\hline SVIBP & Sorts an integer array by nondecreasing absolute value and return the permutation that rearranges the array. \\
\hline SVIGN & Sorts an integer array by algebraically increasing value. \\
\hline SVIGP & Sorts an integer array by algebraically increasing value and return the permutation that rearranges the array. \\
\hline \multicolumn{2}{|l|}{N6a2b...Real} \\
\hline SVRBN & Sorts a real array by nondecreasing absolute value. \\
\hline SVRBP & Sorts a real array by nondecreasing absolute value and return the permutation that rearranges the array. \\
\hline SVRGN & Sorts a real array by algebraically increasing value. \\
\hline SVRGP & Sorts a real array by algebraically increasing value and return the permutation that rearranges the array. \\
\hline \multicolumn{2}{|l|}{N8........Permuting} \\
\hline PERMA & Permutes the rows or columns of a matrix. \\
\hline PERMU & Rearranges the elements of an array as specified by a permutation. \\
\hline Q.......... GRAPHICS & (search also classes L3) \\
\hline
\end{tabular}

PLOTP Prints a plot of up to 10 sets of points.
R...........SERVICE ROUTINES

IDYWK Computes the day of the week for a given date.
IUMAG Sets or retrieves MATH/LIBRARY integer options.
NDAYS Computes the number of days from January 1, 1900, to the given date.
NDY IN Gives the date corresponding to the number of days since January 1, 1900.
SUMAg Sets or retrieves MATH/LIBRARY single-precision options.
TDATE Get stoday's date.
TIMDY Gets time of day.
VERML Obtains IMSL MATH/LIBRARY-related version, system and license numbers.

R1 .........Machine-dependent constants
AMACH Retrieves single-precision machine constants.
IFNAN Checks if a value is NaN (not a number).
IMACH Retrieves integer machine constants.
ISNAN Detects an IEEE NaN (not-a-number).
NAN Returns, as a scalar function, a value corresponding to the IEEE 754 Standard format of floating point (ANSI/IEEE 1985) for NaN .

UMACH Sets or retrieves input or output device unit numbers.

R3 .........Error handling
BUILD_ERROR
_STRUCTURE Fills in flags, values and update the data structure for error conditions that occur in Library routines.
Prepares the structure so that calls to routine
error_post will display the reason for the error.
R3b.......Set unit number for error messages
UMACH Sets or retrieves input or output device unit numbers.
R3c ....... Other utilities ERROR_POST Prints error messages that are generated by IMSL Library routines.
ERSET Sets error handler default print and stop actions.
IERCD Retrieves the code for an informational error.
N1RTY Retrieves an error type for the most recently called IMSL routine.
S. ..........SOFTWARE DEVELOPMENT TOOLS

S3 .........Dynamic program analysis tools
CPSEC Returns CPU time used in seconds.

\section*{Appendix B: Alphabetical Summary of Routines}

\section*{Routines}

\section*{Function/Page}

A
ACBCB see page 1497

ACHAR see page 1698
AMACH see page 1771
ARBRB see page 1495
B
BCLSF see page 1310

BCLSJ see page 1317

BCNLS see page 1324

BCOAH see page 1299

BCODH see page 1293

BCONF see page 1279

BCONG see page 1286

\section*{Purpose Statement}

Adds two complex band matrices, both in band storage mode.
Returns a character given its ASCII value.
Retrieves single-precision machine constants.
Adds two band matrices, both in band storage mode.

Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian.

Solves a nonlinear least squares problem subject to bounds on the variables using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian.
Solves a nonlinear least-squares problem subject to bounds on the variables and general linear constraints.
Minimizes a function of N variables subject to bounds the variables using a modified Newton method and a usersupplied Hessian.
Minimizes a function of N variables subject to bounds the variables using a modified Newton method and a finitedifference Hessian.

Minimizes a function of N variables subject to bounds the variables using a quasi-Newton method and a finitedifference gradient.

Minimizes a function of N variables subject to bounds the variables using a quasi-Newton method and a user-supplied
\begin{tabular}{|c|c|}
\hline BCPOL see page 1306 & Minimizes a function of N variables subject to bounds the variables using a direct search complex algorithm. \\
\hline BLINF see page 1483 & Computes the bilinear form \(x^{T} A y\). \\
\hline BS1GD see page 735 & Evaluates the derivative of a spline on a grid, given its Bspline representation. \\
\hline BS2DR see page 742 & Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation. \\
\hline BS2GD see page 746 & Evaluates the derivative of a two-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid. \\
\hline BS2IG see page 750 & Evaluates the integral of a tensor-product spline on a rectangular domain, given its tensor-product B -spline representation. \\
\hline BS2IN see page 720 & Computes a two-dimensional tensor-product spline interpolant, returning the tensor-product B-spline coefficients. \\
\hline BS2VL see page 741 & Evaluates a two-dimensional tensor-product spline, given its tensor-product B-spline representation. \\
\hline BS3DR see page 756 & Evaluates the derivative of a three-dimensional tensorproduct spline, given its tensor-product B-spline representation. \\
\hline BS3GD see page 760 & Evaluates the derivative of a three-dimensional tensorproduct spline, given its tensor-product B-spline representation on a grid. \\
\hline BS3IG see page 766 & Evaluates the integral of a tensor-product spline in three dimensions over a three-dimensional rectangle, given its tensorproduct B-spline representation. \\
\hline BS3IN see page 725 & Computes a three-dimensional tensor-product spline interpolant, returning the tensor-product B-spline coefficients. \\
\hline BS3VL see page 754 & Evaluates a three-dimensional tensor-product spline, given its tensor-product B -spline representation \\
\hline BSCPP see page 770 & Converts a spline in B-spline representation to piecewise polynomial representation. \\
\hline BSDER see page 732 & Evaluates the derivative of a spline, given its B-spline representation. \\
\hline BSINT see page 711 & Computes the spline interpolant, returning the B-spline coefficients. \\
\hline BSITG see page 738 & Evaluates the integral of a spline, given its B-spline representation. \\
\hline BSLS2 see page 833 & Computes a two-dimensional tensor-product spline approximant using least squares, returning the tensorproduct B-spline coefficients. \\
\hline BSLS3 see page 838 & Computes a three-dimensional tensor-product spline approximant using least squares, returning the tensorproduct B-spline coefficients. \\
\hline
\end{tabular}

BSLSQ see page 815

BSNAK see page 715
BSOPK see page 718
BSVAL see page 731
BSVLS see page 819

BVPFD see page 961

BVPMS see page 973

\section*{C}

CADD

CAXPY

CCBCB see page 1448

CCBCG see page 1455

CCGCB see page 1453

CCGCG see page 1445
CCONV see page 1158
CCOPY
CCORL see page 1168
CDGRD see page 1390
CDOTC
CDOTU
CGBMV

CGEMM

Computes the least-squares spline approximation, and return the B -spline coefficients.
Computes the 'not-a-knot' spline knot sequence.
Computes the 'optimal' spline knot sequence.
Evaluates a spline, given its B-spline representation.
Computes the variable knot B-spline least squares approximation to given data.
Solves a (parameterized) system of differential equations with boundary conditions at two points, using a variable order, variable step size finite-difference method with deferred corrections.

Solves a (parameterized) system of differential equations with boundary conditions at two points, using a multipleshooting method.

Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all complex.
Computes the scalar times a vector plus a vector, \(y \leftarrow a x+y\), all complex.
Copies a complex band matrix stored in complex band storage mode.

Converts a complex matrix in band storage mode to a complex matrix in full storage mode.

Converts a complex general matrix to a matrix in complex band storage mode.

Copies a complex general matrix.
Computes the convolution of two complex vectors.
Copies a vector \(x\) to a vector \(y\), both complex.
Computes the correlation of two complex vectors.
Approximates the gradient using central differences.
Computes the complex conjugate dot product, \(\bar{x}^{T} y\).
Computes the complex dot product \(x^{T} y\).
Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y\), or \(y \leftarrow \alpha \bar{A}^{T}+\beta y\), where \(A\) is a matrix stored in band storage mode.

Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C, C \leftarrow \alpha A^{T} B+\beta C, C \leftarrow \alpha A B^{T}\)
\(+\beta C, C \leftarrow \alpha A^{T} B^{T}+\beta C, C \leftarrow \alpha A \bar{B}^{T}+\beta C\),
or \(C \leftarrow \alpha \bar{A}^{T} B+\beta C, C \leftarrow \alpha A^{T} \bar{B}^{T}+\beta C\),
\(C \leftarrow \alpha \bar{A}^{T} B^{T}+\beta C\), or \(C \leftarrow \alpha \bar{A}^{T} \bar{B}^{T}+\beta C\)

CGEMV Computes one of the matrix-vector operations:
\(y \leftarrow \alpha A x+\beta y, y \leftarrow \alpha A^{T} x+\beta y\), or \(y \leftarrow \alpha \bar{A}^{T}+\beta y\),

CGERC

CGERU

CHBCB see page 1467

CHBMV

CHEMM

CHEMV

CHER

CHER2

CHER2K

CHERK

CHFCG see page 1463

CHGRD see page 1403
CHHES see page 1406
CHJAC see page 1410
CHOL see page 1574

COND see page 1577
CONFT see page 824

CONST see page 1751

CPSEC see page 1705

Computes the rank-one update of a complex general matrix:
\(A \leftarrow A+\alpha x \bar{y}^{T}\).
Computes the rank-one update of a complex general matrix:
\(A \leftarrow A+\alpha x y^{T}\).
Copies a complex Hermitian band matrix stored in band Hermitian storage mode to a complex band matrix stored in band storage mode.
Computes the matrix-vector operation \(y \leftarrow \alpha A x+\beta y\), where \(A\) is an Hermitian band matrix in band Hermitian storage.

Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C\) or \(C \leftarrow \alpha B A+\beta C\),
where \(A\) is an Hermitian matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\), where \(A\) is an Hermitian matrix.
Computes the rank-one update of an Hermitian matrix:
\(A \leftarrow A+\alpha x \bar{x}^{T}\) with \(x\) complex and \(\alpha\) real.
Computes a rank-two update of an Hermitian matrix:
\(A \leftarrow A+\alpha x \bar{y}^{T}+\bar{\alpha} y \bar{x}^{T}\).
Computes one of the Hermitian rank \(2 k\) operations:
\(C \leftarrow \alpha A \bar{B}^{T}+\bar{\alpha} B \bar{A}^{T}+\beta C\) or \(C \leftarrow \alpha \bar{A}^{T} B+\bar{\alpha} \bar{B}^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) Hermitian matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.

Computes one of the Hermitian rank \(k\) operations:
\(C \leftarrow \alpha A \bar{A}^{T}+\beta C\) or \(C \leftarrow \alpha \bar{A}^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) Hermitian matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
Extends a complex Hermitian matrix defined in its upper triangle to its lower triangle.
Checks a user-supplied gradient of a function.
Checks a user-supplied Hessian of an analytic function.
Checks a user-supplied Hessian of an analytic function.
Checks a user-supplied Jacobian of a system of equations with M functions in N unknowns.
Computes the condition number of a matrix.
Computes the condition number of a rectangular matrix, \(A\).

Computes the least-squares constrained spline approximation, returning the B -spline coefficients.
Returns the value of various mathematical and physical constants.
\begin{tabular}{|c|c|}
\hline CRBCB see page 1460 & Returns CPU time used in seconds. \\
\hline CRBRB see page 1447 & Converts a real matrix in band storage mode to a complex matrix in band storage mode. \\
\hline CRBRG see page 1452 & Copies a real band matrix stored in band storage mode. \\
\hline CRGCG see page 1457 & Converts a real matrix in band storage mode to a real general matrix. \\
\hline CRGRB see page 1450 & Copies a real general matrix to a complex general matrix. \\
\hline CRGRG see page 1444 & Converts a real general matrix to a matrix in band storage mode. \\
\hline CRRCR see page 1458 & Copies a real general matrix. \\
\hline CS1GD see page 703 & Copies a real rectangular matrix to a complex rectangular matrix. \\
\hline CSAKM see page 690 & Evaluates the derivative of a cubic spline on a grid. \\
\hline CSBRB see page 1465 & Computes the Akima cubic spline interpolant. \\
\hline CSCAL & Copies a real symmetric band matrix stored in band symmetric storage mode to a real band matrix stored in band storage mode. \\
\hline CSCON see page 692 & Multiplies a vector by a scalar, \(y \leftarrow a y\), both complex. \\
\hline CSDEC see page 682 & Computes a cubic spline interpolant that is consistent with the concavity of the data. \\
\hline CSDER see page 700 & Computes the cubic spline interpolant with specified derivative endpoint conditions. \\
\hline CSET & Evaluates the derivative of a cubic spline. \\
\hline CSFRG see page 1462 & Sets the components of a vector to a scalar, all complex. \\
\hline CSHER see page 687 & Extends a real symmetric matrix defined in its upper triangle to its lower triangle. \\
\hline CSIEZ see page 677 & Computes the cubic spline interpolant with the 'not-a-knot' condition and return values of the interpolant at specified points. \\
\hline CSINT see page 680 & Computes the cubic spline interpolant with the 'not-a-knot' condition. \\
\hline CSITG see page 706 & Evaluates the integral of a cubic spline. \\
\hline CSPER see page 696 & Computes the cubic spline interpolant with periodic boundary conditions. \\
\hline CSROT & Applies a complex Givens plane rotation. \\
\hline CSROTM & Applies a complex modified Givens plane rotation. \\
\hline CSSCAL & Multiplies a complex vector by a single-precision scalar, \(y \leftarrow a y\). \\
\hline CSSCV see page 851 & Computes a smooth cubic spline approximation to noisy data using cross-validation to estimate the smoothing parameter. \\
\hline CSSED see page 844 & Smooths one-dimensional data by error detection. \\
\hline CSSMH see page 848 & Computes a smooth cubic spline approximation to noisy data. \\
\hline CSUB & Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all complex. \\
\hline CSVAL see page 699 & Evaluates a cubic spline. \\
\hline CSVCAL & Multiplies a complex vector by a single-precision scalar and store the result in another complex vector, \(y \leftarrow a x\). \\
\hline CSWAP & Interchanges vectors \(x\) and \(y\), both complex. \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline CSYMM & \begin{tabular}{l}
Computes one of the matrix-matrix operations:
\[
C \leftarrow \alpha A B+\beta C \text { or } C \leftarrow \alpha B A+\beta C
\] \\
where \(A\) is a symmetric matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
\end{tabular} \\
\hline CSYR2K & \begin{tabular}{l}
Computes one of the symmetric rank \(2 k\) operations:
\[
C \leftarrow \alpha A B^{T}+\alpha B A^{T}+\beta C \text { or } C \leftarrow \alpha A^{T} B+\alpha B^{T} A+\beta C,
\] \\
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.
\end{tabular} \\
\hline CSYRK & \begin{tabular}{l}
Computes one of the symmetric rank \(k\) operations:
\[
C \leftarrow \alpha A A^{T}+\beta C \text { or } C \leftarrow \alpha A^{T} A+\beta C
\] \\
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
\end{tabular} \\
\hline CTBMV & \begin{tabular}{l}
Computes one of the matrix-vector operations:
\[
x \leftarrow A x, x \leftarrow A^{T} x, \text { or } x \leftarrow \bar{A}^{T} x
\] \\
where \(A\) is a triangular matrix in band storage mode.
\end{tabular} \\
\hline CTBSV & \begin{tabular}{l}
Solves one of the complex triangular systems:
\[
x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x, \text { or } x \leftarrow\left(\bar{A}^{T}\right)^{-1} x,
\] \\
where \(A\) is a triangular matrix in band storage mode.
\end{tabular} \\
\hline CTRMM & \begin{tabular}{l}
Computes one of the matrix-matrix operations:
\[
\begin{aligned}
& B \leftarrow \alpha A B, B \leftarrow \alpha A^{T} B, B \leftarrow \alpha B A, B \leftarrow \alpha B A^{T}, \\
& B \leftarrow \alpha \bar{A}^{T} B, \text { or } B \leftarrow \alpha B \bar{A}^{T}
\end{aligned}
\] \\
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
\end{tabular} \\
\hline CTRMV & Computes one of the matrix-vector operations: \(x \leftarrow A x, x \leftarrow A^{T} x\), or \(x \leftarrow \bar{A}^{T} x\), where \(A\) is a triangular matrix. \\
\hline CTRSM & \begin{tabular}{l}
Solves one of the complex matrix equations:
\[
\begin{aligned}
& B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{T} B, B \leftarrow \alpha B\left(A^{-1}\right)^{T} \\
& B \leftarrow \alpha\left(\bar{A}^{T}\right)^{-1} B \text {, or } B \leftarrow \alpha B\left(\bar{A}^{T}\right)^{-1}
\end{aligned}
\] \\
where \(A\) is a traiangular matrix.
\end{tabular} \\
\hline CTRSV & Solves one of the complex triangular systems: \(x \leftarrow A^{-1} x, x \leftarrow\left(A^{-1}\right)^{T} x\), or \(x \leftarrow\left(\bar{A}^{T}\right)^{-1} x\), where \(A\) is a triangular matrix. \\
\hline CUNIT see page 1753 & Converts X in units XUNITS to Y in units YUNITS. \\
\hline CVCAL & Multiplies a vector by a scalar and store the result in another vector, \(y \leftarrow a x\), all complex. \\
\hline CVTSI see page 1704 & Converts a character string containing an integer number into the corresponding integer form. \\
\hline CZCDOT & Computes the sum of a complex scalar plus a complex conjugate dot product, \(a+\bar{x}^{T} y\), using a double-precision accumulator. \\
\hline CZDOTA & Computes the sum of a complex scalar, a complex dot product and the double-complex accumulator, which is set to the result \(\mathrm{ACC} \leftarrow \mathrm{ACC}+a+x^{T} y\). \\
\hline
\end{tabular}
\begin{tabular}{ll} 
CZDOTC & \\
CZDOTI & \begin{tabular}{l} 
Computes the complex conjugate dot product, \(\bar{x}^{T} y\), using a \\
double-precision accumulator. \\
Computes the sum of a complex scalar plus a complex dot
\end{tabular} \\
product using a double-complex accumulator, which is set to the \\
result ACC \(\leftarrow a+x^{T} y\).
\end{tabular}

E
EIG see page 1586

EPICG see page 555

EPIHF see page 607

EPIRG see page 548
EPISB see page 589

EPISF see page 571

ERROR_POST see page 1640

ERSET see page 1765
EVAHF see page 596

EVASB see page 578

EVASF see page 561
EVBHF see page 602

EVBSB see page 584

EVBSF see page 566
EVCCG see page 552
EVCCH see page 616

EVCHF see page 593

EVCRG see page 545

EVCRH see page 611
EVCSB see page 575
EVCSF see page 559

EVEHF see page 599

EVESB see page 581

Computes the eigenvalue-eigenvector decomposition of an ordinary or generalized eigenvalue problem.
Computes the performance index for a complex eigensystem.
Computes the performance index for a complex Hermitian eigensystem.
Computes the performance index for a real eigensystem.
Computes the performance index for a real symmetric eigensystem in band symmetric storage mode.
Computes the performance index for a real symmetric eigensystem.
Prints error messages that are generated by IMSL routines using EPACK.
Sets error handler default print and stop actions.
Computes the largest or smallest eigenvalues of a complex Hermitian matrix.
Computes the largest or smallest eigenvalues of a real symmetric matrix in band symmetric storage mode.
Computes the largest or smallest eigenvalues of a real symmetric matrix.
Computes the eigenvalues in a given range of a complex Hermitian matrix.
Computes the eigenvalues in a given interval of a real symmetric matrix stored in band symmetric storage mode.
Computes selected eigenvalues of a real symmetric matrix.
Computes all of the eigenvalues and eigenvectors of a complex matrix.
Computes all of the eigenvalues and eigenvectors of a complex upper Hessenberg matrix.
Computes all of the eigenvalues and eigenvectors of a complex Hermitian matrix.
Computes all of the eigenvalues and eigenvectors of a real matrix.
Computes all of the eigenvalues and eigenvectors of a real upper Hessenberg matrix.
Computes all of the eigenvalues and eigenvectors of a real symmetric matrix in band symmetric storage mode.
Computes all of the eigenvalues and eigenvectors of a real symmetric matrix.
Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a complex Hermitian matrix.
Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix in band symmetric storage mode.
\begin{tabular}{|c|c|}
\hline EVESF see page 563 & Computes the largest or smallest eigenvalues and the corresponding eigenvectors of a real symmetric matrix. \\
\hline EVFHF see page 604 & Computes the eigenvalues in a given range and the corresponding eigenvectors of a complex Hermitian matrix. \\
\hline EVFSB see page 586 & Computes the eigenvalues in a given interval and the corresponding eigenvectors of a real symmetric matrix stored in band symmetric storage mode. \\
\hline EVFSF see page 568 & Computes selected eigenvalues and eigenvectors of a real symmetric matrix. \\
\hline EVLCG see page 550 & Computes all of the eigenvalues of a complex matrix. \\
\hline EVLCH see page 614 & Computes all of the eigenvalues of a complex upper Hessenberg matrix. \\
\hline EVLHF see page 591 & Computes all of the eigenvalues of a complex Hermitian matrix. \\
\hline EVLRG see page 543 & Computes all of the eigenvalues of a real matrix. \\
\hline EVLRH see page 609 & Computes all of the eigenvalues of a real upper Hessenberg matrix. \\
\hline EVLSB see page 573 & Computes all of the eigenvalues of a real symmetric matrix in band symmetric storage mode. \\
\hline EVLSF see page 557 & Computes all of the eigenvalues of a real symmetric matrix. \\
\hline EYE see page 1590 & Creates a rank-2 square array whose diagonals are all the value one. \\
\hline F & \\
\hline FAURE_FREE see page 1736 & Frees the structure containing information about the Faure sequence. \\
\hline FAURE_INIT see page 1736 & Shuffled Faure sequence initialization. \\
\hline FAURE_NEXT see page 1737 & Computes a shuffled Faure sequence. \\
\hline FAST_DFT see page 1086 & Computes the Discrete Fourier Transform of a rank-1 complex array, \(x\). \\
\hline FAST_2DFT see page 1093 & Computes the Discrete Fourier Transform (2DFT) of a rank-2 complex array, \(x\). \\
\hline FAST_3DFT see page 1099 & Computes the Discrete Fourier Transform (2DFT) of a rank-3 complex array, \(x\). \\
\hline FCOSI see page 1124 & Computes parameters needed by FCOST. \\
\hline FCOST see page 1122 & Computes the discrete Fourier cosine transformation of an even sequence. \\
\hline FDGRD see page 1392 & Approximates the gradient using forward differences. \\
\hline FDHES see page 1394 & Approximates the Hessian using forward differences and function values. \\
\hline FDJAC see page 1400 & Approximates the Jacobian of \(M\) functions in \(N\) unknowns using forward differences. \\
\hline FFT see page 1592 & The Discrete Fourier Transform of a complex sequence and its inverse transform. \\
\hline FFT_BOX see page 1594 & The Discrete Fourier Transform of several complex or real \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline & sequences. \\
\hline FFT2B see page 1142 & Computes the inverse Fourier transform of a complex periodic two-dimensional array. \\
\hline FFT2D see page 1139 & Computes Fourier coefficients of a complex periodic twodimensional array. \\
\hline FFT3B see page 1149 & Computes the inverse Fourier transform of a complex periodic three-dimensional array. \\
\hline FFT3F see page 1145 & Computes Fourier coefficients of a complex periodic threedimensional array. \\
\hline FFTCB see page 1113 & Computes the complex periodic sequence from its Fourier coefficients. \\
\hline FFTCF see page 1111 & Computes the Fourier coefficients of a complex periodic sequence. \\
\hline FFTCI see page 1116 & Computes parameters needed by FFTCF and FFTCB. \\
\hline FFTRB see page 1106 & Computes the real periodic sequence from its Fourier coefficients. \\
\hline FFTRF see page 1103 & Computes the Fourier coefficients of a real periodic sequence. \\
\hline FFTRI see page 1109 & Computes parameters needed by FFTRF and FFTRB. \\
\hline FNLSQ see page 811 & Computes a least-squares approximation with user-supplied basis functions. \\
\hline FPS2H see page 1053 & Solves Poisson's or Helmholtz's equation on a twodimensional rectangle using a fast Poisson solver based on the HODIE finite-difference scheme on a uni mesh. \\
\hline FPS3H see page 1059 & Solves Poisson's or Helmholtz's equation on a threedimensional box using a fast Poisson solver based on the HODIE finite-difference scheme on a uniform mesh. \\
\hline FQRUL see page 914 & Computes a Fejér quadrature rule with various classical weight functions. \\
\hline FSINI see page 1120 & Computes parameters needed by FSINT. \\
\hline FSINT see page 1118 & Computes the discrete Fourier sine transformation of an odd sequence. \\
\hline G & \\
\hline GDHES see page 1397 & Approximates the Hessian using forward differences and a user-supplied gradient. \\
\hline GGUES see page 1414 & Generates points in an N-dimensional space. \\
\hline GMRES see page 436 & Uses restarted GMRES with reverse communication to generate an approximate solution of \(A x=b\). \\
\hline GPICG see page 632 & Computes the performance index for a generalized complex eigensystem \(A z=\lambda B z\). \\
\hline GPIRG see page 625 & Computes the performance index for a generalized real eigensystem \(A z=\lambda B z\). \\
\hline GPISP see page 639 & Computes the performance index for a generalized real symmetric eigensystem problem. \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline GQRCF see page 905 & Computes a Gauss, Gauss-Radau or Gauss-Lobatto quadrature rule given the recurrence coefficients for the monic polynomials orthogonal with respect to the weight function. \\
\hline GQRUL see page 901 & Computes a Gauss, Gauss-Radau, or Gauss-Lobatto quadrature rule with various classical weight functions. \\
\hline GVCCG see page 629 & Computes all of the eigenvalues and eigenvectors of a generalized complex eigensystem \(A z=\lambda B z\). \\
\hline GVCRG see page 621 & Computes all of the eigenvalues and eigenvectors of a generalized real eigensystem \(A z=\lambda B z\). \\
\hline GVCSP see page 636 & Computes all of the eigenvalues and eigenvectors of the generalized real symmetric eigenvalue problem \(A z=\lambda B z\), with \(B\) symmetric positive definite. \\
\hline GVLCG see page 627 & Computes all of the eigenvalues of a generalized complex eigensystem \(A z=\lambda B z\). \\
\hline GVLRG see page 618 & Computes all of the eigenvalues of a generalized real eigensystem \(A z=\lambda B z\). \\
\hline GVLSP see page 634 & Computes all of the eigenvalues of the generalized real symmetric eigenvalue problem \(A z=\lambda B z\), with \(B\) symmetric positive definite. \\
\hline \multicolumn{2}{|l|}{H} \\
\hline HRRRR see page 1481 & Computes the Hadamard product of two real rectangular matrices. \\
\hline HYPOT see page 1757 & Computes \(\sqrt{a^{2}+b^{2}}\) without underflow or overflow. \\
\hline \multicolumn{2}{|l|}{I} \\
\hline IACHAR see page 1699 & Returns the integer ASCII value of a character argument. \\
\hline IADD & Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all integer. Finds the smallest index of the component of a complex vector having maximum magnitude. \\
\hline ICAMAX & Finds the smallest index of the component of a complex vector having minimum magnitude. \\
\hline ICAMIN & Returns the ASCII value of a character converted to uppercase. \\
\hline ICASE see page 1700 & Copies a vector \(x\) to a vector \(y\), both integer. \\
\hline ICOPY & Computes the day of the week for a given date. \\
\hline IDYWK see page 1711 & Retrieves the code for an informational error. \\
\hline IERCD and N1RTY see page 1766 & The inverse of the Discrete Fourier Transform of a complex sequence. \\
\hline IFFT see page 1596 & The inverse of the Discrete Fourier Transform of a complex sequence. \\
\hline IFFT_BOX see page 1598 & The inverse Discrete Fourier Transform of several complex or real sequences. \\
\hline IFNAN(X) see page 1772 & Checks if a value is NaN (not a number). \\
\hline
\end{tabular}

IICSR see page 1701
IIDEX see page 1703
IIMAX
IIMIN
IMACH see page 1769
INLAP see page 1172
ISAMAX
ISAMIN

ISET
ISMAX

ISMIN

ISNAN see page 1600

ISRCH see page 1694

ISUB

ISUM
ISWAP
IUMAG see page 1739
IVMRK see page 934

IVPAG see page 944

IVPRK see page 927

\section*{J}

JCGRC see page 433

Compares two character strings using the ASCII collating sequence but without regard to case.
Determines the position in a string at which a given character sequence begins without regard to case.
Finds the smallest index of the maximum component of a integer vector.
Finds the smallest index of the minimum of an integer vector.
Retrieves integer machine constants.
Computes the inverse Laplace transform of a complex function.
Finds the smallest index of the component of a singleprecision vector having maximum absolute value.
Finds the smallest index of the component of a singleprecision vector having minimum absolute value.
Sets the components of a vector to a scalar, all integer.
Finds the smallest index of the component of a singleprecision vector having maximum value.
Finds the smallest index of the component of a singleprecision vector having minimum value.
This is a generic logical function used to test scalars or arrays for occurrence of an IEEE 754 Standard format of floating point (ANSI/IEEE 1985) NaN, or not-a-number.
Searches a sorted integer vector for a given integer and return its index.
Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all integer.
Sums the values of an integer vector.
Interchanges vectors \(x\) and \(y\), both integer.
Sets or retrieves MATH/LIBRARY integer options.
Solves an initial-value problem \(y^{\prime}=f(t, y)\) for ordinary differential equations using Runge-Kutta pairs of various orders.
Solves an initial-value problem for ordinary differential equations using either Adams-Moulton's or Gear's BDF method.
Solves an initial-value problem for ordinary differential equations using the Runge-Kutta-Verner fifth-order and sixth-order method.

Solves a real symmetric definite linear system using the Jacobi preconditioned conjugate gradient method with reverse communication.
\begin{tabular}{|c|c|}
\hline LCHRG see page 489 & Computes the Cholesky decomposition of a symmetric positive semidefinite matrix with optional column pivoting. \\
\hline LCLSQ see page 462 & Solves a linear least-squares problem with linear constraints. \\
\hline LCONF see page 1364 & Minimizes a general objective function subject to linear equality/inequality constraints. \\
\hline LCONG see page 1370 & Minimizes a general objective function subject to linear equality/inequality constraints. \\
\hline LDNCH see page 494 & Downdates the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix after a rank-one matrix is removed \\
\hline LFCCB see page 330 & Computes the \(L U\) factorization of a complex matrix in band storage mode and estimate its \(L_{1}\) condition number. \\
\hline LFCCG see page 127 & Computes the \(L U\) factorization of a complex general matrix and estimate its \(L_{1}\) condition number. \\
\hline LFCCT see page 168 & Estimates the condition number of a complex triangular matrix. \\
\hline LFCDH see page 236 & Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix and estimate its \(L_{1}\) condition number. \\
\hline LFCDS see page 185 & Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix and estimate its \(L_{1}\) condition number. \\
\hline LFCHF see page 263 & Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix and estimate its \(L_{1}\) condition number. \\
\hline LFCQH see page 352 & Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode and estimate its \(L_{1}\) condition number. \\
\hline LFCQS see page 308 & Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode and estimate its \(L_{1}\) condition number. \\
\hline LFCRB see page 287 & Computes the \(L U\) factorization of a real matrix in band storage mode and estimate its \(L_{1}\) condition number. \\
\hline LFCRG see page 93 & Computes the \(L U\) factorization of a real general matrix and estimate its \(L_{1}\) condition number. \\
\hline LFCRT see page 157 & Estimates the condition number of a real triangular matrix. \\
\hline LFCSF see page 214 & Computes the \(U D U^{T}\) factorization of a real symmetric matrix and estimate its \(L_{1}\) condition number. \\
\hline LFDCB see page 342 & Computes the determinant of a complex matrix given the \(L U\) factorization of the matrix in band storage mode. \\
\hline LFDCG see page 148 & Computes the determinant of a complex general matrix given the \(L U\) factorization of the matrix. \\
\hline LFDCT see page 172 & Computes the determinant of a complex triangular matrix. \\
\hline
\end{tabular}

LFDDH see page 256

LFDDS see page 204

LFDHF see page 274

LFDQH see page 362

LFDQS see page 318

LFDRB see page 298

LFDRG see page 113
LFDRT see page 161
LFDSF see page 224

LFICB see page 338

LFICG see page 142
LFIDH see page 251

LFIDS see page 199

LFIHF see page 271

LFIQH see page 360

LFIQS see page 315

LFIRB see page 295

LFIRG see page 107

LFISF see page 221

LFSCB see page 336

Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\) Cholesky factorization of the matrix.
Computes the determinant of a real symmetric positive definite matrix given the \(R^{H} R\) Cholesky factorization of the matrix.
Computes the determinant of a complex Hermitian matrix given the \(U D U^{H}\) factorization of the matrix.
Computes the determinant of a complex Hermitian positive definite matrix given the \(R^{H} R\) Cholesky factorization in band Hermitian storage mode.
Computes the determinant of a real symmetric positive definite matrix given the \(R^{T} R\) Cholesky factorization of the band symmetric storage mode.
Computes the determinant of a real matrix in band storage mode given the \(L U\) factorization of the matrix.
Computes the determinant of a real general matrix given the \(L U\) factorization of the matrix.
Computes the determinant of a real triangular matrix.
Computes the determinant of a real symmetric matrix given the \(U D U^{T}\) factorization of the matrix.
Uses iterative refinement to improve the solution of a complex system of linear equations in band storage mode.
Uses iterative refinement to improve the solution of a complex general system of linear equations.
Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations.
Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations.
Uses iterative refinement to improve the solution of a complex Hermitian system of linear equations.
Uses iterative refinement to improve the solution of a complex Hermitian positive definite system of linear equations in band Hermitian storage mode.
Uses iterative refinement to improve the solution of a real symmetric positive definite system of linear equations in band symmetric storage mode.
Uses iterative refinement to improve the solution of a real system of linear equations in band storage mode.
Uses iterative refinement to improve the solution of a real general system of linear equations.
Uses iterative refinement to improve the solution of a real symmetric system of linear equations.
Solves a complex system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode.

LFSCG see page 138
LFSDH see page 246

LFSDS see page 194

LFSHF see page 269

LFSQH see page 358

LFSQS see page 313

LFSRB see page 293
LFSRG see page 103

LFSSF see page 219

LFSXD see page 404

LFSXG see page 374

LFSZD see page 417

LFSZG see page 387
LFTCB see page 333

LFTCG see page 133
LFTDH see page 241

LFTDS see page 190

LFTHF see page 266

LFTQH see page 355

LFTQS see page 311

Solves a complex general system of linear equations given the \(L U\) factorization of the coefficient matrix.
Solves a complex Hermitian positive definite system of linear equations given the \(R^{H} R\) factorization of the coefficient matrix.
Solves a real symmetric positive definite system of linear equations given the \(R^{T} R\) Choleksy factorization of the coefficient matrix.
Solves a complex Hermitian system of linear equations given the \(U D U^{H}\) factorization of the coefficient matrix.
Solves a complex Hermitian positive definite system of linear equations given the factorization of the coefficient matrix in band Hermitian storage mode.
Solves a real symmetric positive definite system of linear equations given the factorization of the coefficient matrix in band symmetric storage mode.
Solves a real system of linear equations given the \(L U\) factorization of the coefficient matrix in band storage mode.
Solves a real general system of linear equations given the \(L U\) factorization of the coefficient matrix.
Solves a real symmetric system of linear equations given the \(U D U^{T}\) factorization of the coefficient matrix.
Solves a real sparse symmetric positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.

Solves a sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.
Solves a complex sparse Hermitian positive definite system of linear equations, given the Cholesky factorization of the coefficient matrix.
Solves a complex sparse system of linear equations given the \(L U\) factorization of the coefficient matrix.
Computes the \(L U\) factorization of a complex matrix in band storage mode.
Computes the \(L U\) factorization of a complex general matrix.
Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix.
Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix.
Computes the \(U D U^{H}\) factorization of a complex Hermitian matrix.
Computes the \(R^{H} R\) factorization of a complex Hermitian positive definite matrix in band Hermitian storage mode.
Computes the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix in band symmetric storage mode.

LFTRB see page 290

LFTRG see page 98
LFTSF see page 217

LFTXG see page 369

LFTZG see page 382

LINCG see page 149
LINCT see page 174
LINDS see page 205

LINRG see page 114
LINRT see page 163
LIN_EIG_GEN see page 527

LIN_EIG_SELF see page 520
LIN_GEIG_GEN see page 535

LIN_SOL_GEN see page 9
LIN_SOL_LSQ see page 27

LIN_SOL_SELF see page 17

LIN_SOL_SVD see page 36

LIN_SOL_TRI see page 44
LIN_SVD see page 57

LNFXD see page 399

LNFZD see page 412

LQERR see page 473

LQRRR see page 466

LQRRV see page 452

LQRSL see page 478

LSACB see page 324

Computes the \(L U\) factorization of a real matrix in band storage mode.
Computes the \(L U\) factorization of a real general matrix.
Computes the \(U D U^{T}\) factorization of a real symmetric matrix.
Computes the \(L U\) factorization of a real general sparse matrix.
Computes the \(L U\) factorization of a complex general sparse matrix.
Computes the inverse of a complex general matrix.
Computes the inverse of a complex triangular matrix.
Computes the inverse of a real symmetric positive definite matrix.
Computes the inverse of a real general matrix.
Computes the inverse of a real triangular matrix.
Computes the eigenvalues of a self-adjoint
matrix, \(A\).
Computes the eigenvalues of a self-adjoint
matrix, \(A\).
Computes the generalized eigenvalues of an \(n \times n\) matrix pencil, \(A v=\lambda B v\).
Solves a general system of linear equations \(A x=b\).
Solves a rectangular system of linear equations \(A x \cong b\), in a least-squares sense.
Solves a system of linear equations \(A x=b\), where \(A\) is a self-adjoint matrix.
Solves a rectangular least-squares system of linear equations \(A x \cong b\) using singular value decomposition.
Solves multiple systems of linear equations.
Computes the singular value decomposition (SVD) of a rectangular matrix, \(A\).
Computes the numerical Cholesky factorization of a sparse symmetrical matrix \(A\).
Computes the numerical Cholesky factorization of a sparse Hermitian matrix \(A\).
Accumulates the orthogonal matrix \(Q\) from its factored form given the \(Q R\) factorization of a rectangular matrix \(A\).
Computes the \(Q R\) decomposition, \(A P=Q R\), using Householder transformations.
Computes the least-squares solution using Householder transformations applied in blocked form.
Computes the coordinate transformation, projection, and complete the solution of the least-squares problem \(A x=b\).
Solves a complex system of linear equations in band storage mode with iterative refinement.
\begin{tabular}{ll} 
LSACG see page 118 & \begin{tabular}{l} 
Solves a complex general system of linear equations with \\
iterative refinement.
\end{tabular} \\
LSADH see page 226 & \begin{tabular}{l} 
Solves a Hermitian positive definite system of linear \\
equations with iterative refinement.
\end{tabular} \\
LSADS see page 176 & \begin{tabular}{l} 
Solves a real symmetric positive definite system of linear \\
equations with iterative refinement.
\end{tabular} \\
LSAHF see page 258 & \begin{tabular}{l} 
Solves a complex Hermitian system of linear equations with \\
iterative refinement.
\end{tabular} \\
& \begin{tabular}{l} 
Solves a complex Hermitian positive definite system of \\
linear equations in band Hermitian storage mode with
\end{tabular} \\
LSAQH see page 344 & \begin{tabular}{l} 
iterative refinement.
\end{tabular} \\
& \begin{tabular}{l} 
Solves a real symmetric positive definite system of linear \\
equations in band symmetric storage mode with iterative \\
refinement.
\end{tabular} \\
LSAQS see page 300 & \begin{tabular}{l} 
Solves a real system of linear equations in band storage \\
mode with iterative refinement.
\end{tabular} \\
LSARB see page 280 & \begin{tabular}{l} 
Solves a real general system of linear equations with \\
iterative refinement.
\end{tabular} \\
LSAR see page 82 & \begin{tabular}{l} 
Solves a real symmetric system of linear equations with
\end{tabular} \\
iterative refinement.
\end{tabular}
\begin{tabular}{|c|c|}
\hline LSLQB see page 349 & Computes the \(R^{H} D R\) Cholesky factorization of a complex hermitian positive-definite matrix \(A\) in codiagonal band hermitian storage mode. Solve a system \(A x=b\). \\
\hline LSLQH see page 346 & Solves a complex Hermitian positive definite system of linearequations in band Hermitian storage mode without iterative refinement. \\
\hline LSLQS see page 303 & Solves a real symmetric positive definite system of linear equations in band symmetric storage mode without iterative refinement. \\
\hline LSLRB see page 282 & Solves a real system of linear equations in band storage mode without iterative refinement. \\
\hline LSLRG see page 87 & Solves a real general system of linear equations without iterative refinement. \\
\hline LSLRT see page 154 & Solves a real triangular system of linear equations. \\
\hline LSLSF see page 212 & Solves a real symmetric system of linear equations without iterative refinement. \\
\hline LSLTC see page 423 & Solves a complex Toeplitz linear system. \\
\hline LSLTO see page 421 & Solves a real Toeplitz linear system. \\
\hline LSLTQ see page 319 & Solves a complex tridiagonal system of linear equations. \\
\hline LSLTR see page 275 & Solves a real tridiagonal system of linear equations. \\
\hline LSLXD see page 391 & Solves a sparse system of symmetric positive definite linear algebraic equations by Gaussian elimination. \\
\hline LSLXG see page 364 & Solves a sparse system of linear algebraic equations by Gaussian elimination. \\
\hline LSLZD see page 408 & Solves a complex sparse Hermitian positive definite system of linear equations by Gaussian elimination. \\
\hline LSLZG see page 377 & Solves a complex sparse system of linear equations by Gaussian elimination. \\
\hline LSQRR see page 446 & Solves a linear least-squares problem without iterative refinement. \\
\hline LSVCR see page 504 & Computes the singular value decomposition of a complex matrix. \\
\hline LSVRR see page 498 & Computes the singular value decomposition of a real matrix. \\
\hline LUPCH see page 491 & Updates the \(R^{T} R\) Cholesky factorization of a real symmetric positive definite matrix after a rank-one matrix is added. \\
\hline LUPQR see page 484 & Computes an updated \(Q R\) factorization after the rank-one matrix \(\alpha x y{ }^{T}\) is added. \\
\hline \multicolumn{2}{|l|}{M} \\
\hline MCRCR see page 1479 & Multiplies two complex rectangular matrices, \(A B\). \\
\hline MOLCH see page 1038 & Solves a system of partial differential equations of the form \(u_{t}=f\left(x, t, u, u_{x}, u_{x x}\right)\) using the method of lines. The solution is represented with cubic Hermite polynomials. \\
\hline MP_SETUP see page 1758 & Initializes or finalizes MPI. \\
\hline
\end{tabular}

MPS_FREE see page 1343

MRRRR see page 1476
MUCBV see page 1493

MUCRV see page 1491
MURBV see page 1489

MURRV see page 1487
MXTXF see page 1470
MXTYF see page 1472
MXYTF see page 1474

N
NAN see page 1601
IERCD and N1RTY see page
1766
NDAYS see page 1708

NDYIN see page 1710

NEQBF see page 1204

NEQBJ see page 1210

NEQNF see page 1198

NEQNJ see page 1201

NNLPF see page 1377
NNLPG see page 1383
NORM see page 1602

NR1CB see page 1505

NR1RB see page 1504

Deallocates the space allocated for the IMSL derived type S_MPS. This routine is usually used in conjunction with READ_MPS.
Multiplies two real rectangular matrices, \(A B\).
Multiplies a complex band matrix in band storage mode by a complex vector.
Multiplies a complex rectangular matrix by a complex vector.
Multiplies a real band matrix in band storage mode by a real vector.

Multiplies a real rectangular matrix by a vector.
Computes the transpose product of a matrix, \(A^{T} A\).
Multiplies the transpose of matrix \(A\) by matrix \(B, A^{T} B\).
Multiplies a matrx \(A\) by the transpose of a matrix \(B, A B^{T}\).

Returns, as a scalar function, a value corresponding to the IEEE 754 Standard format of floating point (ANSI/IEEE 1985) for NaN .

Retrieves an error type for the most recently called IMSL routine.

Computes the number of days from January 1, 1900, to the given date.
Gives the date corresponding to the number of days since January 1, 1900.
Solves a system of nonlinear equations using factored secant update with a finite-difference approximation to the Jacobian.
Solves a system of nonlinear equations using factored secant update with a user-supplied Jacobian.
Solves a system of nonlinear equations using a modified Powell hybrid algorithm and a finite-difference approximation to the Jacobian.
Solves a system of nonlinear equations using a modified Powell hybrid algorithm with a user-supplied Jacobian.
Uses a sequential equality constrained QP method.
Uses a sequential equality constrained QP method.
Computes the norm of a rank-1 or rank-2 array. For rank-3 arrays, the norms of each rank-2 array, in dimension 3, are computed.
Computes the 1-norm of a complex band matrix in band storage mode.
Computes the 1-norm of a real band matrix in band storage mode.

NR1RR see page 1501
NR2RR see page 1502
NRIRR see page 1499

\section*{0}

OPERATORS:
.h. see page 1556
.hx. see page 1547
.i. see page 1558
.ix. see page 1561
.t. see page 1553
.tx. see page 1541
.x. see page 1537
.xh. see page 1550
.xi. see page 1571
.xt. see page 1544
ORTH see page 1605

\section*{P}

PCGRC see page 427

Computes the 1-norm of a real matrix.
Computes the Frobenius norm of a real rectangular matrix.
Computes the infinity norm of a real matrix.

Computes transpose and conjugate transpose of a matrix. Computes matrix-vector and matrix-matrix products.
Computes the inverse matrix, for square non-singular matrices.

Computes the inverse matrix times a vector or matrix for square non-singular matrices.

Computes transpose and conjugate transpose of a matrix.
Computes matrix-vector and matrix-matrix products.
Computes matrix-vector and matrix-matrix products.
Computes matrix-vector and matrix-matrix products.
Computes the inverse matrix times a vector or matrix for square non-singular matrices.
Computes matrix-vector and matrix-matrix products.
Orthogonalizes the columns of a rank-2 or rank-3 array.

PARALLEL_NONNEGATIVE_LS Q see page \(6 \overline{6}\)

PARALLEL_BOUNDED_LSQ see page 74
PDE_1D_MG see page 1004
PERMA see page 1674
PERMU see page 1673

PGOPT see page 1671
PLOTP see page 1746
POLRG see page 1485
PP1GD see page 776

PPDER see page 774
PPITG see page 780
PPVAL see page 771
PRIME see page 1749

Solves a real symmetric definite linear system using a preconditioned conjugate gradient method with reverse communication.

Solves a linear, non-negative constrained least-squares system.

Solves a linear least-squares system with bounds on the unknowns.

Method of lines with Variable Griddings.
Permutes the rows or columns of a matrix.
Rearranges the elements of an array as specified by a permutation.
Prints a plot of up to 10 sets of points.
Prints a plot of up to 10 sets of points.
Evaluates a real general matrix polynomial.
Evaluates the derivative of a piecewise polynomial on a grid.
Evaluates the derivative of a piecewise polynomial.
Evaluates the integral of a piecewise polynomial.
Evaluates a piecewise polynomial.
Decomposes an integer into its prime factors.

QAND see page 896
QCOSB see page 1135

QCOSF see page 1133

QCOSI see page 1137
QD2DR see page 789

QD2VL see page 786

QD3DR see page 796

QD3VL see page 792

QDAG see page 865

QDAGI see page 872

QDAGP see page 869
QDAGS see page 862

QDAWC see page 886

QDAWF see page 879
QDAWO see page 875
QDAWS see page 883

QDDER see page 784

QDNG see page 889
QDVAL see page 782
QMC see page 899

QPROG see page 1361
QSINB see page 1129

QSINF see page 1126

QSINI see page 1131

Integrates a function on a hyper-rectangle.
Computes a sequence from its cosine Fourier coefficients with only odd wave numbers.
Computes the coefficients of the cosine Fourier transform with only odd wave numbers.
Computes parameters needed by QCOSF and QCOSB.
Evaluates the derivative of a function defined on a rectangular grid using quadratic interpolation.
Evaluates a function defined on a rectangular grid using quadratic interpolation.
Evaluates the derivative of a function defined on a rectangular three-dimensional grid using quadratic interpolation.
Evaluates a function defined on a rectangular threedimensional grid using quadratic interpolation.
Integrates a function using a globally adaptive scheme based on Gauss-Kronrod rules.
Integrates a function over an infinite or semi-infinite interval.
Integrates a function with singularity points given.
Integrates a function (which may have endpoint singularities).
Integrates a function \(\mathrm{F}(\mathrm{X}) /(\mathrm{X}-\mathrm{C})\) in the Cauchy principal value sense.
Computes a Fourier integral.
Integrates a function containing a sine or a cosine.
Integrates a function with algebraic-logarithmic singularities.
Evaluates the derivative of a function defined on a set of points using quadratic interpolation.
Integrates a smooth function using a nonadaptive rule.
Evaluates a function defined on a set of points using quadratic interpolation.
Integrates a function over a hyperrectangle using a quasi-Monte Carlo method.
Solves a quadratic programming problem subject to linear equality/inequality constraints.
Computes a sequence from its sine Fourier coefficients with only odd wave numbers.
Computes the coefficients of the sine Fourier transform with only odd wave numbers.
Computes parameters needed by QSINF and QSINB.

\section*{R}

RAND see page 1608
RAND_GEN see page 1714
RANK see page 1610

RATCH see page 854

RCONV see page 1153
RCORL see page 1163
RCURV see page 806
READ_MPS see page 1333

RECCF see page 908

RECQR see page 911

RLINE see page 803
RNGET see page 1722
RNIN32 see page 1726
RNGE32 see page 1727
RNSE32 see page 1729
RNIN64 see page 1729

RNGE64 see page 1730

RNSE64 see page 1732

RNOPT see page 1724

RNSET see page 1723

RNUN see page 1734
RNUNF see page 1732

Computes a scalar, rank-1, rank-2 or rank-3 array of random numbers.
Generates a rank-1 array of random numbers.
Computes the mathematical rank of a rank-2 or rank-3 array.
Computes a rational weighted Chebyshev approximation to a continuous function on an interval.
Computes the convolution of two real vectors.
Computes the correlation of two real vectors.
Fits a polynomial curve using least squares.
Reads an MPS file containing a linear program problem or a quadratic programming problem.
Computes recurrence coefficients for various monic polynomials.
Computes recurrence coefficients for monic polynomials given a quadrature rule.
Fits a line to a set of data points using least squares.
Retrieves the current value of the seed used in the IMSL random number generators.
Initializes the 32-bit Merseene Twister generator using an array.
Retrieves the current table used in the 32-bit
Mersenne Twister generator.
Sets the current table used in the 32-bit
Mersenne Twister generator.
Initializes the 32-bit Merseene Twister generator using an array.
Retrieves the current table used in the 64-bit Mersenne Twister generator
Sets the current table used in the 64-bit Mersenne Twister generator.
Selects the uniform \((0,1)\) multiplicative congruential pseudorandom number generator.
Initializes a random seed for use in the IMSL random number generators.
Generates pseudorandom numbers from a uniform \((0,1)\) distribution.

Generates a pseudorandom number from a uniform \((0,1)\) distribution.

Adds a scalar to each component of a vector, \(x \leftarrow x+a\), all single precision.
Sums the absolute values of the components of a single-
\begin{tabular}{|c|c|}
\hline & precision vector. \\
\hline SAXPY & Computes the scalar times a vector plus a vector, \(y \leftarrow a x+y\), all single precision. \\
\hline ScaLAPACK_EXIT see page 1640 & Exits ScaLAPACK mode for the IMSL Library routines. \\
\hline ScaLAPACK_GETDIM see page 1624 & Calculates the row and column dimensions of a local distributed array based on the size of the array to be distributed and the row and column blocking factors to be used. \\
\hline ScaLAPACK_MAP see page 1636 & Maps array data from a global array to local arrays in the two-dimensional block-cyclic form required by ScaLAPACK routines. \\
\hline ScaLAPACK_READ see page 1625 & Reads matrix data from a file and transmits it into the twodimensional block-cyclic form required by ScaLAPACK routines. \\
\hline ScaLAPACK_SETUP see page 1622 & Sets up a processor grid and calculates default values for use in mapping arrays to the processor grid \\
\hline ScaLAPACK_UNMAP see page 1637 & Unmaps array data from local distributed arrays to a global array. \\
\hline ScaLAPACK_WRITE see page 1627 & Writes the matrix data to a file. \\
\hline SCASUM & Sums the absolute values of the real part together with the absolute values of the imaginary part of the components of a complex vector. \\
\hline SCNRM2 & Computes the Euclidean norm of a complex vector. \\
\hline SCOPY & Copies a vector \(x\) to a vector \(y\), both single precision. \\
\hline SDDOTA & Computes the sum of a single-precision scalar, a singleprecision dot product and the double-precision accumulator, which is set to the result \(\operatorname{ACC} \leftarrow \operatorname{ACC}+a+x^{T} y\). \\
\hline SDDOTI & Computes the sum of a single-precision scalar plus a singleprecision dot product using a double-precision accumulator, which is set to the result ACC \(\leftarrow a+x^{T} y\). \\
\hline SDOT & Computes the single-precision dot product \(x^{T} y\). \\
\hline SDSDOT & Computes the sum of a single-precision scalar and a single precision dot product, \(a+x^{T} y\), using a double-precision accumulator. \\
\hline SGBMV & \begin{tabular}{l}
Computes one of the matrix-vector operations:
\[
y \leftarrow \alpha A x+\beta y, \text { or } y \leftarrow \alpha A^{T} x+\beta y
\] \\
where \(A\) is a matrix stored in band storage mode.
\end{tabular} \\
\hline SGEMM & Computes one of the matrix-matrix operations:
\[
\begin{aligned}
& C \leftarrow \alpha A B+\beta C, C \leftarrow \alpha A^{T} B+\beta C, C \leftarrow \alpha A B^{T} \\
& +\beta C, \text { or } C \leftarrow \alpha A^{T} B^{T}+\beta C
\end{aligned}
\] \\
\hline SGEMV & Computes one of the matrix-vector operations: \(y \leftarrow \alpha A x+\beta y\), or \(y \leftarrow \alpha A^{T} x+\beta y\), \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline SGER & Computes the rank-one update of a real general matrix: \(A \leftarrow A+\alpha x y^{T}\). \\
\hline SHOW see page 1643 & Prints rank-1 or rank-2 arrays of numbers in a readable format. \\
\hline SHPROD & Computes the Hadamard product of two single-precision vectors. \\
\hline SINLP see page 1175 & Computes the inverse Laplace transform of a complex function. \\
\hline SLCNT see page 1078 & Calculates the indices of eigenvalues of a Sturm-Liouville problem with boundary conditions (at regular points) in a specified subinterval of the real line, \([\alpha, \beta]\). \\
\hline SLEIG see page 1066 & Determines eigenvalues, eigenfunctions and/or spectral density functions for Sturm-Liouville problems in the form with boundary conditions (at regular points). \\
\hline SLPRS see page 1355 & Solves a sparse linear programming problem via the revised simplex algorithm. \\
\hline SNRM2 & Computes the Euclidean length or \(L_{2}\) norm of a singleprecision vector. \\
\hline SORT_REAL see page 1677 & Sorts a rank-1 array of real numbers \(x\) so the \(y\) results are algebraically nondecreasing, \(y_{1} \leq y_{2} \leq \ldots y_{n}\). \\
\hline SPLEZ see page 708 & Computes the values of a spline that either interpolates or fits user-supplied data. \\
\hline SPLINE_CONSTRAINTS see page 652 & Returns the derived type array result. \\
\hline SPLINE_FITTING see page 654 & Weighted least-squares fitting by B-splines to discrete OneDimensional data is performed. \\
\hline SPLINE_VALUES see page 653 & Returns an array result, given an array of input \\
\hline SPRDCT & Multiplies the components of a single-precision vector. \\
\hline SRCH see page 1691 & Searches a sorted vector for a given scalar and return its index. \\
\hline SROT & Applies a Givens plane rotation in single precision. \\
\hline SROTG & Constructs a Givens plane rotation in single precision. \\
\hline SROTM & Applies a modified Givens plane rotation in single precision. \\
\hline SROTMG & Constructs a modified Givens plane rotation in single precision. \\
\hline SSBMV & \begin{tabular}{l}
Computes the matrix-vector operation
\[
y \leftarrow \alpha A x+\beta y
\] \\
where \(A\) is a symmetric matrix in band symmetric storage mode.
\end{tabular} \\
\hline SSCAL & Multiplies a vector by a scalar, \(y \leftarrow a y\), both single precision. \\
\hline SSET & Sets the components of a vector to a scalar, all single precision. \\
\hline SSRCH see page 1696 & Searches a character vector, sorted in ascending ASCII \\
\hline
\end{tabular}
order, for a given string and return its index.

SSUB

SSUM
SSWAP
SSYMM

SSYMV

SSYR

SSYR2

SSYR2K

SSYRK

STBMV

STBSV

STRMM

STRMV

STRSM

Subtracts each component of a vector from a scalar, \(x \leftarrow a-x\), all single precision.
Sums the values of a single-precision vector.
Interchanges vectors \(x\) and \(y\), both single precision.
Computes one of the matrix-matrix operations:
\(C \leftarrow \alpha A B+\beta C\) or \(C \leftarrow \alpha B A+\beta C\),
where \(A\) is a symmetric matrix and \(B\) and \(C\) are \(m\) by \(n\) matrices.
Computes the matrix-vector operation
\(y \leftarrow \alpha A x+\beta y\),
where \(A\) is a symmetric matrix.
Computes the rank-one update of a real symmetric matrix:
\(A \leftarrow A+\alpha x x^{T}\).
Computes the rank-two update of a real symmetric matrix:
\(A \leftarrow A+\alpha x y^{T}+\alpha y x^{T}\).
Computes one of the symmetric rank \(2 k\) operations:
\(C \leftarrow \alpha A B^{T}+\alpha B A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} B+\alpha B^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) and \(B\) are \(n\) by \(k\) matrices in the first case and \(k\) by \(n\) matrices in the second case.
Computes one of the symmetric rank \(k\) operations:
\(C \leftarrow \alpha A A^{T}+\beta C\) or \(C \leftarrow \alpha A^{T} A+\beta C\),
where \(C\) is an \(n\) by \(n\) symmetric matrix and \(A\) is an \(n\) by \(k\) matrix in the first case and a \(k\) by \(n\) matrix in the second case.
Computes one of the matrix-vector operations:
\(x \leftarrow A x\) or \(x \leftarrow A^{T} x\),
where \(A\) is a triangular matrix in band storage mode.
Solves one of the triangular systems:
\(x \leftarrow A^{-1} x\) or \(x \leftarrow\left(A^{-1}\right)^{T} x\),
where \(A\) is a triangular matrix in band storage mode.
Computes one of the matrix-matrix operations:
\(B \leftarrow \alpha A B, B \leftarrow \alpha A^{T} B\), or \(B \leftarrow \alpha B A, B \leftarrow \alpha B A^{T}\),
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
Computes one of the matrix-vector operations:
\(x \leftarrow A x\) or \(x \leftarrow A^{T} x\),
where \(A\) is a triangular matrix.
Solves one of the matrix equations:
\(B \leftarrow \alpha A^{-1} B, B \leftarrow \alpha B A^{-1}, B \leftarrow \alpha\left(A^{-1}\right)^{\mathrm{T}} B\)
or \(B \leftarrow \alpha B\left(A^{-1}\right)^{T}\),
where \(B\) is an \(m\) by \(n\) matrix and \(A\) is a triangular matrix.
STRSV
SUMAG/DUMAG see page
1746

SURF see page 800

SURFACE CONSTRAINTS see page \(6 \overline{64}\)

SURFACE_FITTING see page 666

SURFACE_VALUES see page 665

SVCAL

SVD see page 1612

SVIBN see page 1688
SVIGN see page 1683
SVIGP see page 1684

SVRBN see page 1685
SVRBP see page 1687

SVRGN see page 1679
SVRGP see page 1681

SXYZ

\section*{T}

TDATE see page 1707
TIMDY see page 1706
TRNRR see page 1469
TWODQ see page 891

\section*{U}

UMACH see page 1774
UMAG see page 1743

UMCGF see page 1255

UMCGG see page 1259

Solves one of the triangular linear systems:
\(x \leftarrow A^{-1} x\) or \(x \leftarrow\left(A^{-1}\right)^{T} x\)
where \(A\) is a triangular matrix.
Sets or retrieves MATH/LIBRARY single-precision options.

Computes a smooth bivariate interpolant to scattered data that is locally a quintic polynomial in two variables.

Returns the derived type array result given optional input.

Weighted least-squares fitting by tensor product
B-splines to discrete two-dimensional data is performed.
Returns a tensor product array result, given two arrays of independent variable values.
Multiplies a vector by a scalar and store the result in another vector, \(y \leftarrow a x\), all single precision.
Computes the singular value decomposition of a rank-2 or rank-3 array, \(A=U S V^{T}\).

Sorts an integer array by nondecreasing absolute value.
Sorts an integer array by algebraically increasing value.
Sorts an integer array by algebraically increasing value and returns the permutation that rearranges the array.
Sorts a real array by nondecreasing absolute value.
Sorts a real array by nondecreasing absolute value and returns the permutation that rearranges the array.

Sorts a real array by algebraically increasing value.
Sorts a real array by algebraically increasing value and returns the permutation that rearranges the array.
Computes a single-precision \(x y z\) product.

Gets today's date.
Gets time of day.
Transposes a rectangular matrix.
Computes a two-dimensional iterated integral.

Sets or retrieves input or output device unit numbers.
Handles MATH/LIBRARY and STAT/LIBRARY type REAL and double precision options.
Minimizes a function of N variables using a conjugate gradient algorithm and a finite-difference gradient.
Minimizes a function of N variables using a conjugate
\begin{tabular}{|c|c|}
\hline & gradient algorithm and a user-supplied gradient. \\
\hline UMIAH see page 1249 & Minimizes a function of N variables using a modified Newton method and a user-supplied Hessian. \\
\hline UMIDH see page 1243 & Minimizes a function of \(N\) variables using a modified Newton method and a finite-difference Hessian. \\
\hline UMINF see page 1232 & Minimizes a function of N variables using a modified Newton method and a finite-difference Hessian. \\
\hline UMING see page 1237 & Minimizes a function of N variables using a quasi-New method and a finite-difference gradient. \\
\hline UMPOL see page 1263 & Minimizes a function of N variables using a direct search polytope algorithm. \\
\hline UNIT see page 1614 & Normalizes the columns of a rank-2 or rank-3 array so each has Euclidean length of value one. \\
\hline UNLSF see page 1267 & Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a finite-difference Jacobian. \\
\hline UNLSJ see page 1273 & Solves a nonlinear least squares problem using a modified Levenberg-Marquardt algorithm and a user-supplied Jacobian. \\
\hline UVMGS see page 1229 & Finds the minimum point of a nonsmooth function of a single variable. \\
\hline UVMID see page 1225 & Finds the minimum point of a smooth function of a single variable using both function evaluations and first derivative evaluations. \\
\hline UVMIF see page 1222 & Finds the minimum point of a smooth function of a single variable using only function evaluations. \\
\hline V & \\
\hline VCONC see page 1514 & Computes the convolution of two complex vectors. \\
\hline VCONR see page 1512 & Computes the convolution of two real vectors. \\
\hline VERML see page 1713 & Obtains IMSL MATH/LIBRARY-related version, system and license numbers. \\
\hline W & \\
\hline WRCRL see page 1660 & Prints a complex rectangular matrix with a given format and labels. \\
\hline WRCRN see page 1658 & Prints a complex rectangular matrix with integer row and column labels. \\
\hline WRIRL see page 1655 & Prints an integer rectangular matrix with a given format and labels. \\
\hline WRIRN see page 1653 & Prints an integer rectangular matrix with integer row and column labels. \\
\hline WROPT see page 1664 & Sets or retrieves an option for printing a matrix. \\
\hline WRRRL see page 1649 & Prints a real rectangular matrix with a given format and labels. \\
\hline
\end{tabular}

WRRRN see page 1647

Z
ZANLY see page 1189

ZBREN see page 1192

ZPLRC see page 1184

ZPOCC see page 1188

ZPORC see page 1186

ZQADD

ZQINI

ZQMUL
zeSTO

ZREAL see page 1195

Prints a real rectangular matrix with integer row and column labels.

Finds the zeros of a univariate complex function using Müller's method.
Finds a zero of a real function that changes sign in a given interval.
Finds the zeros of a polynomial with real coefficients using Laguerre's method.

Finds the zeros of a polynomial with complex coefficients using the Jenkins-Traub three-stage algorithm.
Finds the zeros of a polynomial with real coefficients using the Jenkins-Traub three-stage algorithm.

Adds a double complex scalar to the accumulator in extended precision.

Initializes an extended-precision complex accumulator to a double complex scalar.

Multiplies double complex scalars using extended precision.
Stores a double complex approximation to an extendedprecision complex scalar.
Finds the real zeros of a real function using Müller's method.

\title{
Appendix C: References
}

\section*{Aird and Howell}

Aird, Thomas J., and Byron W. Howell (1991), IMSL Technical Report 9103, IMSL, Houston.

\section*{Aird and Rice}

Aird, T.J., and J.R. Rice (1977), Systematic search in high dimensional sets, SIAM Journal on Numerical Analysis, 14, 296-312.

\section*{Akima}

Akima, H. (1970), A new method of interpolation and smooth curve fitting based on local procedures, Journal of the ACM, 17, 589-602.
Akima, H. (1978), A method of bivariate interpolation and smooth surface fitting for irregularly distributed data points, ACM Transactions on Mathematical Software, 4, 148-159.

\section*{Anderson et al.}

Anderson, E., Bai, Z., Bishop, C., Blackford, S., Demmel, J., Dongarra, J., DuCroz, J., Greenbaum, A., Hammarling, S., McKenney, A., and Sorensen, D. (1999), LINPACK Users' Guide, SIAM, \(3^{\text {rd }}\) ed., Philadelphia.

\section*{Arushanian et al.}

Arushanian, O.B., M.K. Samarin, V.V. Voevodin, E.E. Tyrtyshikov, B.S. Garbow, J.M. Boyle, W.R. Cowell, and K.W. Dritz (1983), The TOEPLITZ Package Users' Guide, Argonne National Laboratory, Argonne, Illinois.

\section*{Ashcraft}

Ashcraft, C. (1987), A vector implementation of the multifrontal method for large sparse, symmetric positive definite linear systems, Technical Report ETA-TR-51, Engineering Technology Applications Division, Boeing Computer Services, Seattle, Washington.

\section*{Ashcraft et al.}

Ashcraft, C., R.Grimes, J. Lewis, B. Peyton, and H. Simon (1987), Progress in sparse matrix methods for large linear systems on vector supercomputers. Intern. J. Supercomputer Applic., 1(4), 10-29.

Atkinson
Atkinson, Ken (1978), An Introduction to Numerical Analysis, John Wiley \& Sons, New York.

\section*{Atchison and Hanson}

Atchison, M.A., and R.J. Hanson (1991), An Options Manager for the IMSL Fortran 77 Libraries, Technical Report 9101, IMSL, Houston.

\section*{Bischof et al.}

Bischof, C., J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, D. Sorensen (1988), LAPACK Working Note \#5: Provisional Contents, Argonne National Laboratory Report ANL-88-38, Mathematics and Computer Science.

\section*{Bjorck}

Bjorck, Ake (1967), Iterative refinement of linear least squares solutions I, BIT, 7, 322-337.
Bjorck, Ake (1968), Iterative refinement of linear least squares solutions II, BIT, 8, 8-30.

\section*{Boisvert (1984)}

Boisvert, Ronald (1984), A fourth order accurate fast direct method for the Helmholtz equation, Elliptic Problem Solvers II, (edited by G. Birkhoff and A. Schoenstadt), Academic Press, Orlando, Florida, 35-44.

\section*{Boisvert, Howe, and Kahaner}

Boisvert, Ronald F., Sally E. Howe, and David K. Kahaner (1985), GAMS: A framework for the management of scientific software, ACM Transactions on Mathematical Software, 11, 313-355.

\section*{Boisvert, Howe, Kahaner, and Springmann}

Boisvert, Ronald F., Sally E. Howe, David K. Kahaner, and Jeanne L. Springmann (1990), Guide to Available Mathematical Software, NISTIR 90-4237, National Institute of Standards and Technology, Gaithersburg, Maryland.

\section*{Blackford et al.}

Blackford, L. S., Choi, J., Cleary, A., D'Azevedo, E., Demmel, J., Dhillon, I., Dongarra, J., Hammarling, S., Henry, G., Petitet, A., Stanley, K., Walker, D. and Whaley, R. C., (1997), ScaLAPACK User's Guide, Society for Industrial and Applied Mathematics, Philadephia, PA.

\section*{Brankin et al.}

Brankin, R.W., I. Gladwell, and L.F. Shampine, RKSUITE: a Suite of Runge-Kutta Codes for the Initial Value Problem for ODEs, Softreport 91-1, Mathematics Department, Southern Methodist University, Dallas, Texas, 1991.

\section*{Brenan, Campbell, and Petzold}

Brenan, K.E., S.L. Campbell, L.R. Petzold (1989), Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations, Elseview Science Publ. Co.

\section*{Brenner}

Brenner, N. (1973), Algorithm 467: Matrix transposition in place [F1], Communication of ACM, 16, 692-694.

\section*{Brent}

Brent, R.P. (1971), An algorithm with guaranteed convergence for finding a zero of a function, The Computer Journal, 14, 422-425.
Brent, Richard P. (1973), Algorithms for Minimization without Derivatives, Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

\section*{Brigham}

Brigham, E. Oran (1974), The Fast Fourier Transform, Prentice-Hall, Englewood Cliffs, New Jersey.

\section*{Cheney}

Cheney, E.W. (1966), Introduction to Approximation Theory, McGraw-Hill, New York.

\section*{Cline et al.}

Cline, A.K., C.B. Moler, G.W. Stewart, and J.H. Wilkinson (1979), An estimate for the condition number of a matrix, SIAM Journal of Numerical Analysis, 16, 368-375.

\section*{Cody, Fraser, and Hart}

Cody, W.J., W. Fraser, and J.F. Hart (1968), Rational Chebyshev approximation using linear equations, Numerische Mathematik, 12, 242-251.

\section*{Cohen and Taylor}

Cohen, E. Richard, and Barry N. Taylor (1986), The 1986 Adjustment of the Fundamental Physical Constants, Codata Bulletin, Pergamon Press, New York.

\section*{Cooley and Tukey}

Cooley, J.W., and J.W. Tukey (1965), An algorithm for the machine computation of complex Fourier series, Mathematics of Computation, 19, 297-301.

\section*{Courant and Hilbert}

Courant, R., and D. Hilbert (1962), Methods of Mathematical Physics, Volume II, John Wiley \& Sons, New York, NY.

\section*{Craven and Wahba}

Craven, Peter, and Grace Wahba (1979), Smoothing noisy data with spline functions, Numerische Mathematik, 31, 377-403.

\section*{Crowe et al.}

Crowe, Keith, Yuan-An Fan, Jing Li, Dale Neaderhouser, and Phil Smith (1990), A direct sparse linear equation solver using linked list storage, IMSL Technical Report 9006, IMSL, Houston.

\section*{Crump}

Crump, Kenny S. (1976), Numerical inversion of Laplace transforms using a Fourier series approximation, Journal of the Association for Computing Machinery, 23, 89-96.

\section*{Davis and Rabinowitz}

Davis, Philip F., and Philip Rabinowitz (1984), Methods of Numerical Integration, Academic Press, Orlando, Florida.
de Boor
de Boor, Carl (1978), A Practical Guide to Splines, Springer-Verlag, New York.

\section*{de Hoog, Knight, and Stokes}
de Hoog, F.R., J.H. Knight, and A.N. Stokes (1982), An improved method for numerical inversion of Laplace transforms. SIAM Journal on Scientific and Statistical Computing, 3, 357-366.

\section*{Dennis and Schnabel}

Dennis, J.E., Jr., and Robert B. Schnabel (1983), Numerical Methods for Unconstrained Optimization and Nonlinear Equations, Prentice-Hall, Englewood Cliffs, New Jersey.

\section*{Dongarra et al.}

Dongarra, J.J., and C.B. Moler, (1977) EISPACK - A package for solving matrix eigenvalue problems, Argonne National Laboratory, Argonne, Illinois.
Dongarra, J.J., J.R. Bunch, C.B. Moler, and G.W. Stewart (1979), LINPACK Users' Guide, SIAM, Philadelphia.

Dongarra, J.J., J. DuCroz, S. Hammarling, R. J. Hanson (1988), An Extended Set of Fortran basic linear algebra subprograms, ACM Transactions on Mathematical Software, 14, 1-17.
Dongarra, J.J., J. DuCroz, S. Hammarling, I. Duff (1990), A set of level 3 basic linear algebra subprograms, ACM Transactions on Mathematical Software, 16, 1-17.

\section*{Draper and Smith}

Draper, N.R., and H. Smith (1981), Applied Regression Analysis, second edition, John Wiley \& Sons, New York.

\section*{Du Croz et al.}

Du Croz, Jeremy, P. Mayes, G. and Radicati (1990), Factorization of band matrices using Level-3 BLAS, Proceedings of CONPAR 90 VAPP IV, Lecture Notes in Computer Science, Springer, Berlin, 222.

\section*{Duff and Reid}

Duff, I.S., and J.K. Reid (1983), The multifrontal solution of indefinite sparse symmetric linear equations. ACM Transactions on Mathematical Software, 9, 302-325.
Duff, I.S., and J.K. Reid (1984), The multifrontal solution of unsymmetric sets of linear equations. SIAM Journal on Scientific and Statistical Computing, 5, 633-641.

\section*{Duff et al.}

Duff, I.S., A.M. Erisman, and J.K. Reid (1986), Direct Methods for Sparse Matrices, Clarendon Press, Oxford.

\section*{Enright and Pryce}

Enright, W.H., and J.D. Pryce (1987), Two FORTRAN packages for assessing initial value methods, ACM Transactions on Mathematical Software, 13, 1-22.

\section*{Fabijonas}
B. R. Fabijonas,. Algorithm 838: Airy Functions, ACM Transactions on Mathematical Software, Vol. 30, No. 4, December 2004, Pages 491-501.

\section*{Fabijonas et al.}
B. R. Fabijonas, D. W. Lozier, and F. W. J. Olver Computation of Complex Airy Functions and Their Zeros Using Asymptotics and the Differential Equation, ACM Transactions on Mathematical Software, Vol. 30, No. 4, December 2004, 471-490.

\section*{Forsythe}

Forsythe, G.E. (1957), Generation and use of orthogonal polynomials for fitting data with a digital computer, SIAM Journal on Applied Mathematics, 5, 74-88.

\section*{Fox, Hall, and Schryer}

Fox, P.A., A.D. Hall, and N.L. Schryer (1978), The PORT mathematical subroutine library, ACM Transactions on Mathematical Software, 4, 104-126.

\section*{Garbow}

Garbow, B.S. (1978) CALGO Algorithm 535: The QZ algorithm to solve the generalized eigenvalue problem for complex matrices, ACM Transactions on Mathematical Software, 4, 404-410.

\section*{Garbow et al.}

Garbow, B.S., J.M. Boyle, J.J. Dongarra, and C.B. Moler (1972), Matrix eigensystem Routines: EISPACK Guide Extension, Springer-Verlag, New York.

Garbow, B.S., J.M. Boyle, J.J. Dongarra, and C.B. Moler (1977), Matrix Eigensystem RoutinesEISPACK Guide Extension, Springer-Verlag, New York.

Garbow, B.S., G. Giunta, J.N. Lyness, and A. Murli (1988), Software for an implementation of Weeks' method for the inverse Laplace transform problem, ACM Transactions of Mathematical Software, 14, 163-170.

\section*{Gautschi}

Gautschi, Walter (1968), Construction of Gauss-Christoffel quadrature formulas, Mathematics of Computation, 22, 251-270.

\section*{Gautschi and Milovanofic}

Gautschi, Walter, and Gradimir V. Milovanofic (1985), Gaussian quadrature involving Einstein and Fermi functions with an application to summation of series, Mathematics of Computation, 44, 177-190.

\section*{Gay}

Gay, David M. (1981), Computing optimal locally constrained steps, SIAM Journal on Scientific and Statistical Computing, 2, 186-197.
Gay, David M. (1983), Algorithm 611: Subroutine for unconstrained minimization using a model/trust-region approach, ACM Transactions on Mathematical Software, 9, 503-524.

\section*{Gear}

Gear, C.W. (1971), Numerical Initial Value Problems in Ordinary Differential Equations, Prentice-Hall, Englewood Cliffs, New Jersey.

\section*{Gear and Petzold}

Gear, C.W., and Linda R. Petzold (1984), ODE methods for the solutions of differential/algebraic equations, SIAM Journal Numerical Analysis, 21, \#4, 716.

\section*{George and Liu}

George, A., and J.W.H. Liu (1981), Computer Solution of Large Sparse Positive-definite Systems, Prentice-Hall, Englewood Cliffs, New Jersey.

\section*{Gill et al.}

Gill, Philip E., and Walter Murray (1976), Minimization subject to bounds on the variables, NPL Report NAC 72, National Physical Laboratory, England.
Gill, Philip E., Walter Murray, and Margaret Wright (1981), Practical Optimization, Academic Press, New York.

Gill, P.E., W. Murray, M.A. Saunders, and M.H. Wright (1985), Model building and practical aspects of nonlinear programming, in Computational Mathematical Programming, (edited by K. Schittkowski), NATO ASI Series, 15, Springer-Verlag, Berlin, Germany.

\section*{Goldfarb and Idnani}

Goldfarb, D., and A. Idnani (1983), A numerically stable dual method for solving strictly convex quadratic programs, Mathematical Programming, 27, 1-33.

\section*{Golub}

Golub, G.H. (1973), Some modified matrix eigenvalue problems, SIAM Review, 15, 318-334.

\section*{Golub and Van Loan}

Golub, Gene H., and Charles F. Van Loan (1983), Matrix Computations, Johns Hopkins University Press, Baltimore, Maryland.

Golub, Gene H., and Charles F. Van Loan (1989), Matrix Computations, 2d ed., Johns Hopkins University Press, Baltimore, Maryland.
Golub, Gene H., and Charles F. Van Loan (1996), Matrix Computations, 3rd ed., Johns Hopkins University Press, Baltimore, Maryland.

\section*{Golub and Welsch}

Golub, G.H., and J.H. Welsch (1969), Calculation of Gaussian quadrature rules, Mathematics of Computation, 23, 221-230.

\section*{Gregory and Karney}

Gregory, Robert, and David Karney (1969), A Collection of Matrices for Testing Computational Algorithms, Wiley-Interscience, John Wiley \& Sons, New York.

\section*{Griffin and Redish}

Griffin, R., and K.A. Redish (1970), Remark on Algorithm 347: An efficient algorithm for sorting with minimal storage, Communications of the ACM, 13, 54.

\section*{Grosse}

Grosse, Eric (1980), Tensor spline approximation, Linear Algebra and its Applications, 34, 29-41.

\section*{Guerra and Tapia}

Guerra, V., and R. A. Tapia (1974), A local procedure for error detection and data smoothing, MRC Technical Summary Report 1452, Mathematics Research Center, University of Wisconsin, Madison.

\section*{Hageman and Young}

Hageman, Louis A., and David M.Young (1981), Applied Iterative Methods, Academic Press, New York.

\section*{Hanson}

Hanson, Richard J. (1986), Least squares with bounds and linear constraints, SIAM Journal Sci. Stat. Computing, 7, \#3.

Hanson, Richard.J. (1990), A cyclic reduction solver for the IMSL Mathematics Library, IMSL Technical Report 9002, IMSL, Houston.

\section*{Hanson et al.}

Hanson, Richard J., R. Lehoucq, J. Stolle, and A. Belmonte (1990), Improved performance of certain matrix eigenvalue computations for the IMSL/MATH Library, IMSL Technical Report 9007, IMSL, Houston.

\section*{Hartman}

Hartman, Philip (1964) Ordinary Differential Equations, John Wiley and Sons, New York, NY.

\section*{Hausman}

Hausman, Jr., R.F. (1971), Function Optimization on a Line Segment by Golden Section, Lawrence Radiation Laboratory, University of California, Livermore.

\section*{Hindmarsh}

Hindmarsh, A.C. (1974), GEAR: Ordinary differential equation system solver, Lawrence Livermore Laboratory Report UCID-30001, Revision 3.

\section*{Hull et al.}

Hull, T.E., W.H. Enright, and K.R. Jackson (1976), User's guide for DVERK - A subroutine for solving non-stiff ODEs, Department of Computer Science Technical Report 100, University of Toronto.

\section*{IEEE}

ANSI/IEEE Std 754-1985 (1985), IEEE Standard for Binary Floating-Point Arithmetic, The IEEE, Inc., New York.

\section*{IMSL (1991)}

IMSL (1991), IMSL STAT/LIBRARY User's Manual, Version 2.0, IMSL, Houston.

\section*{Irvine et al.}

Irvine, Larry D., Samuel P. Marin, and Philip W. Smith (1986), Constrained interpolation and smoothing, Constructive Approximation, 2, 129-151.

\section*{Jenkins}

Jenkins, M.A. (1975), Algorithm 493: Zeros of a real polynomial, ACM Transactions on Mathematical Software, 1, 178-189.

\section*{Jenkins and Traub}

Jenkins, M.A., and J.F. Traub (1970), A three-stage algorithm for real polynomials using quadratic iteration, SIAM Journal on Numerical Analysis, 7, 545-566.

Jenkins, M.A., and J.F. Traub (1970), A three-stage variable-shift iteration for polynomial zeros and its relation to generalized Rayleigh iteration, Numerische Mathematik, 14, 252-263.
Jenkins, M.A., and J.F. Traub (1972), Zeros of a complex polynomial, Communications of the \(A C M, 15,97-99\).

\section*{Kennedy and Gentle}

Kennedy, William J., Jr., and James E. Gentle (1980), Statistical Computing, Marcel Dekker, New York.

\section*{Kershaw}

Kershaw, D. (1982), Solution of tridiagonal linear systems and vectorization of the ICCG algorithm on the Cray-1, Parallel Computations, Academic Press, Inc., 85-99.

\section*{Knuth}

Knuth, Donald E. (1973), The Art of Computer Programming, Volume 3: Sorting and Searching, Addison-Wesley Publishing Company, Reading, Mass.

\section*{Krogh}

Krogh, Fred T. (2005), An Algorithm for Linear Programming, http://mathalacarte.com/fkrogh/pub/lp.pdf, Tojunga, CA.

\section*{Lawson et al.}

Lawson, C.L., R.J. Hanson, D.R. Kincaid, and F.T. Krogh (1979), Basic linear algebra subprograms for Fortran usage, ACM Transactions on Mathematical Software, 5, 308-323.

\section*{Leavenworth}

Leavenworth, B. (1960), Algorithm 25: Real zeros of an arbitrary function, Communications of the ACM, 3, 602.

\section*{Levenberg}

Levenberg, K. (1944), A method for the solution of certain problems in least squares, Quarterly of Applied Mathematics, 2, 164-168.

\section*{Lewis et al.}

Lewis, P.A. W., A.S. Goodman, and J.M. Miller (1969), A pseudo-random number generator for the System/360, IBM Systems Journal, 8, 136-146.

\section*{Liepman}

Liepman, David S. (1964), Mathematical constants, in Handbook of Mathematical Functions, Dover Publications, New York.

\section*{Liu}

Liu, J.W.H. (1986), On the storage requirement in the out-of-core multifrontal method for sparse factorization. ACM Transactions on Mathematical Software, 12, 249-264.

Liu, J.W.H. (1987), A collection of routines for an implementation of the multifrontal method, Technical Report CS-87-10, Department of Computer Science, York University, North York, Ontario, Canada.

Liu, J.W.H. (1989), The multifrontal method and paging in sparse Cholesky factorization. \(A C M\) Transactions on Mathematical Software, 15, 310-325.

Liu, J.W.H. (1990), The multifrontal method for sparse matrix solution: theory and practice, Technical Report CS-90-04, Department of Computer Science, York University, North York, Ontario, Canada.

\section*{Liu and Ashcraft}

Liu, J., and C. Ashcraft (1987), A vector implementation of the multifrontal method for large sparse, symmetric positive definite linear systems, Technical Report ETA-TR-51, Engineering Technology Applications Division, Boeing Computer Services, Seattle, Washington.

\section*{Lyness and Giunta}

Lyness, J.N. and G. Giunta (1986), A modification of the Weeks Method for numerical inversion of the Laplace transform, Mathmetics of Computation, 47, 313-322.

\section*{Madsen and Sincovec}

Madsen, N.K., and R.F. Sincovec (1979), Algorithm 540: PDECOL, General collocation software for partial differential equations, ACM Transactions on Mathematical Software, 5, \#3, 326-351.

\section*{Marquardt}

Marquardt, D. (1963), An algorithm for least-squares estimation of nonlinear parameters, SIAM Journal on Applied Mathematics, 11, 431-441.

\section*{Martin and Wilkinson}

Martin, R.S., and J.W. Wilkinson (1968), Reduction of the symmetric eigenproblem \(A x=\lambda B x\) and related problems to standard form, Numerische Mathematik, 11, 99-119.

\section*{Matsumoto and Nishimure}

Makoto Matsumoto and Takuji Nishimura, ACM Transactions on Modeling and Computer Simulation, Vol. 8, No. 1, January 1998, Pages 3-30.

\section*{Micchelli et al.}

Micchelli, C.A., T.J. Rivlin, and S. Winograd (1976), The optimal recovery of smooth functions, Numerische Mathematik, 26, 279-285

Micchelli, C.A., Philip W. Smith, John Swetits, and Joseph D. Ward (1985), Constrained \(L_{p}\) approximation, Constructive Approximation, 1, 93-102.

\section*{Moler and Stewart}

Moler, C., and G.W. Stewart (1973), An algorithm for generalized matrix eigenvalue problems, SIAM Journal on Numerical Analysis, 10, 241-256.

\section*{More et al.}

More, Jorge, Burton Garbow, and Kenneth Hillstrom (1980), User guide for MINPACK-1, Argonne National Labs Report ANL-80-74, Argonne, Illinois.

\section*{Muller}

Muller, D.E. (1956), A method for solving algebraic equations using an automatic computer, Mathematical Tables and Aids to Computation, 10, 208-215.

\section*{Murtagh}

Murtagh, Bruce A. (1981), Advanced Linear Programming: Computation and Practice, McGrawHill, New York.

\section*{Murty}

Murty, Katta G. (1983), Linear Programming, John Wiley and Sons, New York.

\section*{Nelder and Mead}

Nelder, J.A., and R. Mead (1965), A simplex method for function minimization, Computer Journal 7, 308-313.

\section*{Neter and Wasserman}

Neter, John, and William Wasserman (1974), Applied Linear Statistical Models, Richard D. Irwin, Homewood, Ill.

\section*{Park and Miller}

Park, Stephen K., and Keith W. Miller (1988), Random number generators: good ones are hard to find, Communications of the ACM, 31, 1192-1201.

\section*{Parlett}

Parlett, B.N. (1980), The Symmetric Eigenvalue Problem, Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

\section*{Pereyra}

Pereyra, Victor (1978), PASVA3: An adaptive finite-difference FORTRAN program for first order nonlinear boundary value problems, in Lecture Notes in Computer Science, 76, SpringerVerlag, Berlin, 67-88.

\section*{Petro}

Petro, R. (1970), Remark on Algorithm 347: An efficient algorithm for sorting with minimal storage, Communications of the ACM, 13, 624.

\section*{Petzold}

Petzold, L.R. (1982), A description of DASSL: A differential/ algebraic system solver, Proceedings of the IMACS World Congress, Montreal, Canada.

\section*{Piessens et al.}

Piessens, R., E. deDoncker-Kapenga, C.W. Uberhuber, and D.K. Kahaner (1983), QUADPACK, Springer-Verlag, New York.

\section*{Powell}

Powell, M.J.D. (1977), Restart procedures for the conjugate gradient method, Mathematical Programming, 12, 241-254.

Powell, M.J.D. (1978), A fast algorithm for nonlinearly constrained optimization calculations, in Numerical Analysis Proceedings, Dundee 1977, Lecture Notes in Mathematics, (edited by G.A. Watson), 630, Springer-Verlag, Berlin, Germany, 144-157.
Powell, M.J.D. (1983), ZQPCVX a FORTRAN subroutine for convex quadratic programming, DAMTP Report NA17, Cambridge, England.

Powell, M.J.D. (1985), On the quadratic programming algorithm of Goldfarb and Idnani, Mathematical Programming Study, 25, 46-61.
Powell, M.J.D. (1988), A tolerant algorithm for linearly constrained optimization calculations, DAMTP Report NA17, University of Cambridge, England.

Powell, M.J.D. (1989), TOLMIN: A fortran package for linearly constrained optimization calculations, DAMTP Report NA2, University of Cambridge, England.

\section*{Pruess and Fulton}

Pruess, S. and C.T. Fulton (1993), Mathematical Software for Sturm-Liouville Problems, \(A C M\) Transactions on Mathematical Software, 17, 3, 360-376.

\section*{Reinsch}

Reinsch, Christian H. (1967), Smoothing by spline functions, Numerische Mathematik, 10, 177-183.

\section*{Rice}

Rice, J.R. (1983), Numerical Methods, Software, and Analysis, McGraw-Hill, New York.

\section*{Saad and Schultz}

Saad, Y., and M.H. Schultz (1986), GMRES: a generalized minimal residual residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Stat. Comput., 7, 856-869.

\section*{Schittkowski}

Schittkowski, K. (1987), More test examples for nonlinear programming codes, SpringerVerlag, Berlin, 74.

\section*{Schnabel}

Schnabel, Robert B. (1985), Finite Difference Derivatives - Theory and Practice, Report, National Bureau of Standards, Boulder, Colorado.

\section*{Schreiber and Van Loan}

Schreiber, R., and C. Van Loan (1989), A Storage-Efficient WY Representation for Products of Householder Transformations, SIAM J. Sci. Stat. Comp., Vol. 10, No. 1, pp. 53-57, January (1989).

\section*{Scott et al}

Scott, M.R., L.F. Shampine, and G.M. Wing (1969), Invariant Embedding and the Calculation of Eigenvalues for Sturm-Liouville Systems, Computing, 4, 10-23.

\section*{Sewell}

Sewell, Granville (1982), IMSL software for differential equations in one space variable, IMSL Technical Report 8202, IMSL, Houston.

\section*{Shampine}

Shampine, L.F. (1975), Discrete least-squares polynomial fits, Communications of the ACM, 18, 179-180.

\section*{Shampine and Gear}

Shampine, L.F. and C.W. Gear (1979), A user's view of solving stiff ordinary differential equations, SIAM Review, 21, 1-17.

\section*{Sincovec and Madsen}

Sincovec, R.F., and N.K. Madsen (1975), Software for nonlinear partial differential equations, ACM Transactions on Mathematical Software, 1, \#3, 232-260.

\section*{Singleton}

Singleton, R.C. (1969), Algorithm 347: An efficient algorithm for sorting with minimal storage, Communications of the ACM, 12, 185-187.

\section*{Smith}

Smith, B.T. (1967), ZERPOL, A Zero Finding Algorithm for Polynomials Using Laguerre's Method, Department of Computer Science, University of Toronto.

\section*{Smith et al.}

Smith, B.T., J.M. Boyle, J.J. Dongarra, B.S. Garbow, Y. Ikebe, V.C. Klema, and C.B. Moler (1976), Matrix Eigensystem Routines - EISPACK Guide, Springer-Verlag, New York.

\section*{Spang}

Spang, III, H.A. (1962), A review of minimization techniques for non-linear functions, SIAM Review, 4, 357-359.

\section*{Stewart}

Stewart, G.W. (1973), Introduction to Matrix Computations, Academic Press, New York.
Stewart, G.W. (1976), The economical storage of plane rotations, Numerische Mathematik, 25, 137-139.

\section*{Stoer}

Stoer, J. (1985), Principles of sequential quadratic programming methods for solving nonlinear programs, in Computational Mathematical Programming, (edited by K. Schittkowski), NATO ASI Series, 15, Springer-Verlag, Berlin, Germany.

\section*{Stroud and Secrest}

Stroud, A.H., and D.H. Secrest (1963), Gaussian Quadrature Formulae, Prentice-Hall, Englewood Cliffs, New Jersey.

\section*{Titchmarsh}

Titchmarsh, E. Eigenfunction Expansions Associated with Second Order Differential Equations, Part I, 2d Ed., Oxford University Press, London, 1962.

\section*{Trench}

Trench, W.F. (1964), An algorithm for the inversion of finite Toeplitz matrices, Journal of the Society for Industrial and Applied Mathematics, 12, 515-522.

\section*{Walker}

Walker, H.F. (1988), Implementation of the GMRES method using Householder transformations, SIAM J. Sci. Stat. Comput., 9, 152-163.

\section*{Washizu}

Washizu, K. (1968), Variational Methods in Elasticity and Plasticity, Pergamon Press, New York.

\section*{Watkins and Elsner}

Watkins, D.S., and L. Elsner (1990), Convergence of algorithms of decomposition type for the eigenvalue problem, Linear Algebra and Applications (to appear).

\section*{Weeks}

Weeks, W.T. (1966), Numerical inversion of Laplace transforms using Laguerre functions, \(J\). ACM, 13, 419-429.

\section*{Wilkinson}

Wilkinson, J.H., and Howinson, S., and Dewynne, J (1965), The Algebraic Eigenvalue Problem, Oxford University Press, London, 635.

\section*{Wilmot et al.}

Wilkinson, J.H. (1965), The Mathematics of Financial Derivatives: A Student Introduction, Cambridge University Press, NY, 41-57.

\section*{Appendix D: Benchmarking or Timing Programs}

\section*{Scalar Program Descriptions}

A set of benchmark programs is provided to allow the user to compare performance of certain routines with similar functionality. For example, the user may wish to compare the performance of lin_sol_gen with LFTRG and LFSRG. Since performance is dependent on problem size and platform, the user can run the time_sol_gen benchmark to determine which of these routines is likely to perform better with the user's specific configuration.

The benchmark programs are supplied with the product in the examples benchmark subdirectory and are summarized in Table B. These programs call Fortran 90 array functions, in single and double precision, to compare the routines shown in columns A and B of Table B. The main program reads single lines of input:
\begin{tabular}{lll} 
NSIZE & NTRIES PREC & Description" \\
NSIZE & NTRIES PREC
\end{tabular}

QUIT
The parameters of NSIZE and NTRIES appear in summary tables. The parameter PREC has values 1,2 or 3 . The choice depends on whether the user wants precision of single, double or both versions timed. The array functions return a summary table with these 6 values:
1. Average time
2. Standard deviation
3. Total time
4. nsize
5. ntries
6. Time Units/Sec.

As an example, the program time_rand_gend is compiled and linked with the single and double precision timing functions s_rand_gen_bench and d_rand_gen_bench.
The two lines of input are:
```

100000
5 3 "Random Number Benchmarks"
QUIT

```

This routine evaluates the elapsed time to compute 100,000 random numbers obtained with rand_gen and rnun(drnun). The "Average" is the mean of the individual elapsed times for 5 calls to the routines, obtaining 100,000 random numbers in each call. The "St. Dev." is the standard deviation for that "Average". This value indicates the variability of the "Average". In order for this value to provide any useful information it is necessary for \(\mid\) NTRIES \(\mid>1\). The value \(\mid\) NTRIES \(\mid=1\) is acceptable, but only one time sample and no standard deviation is obtained. Values of ntries \(>0\) result in the printing of results as shown in Table A. The numbers in the table will vary depending on the machine and other factors that impact performance of Fortran codes.
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{3}{|l|}{\begin{tabular}{l}
Benchmark of rand_gend and rnun: \\
Date of benchmark, (Y, Mo, D, H, M, S): 2006511
\end{tabular}} & \multirow[b]{2}{*}{\[
\begin{array}{|l|}
\hline 85858 \\
\hline \text { Average }
\end{array}
\]} \\
\hline 1 & \(3.6000 \mathrm{E}+00\) & \(3.2000 \mathrm{E}+00\) & \\
\hline 2 & \(4.8990 \mathrm{E}-01\) & \(4.0000 \mathrm{E}-01\) & St. Dev. \\
\hline 3 & \(1.8000 \mathrm{E}+01\) & \(1.6000 \mathrm{E}+01\) & Total Ticks \\
\hline 4 & \(1.0000 \mathrm{E}+04\) & \(1.0000 \mathrm{E}+04\) & Size \\
\hline 5 & \(5.0000 \mathrm{E}+00\) & \(5.0000 \mathrm{E}+00\) & Repeats \\
\hline 6 & \(5.0000 \mathrm{E}+01\) & \(5.0000 \mathrm{E}+01\) & Ticks per sec. \\
\hline \multicolumn{4}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
Benchmark of rand_gend and rnun: \\
Date of benchmark, (Y, Mo, D, H, M, S): 200651185858
\end{tabular}}} \\
\hline & & & \\
\hline 1 & \(2.8000 \mathrm{E}+00\) & \(3.2000 \mathrm{E}+00\) & Average \\
\hline 2 & \(4.0000 \mathrm{E}-01\) & \(4.0000 \mathrm{E}-01\) & St. Dev. \\
\hline 3 & \(1.4000 \mathrm{E}+01\) & \(1.6000 \mathrm{E}+01\) & Total Ticks \\
\hline 4 & \(1.0000 \mathrm{E}+04\) & \(1.0000 \mathrm{E}+01\) & Size \\
\hline 5 & \(5.0000 \mathrm{E}+00\) & \(5.0000 \mathrm{E}+00\) & Repeats \\
\hline 6 & \(5.0000 \mathrm{E}+01\) & \(5.0000 \mathrm{E}+01\) & Ticks per sec. \\
\hline
\end{tabular}

Table A: Benchmark Summary: rand_gen, rnun, (drnun)

If NTRIES \(<0\) the \(6 \times 2\) functions return the tabular values shown, with \(\mid\) NTRIES \(\mid\) samples. No printing is performed with NTRIES \(<0\).

To compute a related benchmark such as the rate "random numbers per second" for single precision rand_gen, separately calculate
```

rate = size x ticks per sec./average
= 104 x 50/3.6
= 138,889. numbers/sec.
= 0.139 million numbers/sec.

```
\begin{tabular}{|c|c|c|c|}
\hline \multirow[t]{2}{*}{Number} & \multirow[t]{2}{*}{Program Units} & \multicolumn{2}{|r|}{Routines Timed for Comparison} \\
\hline & & A & B \\
\hline 1 & ```
time_dft.f90,
s_dft_bench.f90,
d_dft_bench.f90
``` & fast_dft & \[
\begin{gathered}
\text { fftcf, fftcb } \\
\text { dfftcf, } \mathrm{dfftcb}
\end{gathered}
\] \\
\hline 2 & ```
time_eig_gen.f90,
s_eig_gen_bench.f90,
d_eig_gen_bench.f90
``` & lin_eig_gen & e8crg, de8crg \\
\hline 3 & ```
time_eig_self.f90,
s_eig_self_bench.f90,
d_eig_self_bench.f90
``` & lin_eig_self & e5csf, de5csf \\
\hline 4 & ```
time_geig_gen.f90,
s_geig_gen_bench.f90,
d_geig_gen_bench.f90
``` & lin_geig_gen & g8crg, dg8crg \\
\hline 5 & ```
time_inv_chol.f90,
s_inv_chol_bench.f90,
d_inv_chol_bench.f90
``` & lin_sol_self & 12 nds , dl2nds \\
\hline 6 & ```
time_inv_gen.f90,
s_inv_gen_bench.f90,
d_inv gen bench.f90
``` & lin_sol_gen & l2nrg, dl2nrg \\
\hline 7 & ```
time_inv_lsq.f90,
s_inv_lsq_bench.f90,
d_inv_lsq_bench.f90
``` & lin_sol_lsq & lsgrr, dlsgrr \\
\hline 8 & ```
time_inv_self.f90,
s_inv_self_bench.f90,
d_inv_self_bench.f90
``` & lin_sol_self & \begin{tabular}{l}
lftsf, lfssf \\
dlftsf, dlfssf
\end{tabular} \\
\hline 9 & ```
time_rand_gen.f90,
s_inv_rand_bench.f90,
d_inv_rand_bench.f90
``` & rand_gen & rnun, drnun \\
\hline
\end{tabular}

Table B: Scalar Benchmark Comparisons
\begin{tabular}{|c|c|c|c|}
\hline \multirow[t]{2}{*}{Number} & \multirow[t]{2}{*}{Program Units} & \multicolumn{2}{|r|}{Routines Timed for Comparison} \\
\hline & & A & B \\
\hline 10 & \[
\begin{aligned}
& \text { time_sol_chol.f90, } \\
& \text { s_inv_sol_chol.f90, } \\
& \text { d_inv_sol_chol.f90 }
\end{aligned}
\] & lin_sol_self & \[
\begin{gathered}
\text { lftds, lfsds } \\
\text { dlftds, dlfsds }
\end{gathered}
\] \\
\hline 11 & \[
\begin{aligned}
& \text { time_sol_gen.f90, } \\
& \text { s_sol_gen_bench.f90, } \\
& \text { d_sol_gen_bench.f90 }
\end{aligned}
\] & lin_sol_gen & lftrg, lfsrg dftrg, dlfsrg \\
\hline 12 & \[
\begin{aligned}
& \text { time_sol_lsq.f90, } \\
& \text { s_sol_lsq_bench.f90, } \\
& \text { d_sol_lsq_bench.f90 }
\end{aligned}
\] & lin_sol_lsq & l2rrv, dl2rrv \\
\hline 13 & ```
time_sol_self.f90,
s_sol_self_bench.f90,
d_sol_self bench.f90
``` & lin_sol_self & lftsf, lfssf, dlftsf, dlfssf \\
\hline 14 & ```
time_svd.f90,
s_svd_bench.f90,
d_svd_bench.f90
``` & lin_svd & lsvrr, dlsvrr \\
\hline 15 & ```
time_tri.f90,
s_tri_bench.f90,
d_tri_bench.f90
``` & lin_sol_tri & lslcr, dlslcr \\
\hline 16 & \[
\begin{aligned}
& \text { time_mult.f90 } \\
& \text { s_mult_bench.f90 } \\
& \text { d_mult_bench.f90 }
\end{aligned}
\] & A . x . B & matmul (D, E) \\
\hline
\end{tabular}

Table B- continued: Scalar Benchmark Comparisons

Notes on the comparable problems:
1. Perform forward and backward DFT of a random complex sequence of size NSIZE.
2. Compute eigenexpansion of a random real matrix of dimension NSIZE \(\times\) NSIZE.
3. Compute eigenexpansion of a random symmetric real matrix of dimension NSIZE \(\times\) NSIZE.
4. Compute generalized eigenexpansion of a random matrix pencil of dimension NSIZE \(\times\) NSIZE.
5. Compute the inverse of a positive definite real matrix of dimension NSIZE \(\times\) NSIZE. Uses Cholesky method.
6. Compute the inverse of a general real random matrix of dimension NSIZE \(\times\) NSIZE. Uses LU factorization.
7. Compute the generalized inverse of a general real random matrix of dimension \((2 \times\) NSIZE \() \times\) NSIZE. Uses QR factorization for lin_sol_lsq and SVD for LSGRR.
8. Compute the inverse of a real, symmetric random matrix of dimension NSIZE \(\times\) NSIZE. Uses Aasen's decomposition for lin_sol_self and BunchKaufman decomposition for LFTSF.
9. Generate NSIZE random numbers.
10. Solve a single system of linear equations with a positive definite real random matrix of dimension NSIZE \(\times\) NSIZE.
11. Solve a single system of linear equations with a general real random matrix of dimension NSIZE \(\times\) NSIZE.
12. Solve a single least-squares system of linear equations with a real random matrix of dimension \((2 \times\) NSIZE \() \times\) NSIZE.
13. Solve a single system of linear equations with a symmetric real random matrix of dimension NSIZE \(\times\) NSIZE.
14. Compute the full singular value decomposition of a general real random matrix of dimension NSIZE \(\times\) NSIZE.
15. Solve NSIZE systems of linear equations of a nonsymmetric NSIZE \(\times\) NSIZE tridiagonal matrix. Uses cyclic reduction.
16. Compute products of square matrices of size NSIZE \(\times\) NSIZE. Compare the IMSL defined operation \(C=A . x . B\) with \(F=\operatorname{matmul}(D, E)\). The arrays are assumed shape. Identical problems \(A=D\) and \(B=E\) are timed.
17. Compare times to use SHOW () for writing a random array of size NSIZE to a CHARACTER buffer vs. writing the same array to a scratch file.

\section*{Parallel Program Descriptions}

A set of parallel benchmark programs is shown in Table D. These main programs call Fortran 90 box data type functions, in single and double precision. They compare our parallel allocation algorithm to a scalar sequential method. The main program reads single lines of input:
```

NSIZE NTIMES NRACKS PREC ROOT_WORKS "Description"
QUIT to Stop

```

Two initial lines of output echo the "Description" field, whether or not the root is working, and the number of processors in the MPI communicator. The parameters NSIZE, NTRIES and NRACKS appear in the summary tables. The parameter PREC has values 1,2 or 3 . The choice depends on
whether the user wants precision of single, double or both versions timed. The array functions return a \(7 \times 2\) summary table of values. The \((1: 6,1)\) and \((1: 6,2)\) elements of this array represent the results and parameters of the benchmark for the parallel and non-parallel versions. The \((7,1)\) and \((7,2)\) elements of this array represent the ratio of the parallel to the scalar times and a first-order approximation to the variation in the ratio.

\section*{Parallel Box Version}
1. Average time
2. Standard deviation
3. Total Seconds
4. nsize
5. nracks
6. ntries
7. Parallel/Scalar Ratio

\section*{Scalar Box Equivalent}

Average time
Standard deviation
Total Seconds
nsize
nracks
ntries
Variation in Ratio

As an example, the program time_parallel_i is compiled and linked with the single and double precision timing functions s_parallel_i_bench and d_parallel_i_bench.

This routine evaluates the time to compute 4 inverse matrices of size 600 by 600 using the defined operator .i. The "Average" is the mean of the individual elapsed times for 5 calls to the routines, obtaining 4 inverses in each call. The "St. Dev." is the standard deviation for that "Average". This value indicates the variability of the "Average". In order for this value to provide any useful information it is necessary for \(\mid\) NTRIES \(\mid>1\). The value \(\mid\) NTRIES \(\mid=1\) is acceptable, but only one time sample and no standard deviation is obtained. Values of NTRIES \(>0\) result in the printing of results as shown in Table C. The numbers in the table will vary depending on the machine and other factors that impact performance of Fortran codes. If NTRIES \(<0\) the \(7 \times 2\) functions return


Single precision benchmark of parallel .i. and non-parallel .i.:
Date of benchmark, (Y, Mo, D, H, M, S): 200651185858
\begin{tabular}{|l|c|c|l|}
\hline 1 & \(1.5815 \mathrm{E}+00\) & \(4.0241 \mathrm{E}+00\) & Average \\
\hline 2 & \(2.5031 \mathrm{E}-01\) & \(1.8035 \mathrm{E}-02\) & St. Dev. \\
3 & \(7.9077 \mathrm{E}+00\) & \(2.0121 \mathrm{E}+01\) & Total Seconds \\
4 & \(5.0000 \mathrm{E}+01\) & \(5.0000 \mathrm{E}+01\) & Size \\
5 & \(5.0000 \mathrm{E}+00\) & \(5.0000 \mathrm{E}+00\) & Racks per box \\
6 & \(5.0000 \mathrm{E}+00\) & \(5.0000 \mathrm{E}+00\) & Repeats \\
\hline
\end{tabular}

Non-parallel/parallel averages and variation:
\begin{tabular}{|l|l|l|l|}
\hline & \(2.5444 \mathrm{E}+00\) & \(3.9129 \mathrm{E}-01\) & \\
\hline
\end{tabular}

Double precision benchmark of parallel .i. and non-parallel .i.:
Date of benchmark, (Y, Mo, D, H, M, S): 200651185859
\begin{tabular}{|c|c|c|c|}
\hline 1 & \(1.6985 \mathrm{D}+00\) & \(4.0372 \mathrm{D}+00\) & Average \\
\hline 2 & 9.8576D-01 & \(2.3836 \mathrm{D}-02\) & St. Dev. \\
\hline 3 & \(8.4923 \mathrm{D}+00\) & \(2.0186 \mathrm{D}+01\) & Total Seconds \\
\hline 4 & \(5.0000 \mathrm{D}+01\) & \(5.0000 \mathrm{D}+01\) & Size \\
\hline 5 & \(5.0000 \mathrm{D}+00\) & \(5.0000 \mathrm{D}+00\) & Racks per box \\
\hline 6 & \(5.0000 \mathrm{D}+00\) & \(5.0000 \mathrm{D}+00\) & Repeats \\
\hline \multicolumn{4}{|l|}{Non-parallel/parallel averages and variation:} \\
\hline & \(2.3770 \mathrm{D}+00\) & 1.2392D-01 & \\
\hline
\end{tabular}

Table C: Performance Summary: Box operator .i.

Below is a list of the performance evaluation programs that time the box data computations using parallel and non-parallel resources.
\begin{tabular}{|c|c|c|}
\hline Number & Program Units & Function Timed \\
\hline 1 & \begin{tabular}{l}
time_parallel_i.f90, \\
s_parallel_i_bench.f90, \\
d parallel i bench.f90
\end{tabular} & .i. A \\
\hline 2 & time_parallel_ix.f90, s_parallel_ix_bench.f90, d_parallel_ix_bench.f90 & A .ix. B \\
\hline 3 & \begin{tabular}{l}
time_parallel_xi.f90, s_parallel_xi_bench.f90, \\
d_parallel_xi_bench.f90
\end{tabular} & B .xi. A \\
\hline 4 & \begin{tabular}{l}
time_parallel_x.f90, \\
s_parallel_x_bench.f90, \\
d_parallel_x_bench.f90
\end{tabular} & A .x. B \\
\hline 5 & time_parallel_tx.f90, s_parallel_tx_bench.f90, d_parallel_tx_bench.f90 & A .tx. B \\
\hline 6 & \begin{tabular}{l}
time_parallel_xt.f90, s_parallel_xt_bench.f90, \\
d_parallel_xt_bench.f90
\end{tabular} & A . xt. B \\
\hline 7 & \begin{tabular}{l}
time_parallel_hx.f90, \\
s_parallel_hx_bench.f90, \\
d_parallel_hx_bench.f90
\end{tabular} & A .hx. B \\
\hline 8 & \begin{tabular}{l}
time_parallel_xh.f90, s_parallel_xh_bench.f90, \\
d_parallel_xh_bench.f90
\end{tabular} & A .xh. B \\
\hline 9 & ```
time_parallel_chol.f90,
s_parallel_chol_bench.f90,
d_parallel_chol_bench.f90
``` & CHOL (A) \\
\hline 10 & \begin{tabular}{l}
time_parallel_cond.f90, s_parallel_cond_bench.f90, \\
d_parallel_cond_bench.f90
\end{tabular} & COND (A) \\
\hline 11 & time_parallel_rank.f90, s_parallel_rank_bench.f90, d_parallel_rank_bench.f90 & RANK ( A ) \\
\hline
\end{tabular}

Table D: Parallel and non-Parallel Box Comparisons
\begin{tabular}{|c|c|c|}
\hline Number & Program Units & Function Timed \\
\hline 12 & \begin{tabular}{l}
time_parallel_det.f90, s_parallel_det_bench.f90, \\
d parallel det bench.f90
\end{tabular} & DET (A) \\
\hline 13 & time_parallel_orth.f90, s_parallel_orth_bench.f90, d_parallel_orht_bench.f90 & ORTH ( \(\mathrm{A}, \mathrm{R}=\mathrm{R}\) ) \\
\hline 14 & ```
time_parallel_svd.f90,
s_parallel_svd_bench.f90,
d_parallel_svd_bench.f90
``` & SVD ( \(\mathrm{A}, \mathrm{U}=\mathrm{U}, \mathrm{V}=\mathrm{V}\) ) \\
\hline 15 & ```
time_parallel_norm.f90,
s_parallel_norm_bench.f90,
d_parallel_norm_bench.f90
``` & \(\operatorname{NORM}(\mathrm{A}, \mathrm{TYPE}=\mathrm{I})\) \\
\hline 16 & ```
time_parallel_eig.f90,
s_parallel_eig_bench.f90,
d_parallel_eig_bench.f90
``` & EIG ( \(\mathrm{A}, \mathrm{W}=\mathrm{W}\) ) \\
\hline 17 & \begin{tabular}{l}
time_parallel_fft.f90, s_parallel_fft_bench.f90, \\
d_parallel_fft_bench.f90
\end{tabular} & \[
\begin{aligned}
& \text { FFT_BOX(A) } \\
& \text { IFFT_BOX(A) }
\end{aligned}
\] \\
\hline
\end{tabular}

Table D continued: Parallel and non-Parallel Box Comparisons

\section*{Product Support}

\section*{Contacting Visual Numerics Support}

Users within support warranty may contact Visual Numerics regarding the use of the IMSL Fortran Numerical Library. Visual Numerics can consult on the following topics:
- Clarity of documentation
- Possible Visual Numerics-related programming problems
- Choice of IMSL Libraries functions or procedures for a particular problem

Not included in these topics are mathematical/statistical consulting and debugging of your program.

\section*{Refer to the following for Visual Numerics Product Support contact information:}
- http://www.vni.com/tech/imsl/phone.html

The following describes the procedure for consultation with Visual Numerics:
1. Include your Visual Numerics license number
2. Include the product name and version number: IMSL Fortran Numerical Library Version 6.0
3. Include compiler and operating system version numbers
4. Include the name of the routine for which assistance is needed and a description of the problem
band symmetric storage mode 300 , 303, 308, 311, 313, 315, 318, \(319,321,324,327,330,333\), \(336,338,342,344,346,349\), 352, 355, 358, 360, 362, 364, \(369,374,573,575,578,581\), 584, 586, 589, 1465, 1778
band triangular storage mode 1780
Basic Linear Algebra Subprograms 1422
basis functions 811
bidiagonal matrix 59
bilinear form 1483
BLACS 1619
BLAS 1422, 1423, 1433, 1434, 1435
Level 1 1422, 1423
Level 2 1433, 1434
Level 3 1433, 1434, 1435
block-cyclic decomposition
reading, writing utility 1620
boundary conditions 961
boundary value problem 53
Brenan 54
Broyden's update 1184
B-spline coefficients 711, 815, 824
B-spline representation 731, 732, 735, 738, 770
B-splines 646

\section*{C}

Campbell 54
Cauchy principal value 860,886
central differences 1390
changing messages 1642
character arguments 1699
character sequence 1703
character string 1704
character workspace 1787
Chebyshev approximation 649, 854
Chebyshev polynomials 30
Cholesky
algorithm 20
decomposition 18, 525, 538
factorization 1574, 1575
method 22
Cholesky decomposition 489
Cholesky factorization 185, 190,
194, 204, 305, 308, 311, 318,
349, 362, 395, 399, 404, 412, 417, 421, 491, 494
circulant linear system 425
circulant matrices 8
classical weight functions 901,914
codiagonal band hermitian storage mode 349
codiagonal band Hermitian storage mode 1782
codiagonal band symmetric storage mode 305, 1781
coefficient matrix \(293,313,336\), \(358,374,377,382,387,391\), 395, 399, 404, 408, 417, 421, 423, 425, 427, 433, 436, 446, 452, 458, 462, 466, 473, 478, 484, 489, 491, 498, 504, 508
coefficients 1126,1133
column pivoting 489
companion matrix 531
complex function 1172, 1175
complex periodic sequence 1111, 1113
complex sparse Hermitian positive definite system 408, 417, 421
complex sparse system 377,387
complex triangular system 164
complex tridiagonal system 319
complex vectors 1158,1168
computing
eigenvalues, example 522
the rank of A 36
the SVD 60
computing eigenvalues, example 530
condition number 157, 168, 534
conjugate gradient algorithm 1255, 1259
conjugate gradient method 427, 433
continuous Fourier transform 1085
continuous function 854
convolution 1153, 1158, 1512, 1514
convolutions, real or complex periodic sequences 1092
coordinate transformation 478
correlation 1163, 1168
cosine 875
cosine Fourier coefficients 1135
cosine Fourier transform 1133
covariance matrix 22, 27, 28
CPU time 1705
crossvalidation 851
cross-validation with weighting, example 64
cubic spline \(699,700,703,706\)
cubic spline approximation 848, 851
cubic spline interpolant 677, 680, 682, 687, 690, 692, 696
cubic splines 647
cyclic reduction 44,47
cyclic reduction algorithm 321
cyclical 2D data, linear trend 1097
cyclical data, linear trend 1089

\section*{D}

DASPG routine 54
data fitting
polynomial 30
two dimensional 33
data points 803
data, optional xxvii
date \(1707,1708,1710,1711\)
decomposition, singular value 36,16
degree of accuracy 1763
DENSE_LP 1346
deprecated routines 1787
determinant 1581, 1582, 7
determinant of A 9
determinants \(113,148,161,163\), 204, 224, 274, 298, 318, 342, 362
determinants 7
DFT (Discrete Fourier Transform) 1086, 1099
differential algebraic equations 924
Differential Algebraic Equations 540
differential equations 923, 961
differential-algebraic solver 54
diffusion equation 53
direct- access message file 1643
direct search complex algorithm 1306
direct search polytope algorithm 1263
discrete Fourier cosine transformation 1122
discrete Fourier sine transformation 1118
discrete Fourier transform 1085, 1592, 1593, 1594, 1595, 1597, 1598, 1599, 9, 11
inverse 1596, 11
dot product 1426, 1427, 1428
double precision xix, 1517
DOUBLE PRECISION types xxii

\section*{E}
efficient solution method 532
eigensystem
complex 555, 627, 629, 632
Hermitian 607
real 548, 571, 618, 621, 625
symmetric 589,639
eigenvalue \(1586,1587,8\)
eigenvalue-eigenvector
decomposition 521, 525, 1586, 1587, 8
expansion (eigenexpansion) 523
eigenvalues 543, 545, 550, 552, 557, 559, 561, 563, 566, 568, 573, 575, 578, 581, 584, 586, 591, 593, 596, 599, 602, 604, 609, 611, 614, 616, 618, 621, 627, 629, 634, 636
eigenvalues, self-adjoint matrix 23 , 520, 527, 16
eigenvectors \(50,520,523,525,527\), \(545,552,559,563,568,575\), 581, 586, 593, 599, 604, 611, 616, 621, 629, 636
EISPACK xxxii
endpoint singularities 862
equality constraint, least squares 35
error detection 844
error handling xxix, 1766
errors \(1763,1764,1765\)
alert 1764
detection 1763
fatal 1764
informational 1764
multiple 1763
note 1764
printing error messages 1640
severity 1763
terminal 1763, 1765
warning 1764
Euclidean (2-norm) distance 1507
Euclidean length 1614,1615
even sequence 1122
example
least-squares, by rows
distributed 71
linear constraints
distributed 77
linear inequalities distributed 68
linear system
distributed, ScaLAPACK 1633, 1638
matrix product distributed, PBLAS 1631
Newton's Method
distributed 79
transposing matrix distributed 1628
examples
accuracy estimates of eigenvalues 534
accurate least-squares solution with iterative refinement 25
analysis and reduction of a generalized eigensystem 525
complex polynomial equation Roots 531
computing eigenvalues 522,530
computing eigenvectors with inverse iteration 523
computing generalized eigenvalues 538
computing the SVD 60
constraining a spline surface to be non-negative interpolation to data 675
constraining points using spline surface 673
convolution with Fourier Transform 1092
cross-validation with weighting 64
cyclical 2D data with a linear trend 1097
cyclical data with a linear trend 1089
eigenvalue-eigenvector expansion of a square matrix 523
evaluating the matrix exponential 14,16
Generalized Singular Value Decomposition 62
generating strategy with a histogram 1718
generating with a Cosine distribution 1720
internal write of an array 1646
iterative refinement and use of partial pivoting 48
Laplace transform solution 41
larger data uncertainty 541
least squares with an equality constraint 35
least-squares solution of a rectangular system 38
linear least squares with a quadratic constraint 60
matrix inversion and determinant 13
natural cubic spline interpolation to data 655
parametric representation of a sphere 671
periodic curves 662
polar decomposition of a square matrix 39
printing an array 1645
reduction of an array of black and white 40
ridge regression 64
running mean and variance 1716
seeding, using, and restoring the generator 1717
selected eigenvectors of tridiagonal matrices 50
self-adjoint, positive definite generalized eigenvalue problem 539
several 2D transforms with initialization 1098
several transforms with initialization 1091
shaping a curve and its derivatives 657
solution of multiple tridiagonal systems 47
solving a linear least squares system of equations 21, 30
solving a linear system of equations 12
solving parametric linear systems with scalar change 532
sort and final move with a permutation 1678
sorting an array 1678
splines model a random number generator 659
system solving with Cholesky method 22
system solving with the generalized inverse 32
tensor product spline fitting of data 669
test for a regular matrix pencil 540
transforming array of random complex numbers 1089, 1096, 1102
tridiagonal matrix solving 53
two-dimensional data fitting 33
using inverse iteration for an eigenvector 23
exclusive OR 1715
extended precision arithmetic 1516

\section*{F}
factored secant update 1204, 1210
factorization, LU 9
Fast Fourier Transforms 1084
Faure 1736, 1738, 37, 9
Faure sequence 1736, 1737, 37, 9
Fejer quadrature rule 914
FFT (Fast Fourier Transform) 1088, 1096, 1102
finite difference gradient 1377
finite-difference approximation 1198, 1204
finite-difference gradient 1232, 1255, 1279
finite-difference Hessian 1243
finite-difference Jacobian 1267
first derivative 918
first derivative evaluations 1225
first order differential 980
FORTRAN 77
combining with Fortran 90 xix
Fortran 90
language xix
rank-2 array xxvi
real-time clock 1716
forward differences 1392, 1394, 1397, 1400
Fourier coefficients 1103, 1106, 1111, 1113, 1139, 1145
Fourier integral 879
Fourier transform 1142, 1149
Frobenius norm 1502
full storage mode 1455
Fushimi 1715, 1716

\section*{G}

Galerkin principle 54
Gauss quadrature 861
Gauss quadrature rule 901, 905
Gaussian elimination 364, 369, 374, 377, 391, 408, 412
Gauss-Kronrod rules 865
Gauss-Lobatto quadrature rule 901, 905
Gauss-Radau quadrature rule 901, 905
Gear's BDF method 944
generalized
eigenvalue \(525,538,1586,1587,8\)
feedback shift register (GFSR) 1714
inverse matrix 27, 28, 32
generalized inverse
system solving 32
generator 1717, 1720
getting started xxvi
GFSR algorithm 1715
Givens plane rotation 1430
Givens transformations 1432, 1433
globally adaptive scheme 865
Golub 12, 20, 30, 35, 59, 62, 64, 521, 525, 530
gradient 1390, 1392, 1397, 1403
Gray code 1739
GSVD 62

\section*{H}

Hadamard product 1428, 1481
Hanson 521
harmonic series 1089, 1097
Helmholtz's equation 1053
Helmholtz's equation 1059
Hermite interpolant 687
Hermite polynomials 1038
Hermitian positive definite system 226, 231, 246, 251, 344, 346, 358, 360
Hermitian system 258, 261, 269, 271
Hessenberg matrix, upper 527, 531

Hessian 1249, 1293, 1299, 1394, 1397, 1406
High Performance Fortran
HPF 1620
histogram 1718
Horner's scheme 1486
Householder 538
Householder transformations 452, 466
hyper-rectangle 896

\section*{I}

IEEE 1601, 1602, 12
infinite eigenvalues 538
infinite interval 872
infinity norm 1499
infinity norm distance 1510
informational errors 1764
initialization, several 2D transforms 1098
initialization, several transforms 1091
initial-value problem 927, 934, 944
integer options 1739
INTEGER types xxii
integrals 706
integration \(862,865,869,872,875\), 883, 886, 889, 896
interface block xix
internal write 1646
interpolation 651
cubic spline 677,680
quadratic 649
scattered data 649
inverse 9
iteration, computing eigenvectors 23, 51, 523
matrix xxvi, 10, 17, 18, 22
generalized 27, 28
transform 1087, 1094, 1100
inverse matrix 9
isNaN 1601, 1602
ISO xix
iterated integral 891
iterative refinement xxvii, 6, 7, 48, 82, 107, 142, 176, 181, 185, 190, 194, 199, 204, 205, 209, 212, 221, 251, 271, 295, 315, 338, 344, 360, 446, 458
IVPAG routine 54

\section*{J}

Jacobian 1184, 1198, 1201, 1204, 1210, 1273, 1310, 1317, 1400, 1410

Jenkins-Traub three-stage algorithm 1186

K
Kershaw 47

L
Laguerre's method 1184
LAPACK xxxii, \(543,546,551,553\), 557, 560, 592, 594, 619, 622, 628
Laplace transform 1172, 1175
Laplace transform solution 41
larger data uncertainty, example 541
LDU factorization 321
least squares \(21,27,33,35,36,41\), \(42,649,803,806,824,1090\), 1097, 16
least-squares approximation 811,819
least-squares problem 478
least-squares solution 452
Lebesque measure 1738
Level 1 BLAS 1422, 1423
Level 2 BLAS 1433, 1434
Level 3 BLAS 1433, 1434, 1435
Levenberg-Marquardt algorithm 1218, 1267, 1273, 1310, 1317
library subprograms xxiii
linear algebraic equations 364,391
linear constraints 462
linear equality/inequality constraints 1364, 1370
linear equations 17
solving \(82,87,103,118,123,138\), 164, 176, 181, 194, 199, 209, 212, 219, 221, 226, 231, 246, 251, 258, 261, 269, 271, 275, 280, 282, 293, 295, 300, 303, 313, 315, 319, 338, 344, 346, \(358,360,374,377,387,391\), 404, 408, 417, 421, 427
linear least-squares problem 446, 458, 462
linear least-squares with nonnegativity constraints 66,67 , 69, 76
linear programming problem 1346, 1351, 1355
linear solutions
packaged options 11
linear trend, cyclical 2D data 1097
linear trend, cyclical data 1089
LINPACK xxxii, 543, 546, 551, 553, 557, 560, 592, 594, 619, 622, 628
low-discrepancy 1739
\(L U\) factorization 93, 98, 103, 113, 127, 133, 138, 148, 287, 290, 293, 298, 330, 333, 336, 342, 369, 374, 382, 387
LU factorization of A 9, 10, 11, 18, 1522

M
machine-dependent constants 1769
mathematical constants 1751
matrices \(1444,1445,1447,1448\), 1450, 1452, 1453, 1455, 1457, 1458, 1460, 1465, 1467, 1469, 1476, 1479, 1487, 1489, 1491, 1497, 1502, 1504, 1505, 1647, 1649, 1653, 1655, 1658, 1660, 1664
adjoint xxiii
complex 330, 333, 342, 504, 550, 552, 1455, 1460
band 1448, 1493, 1497, 1505
general \(127,148,149,1445\), 1453, 1457
general sparse 382
Hermitian 236, 263, 266, 274, \(349,352,355,362,591,593\), 596, 599, 602, 604, 1463, 1467
rectangular \(1458,1479,1491\), 1658, 1660
tridiagonal 321
upper Hessenberg 614, 616
copying \(1444,1445,1447,1448\), 1457, 1458, 1465, 1467
covariance \(22,27,28\)
general 1775
Hermitian 1776
inverse xxvi, \(9,10,17,18,22\) generalized 27, 28, 32
inversion and determinant 13
multiplying 1474, 1476, 1479, 1487, 1489, 1491
orthogonal xxiii
permutation 1674
poorly conditioned 38
printing 1647, 1649, 1653, 1655, 1658, 1660, 1664
real 287, 290, 298, 508, 543, 545, 1452, 1460
band 1447, 1489, 1504
general 93, 98, 113, 114, 1444, 1450, 1457
general sparse 369
rectangular 1458, 1476, 1481, 1487, 1502, 1647, 1649
sparse 6
symmetric \(185,190,204,205\),
214, 217, 224, 305, 308, 311,
318, 491, 494, 557, 559, 561,
\(563,566,568,573,575,578\),
581, 584, 586, 1462, 1465
tridiagonal 277
upper Hessenberg 609, 611
rectangular 1469, 1775
sparse
Hermitian 412
symmetric 395
symmetrical 399
symmetric 489,1776
transposing 1469, 1470, 1472
triangular 1776
unitary xxiii
upper Hessenberg 531
matrix
inversion 7
types 5
matrix pencil 538, 540
matrix permutation 1674
matrix storage modes 1775
matrix/vector operations 1443
matrix-matrix multiply 1440, 1441, 1442
matrix-matrix solve 1443
matrix-vector multiply 1437, 1438, 1439
means 1716
Mersenne Twister 1726, 1727, 1729, 1730, 1732
message file
building new direct-access
message file 1643
changing messages 1642
management 1642
private message files 1643
Metcalf xix
method of lines 54, 1038
minimization \(1218,1219,1220\), 1222, 1225, 1229, 1232, 1237, 1243, 1249, 1255, 1259, 1263, 1279, 1286, 1293, 1299, 1306, 1310, 1346, 1351, 1364, 1370, 1377, 1383, 1390, 1392, 1394, 1397, 1400, 1403, 1406, 1410, 1414
minimum degree ordering 395
minimum point \(1222,1225,1229\)
mistake
missing argument 1622
Type, Kind or Rank
TKR 1622
Modified Gram-Schmidt algorithm 1606
modified Powell hybrid algorithm 1198, 1201
monic polynomials 908, 911
Moore-Penrose 1558, 1559, 1562,
1571, 1572
MPI 1528, 1759
parallelism xxxvi
Muller's method 1184, 1189
multiple right sides 7
multivariate functions 1218
multivariate quadrature 861

\section*{N}
naming conventions xxii
NaN (Not a Number) 1601
quiet 1601
signaling 1601
Newton algorithm 1218
Newton method 1243, 1249, 1293, 1299
Newton' s method 42, 60
noisy data 848,851
nonadaptive rule 889
nonlinear equations 1198, 1201, 1204, 1210
nonlinear least-squares problem 1218, 1267, 1273, 1310, 1317, 1324
nonlinear programming 1377, 1383
norm 1602
normalize 1614
not-a-knot condition 677, 680
numerical differentiation 862

\section*{0}
odd sequence 1118
odd wave numbers 1126,1129 , 1133, 1135
optional argument xxvi
optional data xxvi, xxvii
optional subprogram arguments xxvi ordinary differential equations 923 , 924, 927, 934, 944
ordinary eigenvectors, example 534
orthogonal
decomposition 59
factorization 29
matrix xxiii
orthogonal matrix 473
orthogonalized 51, 523
overflow xxiv

\section*{P}
page length 1671
page width 1671
parameters \(1109,1116,1120,1124\), 1131, 1137
parametric linear systems with scalar change 532
parametric systems 532
partial differential equations 923 , 925, 1038
partial pivoting 44, 47
PBLAS 1619
performance index \(548,555,571\), 589, 607, 625, 632, 639
periodic boundary conditions 696
permutation 1678
Petzold 54, 980
physical constants 1751
piecewise polynomial 645, 770, 771, 774, 776, 780
piecewise-linear Galerkin 54
pivoting
partial 9, 12, 19
row and column 27, 30
symmetric 17
plane rotation 1431
plots 1746
Poisson solver 1053, 1059
Poisson's equation 1053, 1059
polar decomposition 39, 48
polynomial 1485
polynomial curve 806
prime factors 1749
printing \(1671,1746,1765\)
printing an array, example 1645
printing arrays 1643
printing results xxix
private message files 1643
programming conventions xxiv
pseudorandom number generators 1724
pseudorandom numbers 1732, 1734
PV_WAVE 1013

Q
QR algorithm 59, 521
double-shifted 530
\(Q R\) decomposition 8, 466, 1582
\(Q R\) factorization 473, 484
quadratic interpolation 782,784 , 786, 789, 792, 796
quadratic polynomial interpolation 649
quadrature formulas 861
quadrature rule 911
quadruple precision 1516
quasi-Monte Carlo 899
quasi-Newton method 1232, 1237, 1279, 1286
quintic polynomial 800

R
radial-basis functions 33
random complex numbers, transforming an array 1089, 1096, 1102
random number generator 1727 , \(1729,1730,1732\)
random number generators 1722 , 1723
random numbers 1714
rank-2k update 1442
rank-k update 1441
rank-one matrix 484, 491, 494
rank-one matrix update 1439,1440
rank-two matrix update 1440
rational weighted Chebyshev approximation 854
READ_MPS 1333, 1343
real numbers, sorting 1677
real periodic sequence 1103, 1106
real sparse symmetric positive definite system 404
real symmetric definite linear system 427, 433
real symmetric positive definite system 176, 181, 194, 199, 300, 303, 313, 315
real symmetric system 209, 212, 219, 221
real triangular system 154
real tridiagonal system 275
REAL types xxii
real vectors 1153, 1163
record keys, sorting 1679
rectangular domain 750
rectangular grid 786, 789, 792, 796
recurrence coefficients 905, 908, 911
reduction
array of black and white 40
regularizing term 47
Reid xix
required arguments xxvi
reserved names 1784
reverse communication 54
ridge regression 64
cross-validation
example 64
Rodrigue 47
row and column pivoting 27, 30
row vector, heavily weighted 35
Runge-Kutta-order method 934
Runge-Kutta-Verner fifth-order method 927
Runge-Kutta-Verner sixth-order method 927

S
ScaLAPACK
contents 1620
data types 1620
definition of library 1619
interface modules 1622
reading utility
block-cyclic distributions 1625, 1636, 1637
scattered data 800
scattered data interpolation 649
Schur form 527, 532
search \(1691,1694,1696\)
second derivative 918
self-adjoint
eigenvalue problem 525
linear system 25
matrix \(17,20,521,523,525,16\)
eigenvalues \(23,520,527,16\)
tridiagonal 20
semi-infinite interval 872
sequence 1129,1135
serial number 1713
Shared-Memory Multiprocessors and xxx
simplex algorithm 1351, 1355
sine 875
sine Fourier coefficients 1129
sine Fourier transform 1126
single precision xix
SINGLE PRECISION options 1743
Single Program, Multiple Data SPMD 1619
singular value decomposition 504
singular value decomposition (SVD) \(36,1612,1613,16\)
singularity 7
singularity points 869
smooth bivariate interpolant 800
smoothing 844
smoothing formulas 32
smoothing spline routines 649
solvable 540
solving
general system 9
linear equations 17
rectangular
least squares 36
system 27
solving linear equations 5
sorting 1679, 1681, 1683, 1684, \(1685,1687,1688,1690,1691\), 1694, 1696
sorting an array, example 1678
Sparse <atrix, Complex
Harwell-Boeing column-oriented sparse form 1533
sparse linear programming 1355
Sparse Matrix Computations, Examples
Plane Poisson Problem with Dirichlet Boundary Conditions 1563
sparse matrix storage mode 1783
Sparse Matrix, Complex 1530, 1532
Accumulate entries of sparse matrix 1535
Collection of Triplets 1532, 1535
Compressed Sparse Column Format 1533
Converstion of Triplets to HarwellBoeing form 1535
Derived types for sparse matrices 1532
Triplets types for sparse matrices 1532
Sparse Matrix, Real 1530
sparse system 364,374
spline approximation 815,824
spline interpolant 711, 720
spline knot sequence 715,718
splines \(649,708,731,732,735,738\)
cubic 647
tensor product 648
square matrices
eigenvalue-eigenvector expansion 523
polar decomposition 39, 48
square root 1757
Sturm-Liouville problem 1066, 1078
subprograms
library xxiii
optional arguments xxvi
SVD 57, 62, 16
SVRGN 1678
symmetric Markowitz strategy 370

\section*{T}
tensor product splines 648
tensor-product B-spline coefficients \(720,725,833,838\)
tensor-product B-spline representation 741, 742, 746, 750, 754, 756, 760, 766
tensor-product spline 741, 742, 746, 750, 754, 756, 760, 766
tensor-product spline approximant 833, 838
tensor-product spline interpolant 725
terminal errors 1763
third derivative 918
time 1706
Timing
Benchmarking
list, parallel codes 8, 9
list,scalar version 4
parallel version 1,5
Toeplitz linear system 423
Toeplitz matrices 8
traceback 1768
transfer 1602
transpose 1553, 1554, 1556
tridiagonal 44
matrix 47
matrix solving, example 53
triple inner product 1428
two-dimensional data fitting 33

\section*{U}
unconstrained minimization 1218
underflow xxiv
uniform \((0,1)\) distribution 1732, 1734
uniform mesh 1059
unitary matrix xxiii
univariate functions 1218
univariate quadrature 860
upper Hessenberg matrix 531
user errors 1763
user interface xix
user-supplied function 918
user-supplied gradient 1259, 1286, 1383
Using LAPACK, LINPACK, and EISPACK xxxii
using library subprograms xxiii

V
Van Loan 12, 20, 30, 35, 59, 62, 64, 521, 525, 530
variable knot B-spline 819
variable order 961
variances 1716
variational equation 53
vectors \(1425,1426,1428,1429\), 1437, 1491, 1493, 1512, 1514
complex 1514
real 1512
version 1713

\section*{W}
workspace allocation 1785, 1786
World Wide Web
URL for ScaLAPACK User's Guide 1620
zero of a real function 1192
zeros of a polynomial 1184, 1186, 1188
zeros of a univariate complex function 1189
zeros of the polynomial 1183```


[^0]:    ${ }^{1}$ The assign-to equality, $a:=b$, used here and below, is read "the expression $b$ is evaluated and then assigned to the location $a$."

[^1]:    ${ }^{2}$ The three-tiered equal sign, used here and below, is read " $a \equiv b$ or $a$ and $b$ are exactly the same object or value."

