

Showing that Android's, Java's and Python's sorting algorithm is broken and fixing it formally

Stijn de Gouw

Jurriaan Rot, Frank de Boer, Richard Bubel, Reiner Hähnle

CWI Amsterdam / SDL Fredhopper

Formal Methods 2015

Oslo, June 22, 2015



<http://www.envisage-project.eu>

Library

Collection of commonly used algorithms that are invoked through a well-defined interface

Library

Collection of commonly used algorithms that are invoked through a well-defined interface

Example: Java standard library functions

Programming to interfaces:

- ▶ Sorting a given array `a`

```
static void sort(Object[] a)
```

- ▶ Searching a value `key` in the array `a`

```
static int binarySearch(Object[] a, Object key)
```

Usability of programming language partially depends on good libraries

Library

Collection of commonly used algorithms that are invoked through a well-defined interface

Example: Java standard library functions

Programming to interfaces:

- ▶ Sorting a given array `a`

```
static void sort(Object[] a)
```

- ▶ Searching a value `key` in the array `a`

```
static int binarySearch(Object[] a, Object key)
```

Usability of programming language partially depends on good libraries

Correctness of library functions is crucial:
used as building blocks in millions of programs

Description

Timsort: a hybrid sorting algorithm (insertion sort + merge sort)
optimized for partially sorted arrays (often encountered in real-world data).

Timsort (I)

Description

Timsort: a hybrid sorting algorithm (insertion sort + merge sort)
optimized for partially sorted arrays (often encountered in real-world data).

Timsort is used in

- ▶ Java (standard library), used by Oracle
- ▶ Python (standard library), used by Google
- ▶ Android (standard library), used by Google
- ▶ Hadoop (Big data), used by Apache, Facebook and Yahoo
- ▶ ... and many more languages / frameworks!

`TimSort.rangeCheck` appeared in court case between Oracle and Google

Timsort (I)

Description

Timsort: a hybrid sorting algorithm (insertion sort + merge sort)
optimized for partially sorted arrays (often encountered in real-world data).

Timsort is used in

- ▶ Java (standard library), used by Oracle
- ▶ Python (standard library), used by Google
- ▶ Android (standard library), used by Google
- ▶ Hadoop (Big data), used by Apache, Facebook and Yahoo
- ▶ ... and many more languages / frameworks!

`TimSort.rangeCheck` appeared in court case between Oracle and Google

Why analyze Timsort?

- ▶ Complex algorithm, widely used
- ▶ Extensively tested + manual code reviews: bugs unlikely!?

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the invariant”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 1 | 1 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the invariant”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 1 | 1 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| |
|---|
| 5 |
|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the invariant”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 1 | 1 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| | |
|---|---|
| 5 | 3 |
|---|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the invariant”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 1 | 1 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| | | |
|---|---|---|
| 5 | 3 | 2 |
|---|---|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the **invariant**”)
 - 1 `runLen[n-2] > runLen[n-1] + runLen[n]`
 - 2 `runLen[n-1] > runLen[n]`

Merging: if **(1) is false** and `runLen[n-2] < runLen[n]`, merge runs at `n-2` and `n-1`, **otherwise at `n-1` and `n`**

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 1 | 1 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| | | |
|---|---|---|
| 5 | 3 | 2 |
|---|---|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the **invariant**”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 0 | 1 | 1 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| | |
|---|---|
| 5 | 5 |
|---|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the **invariant**”)
 - 1 `runLen[n-2] > runLen[n-1] + runLen[n]`
 - 2 `runLen[n-1] > runLen[n]`

Merging: if (1) is false and `runLen[n-2] < runLen[n]`, merge runs at `n-2` and `n-1`, **otherwise at `n-1` and `n`**

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 0 | 0 | 1 | 1 | 4 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| | |
|---|---|
| 5 | 5 |
|---|---|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“the **invariant**”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$

Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 0 | 0 | 1 | 1 | 1 | 2 | 3 | 4 | 4 | 5 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| |
|----|
| 10 |
|----|

Timsort (II)

The algorithm

- ▶ Find next already sorted segment (“runs”) extending to length ≥ 16 with insertion sort.
- ▶ Add length of new run to `runLen` array
- ▶ Merge until last 3 runs satisfy two conditions (“invariant”)
 - 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
 - 2 $\text{runLen}[n-1] > \text{runLen}[n]$



Merging: if (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$, merge runs at $n-2$ and $n-1$, otherwise at $n-1$ and n

- ▶ At the end: merge all runs, resulting in a sorted array

Example, ignoring length ≥ 16 requirement

Input

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 0 | 0 | 1 | 1 | 1 | 2 | 3 | 4 | 4 | 5 |
|---|---|---|---|---|---|---|---|---|---|

runLen

| |
|----|
| 10 |
|----|

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant

Size of `runLen`

- 1 `runLen[n-2] > runLen[n-1] + runLen[n]`
- 2 `runLen[n-1] > runLen[n]`

If the above invariant is true for **all** n and `runLen[n] >= 16`, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52+34+17+16)
- ▶ Runs do not overlap: few runs required to cover input array

```
int stackLen = (len < 120 ? 4 :  
                len < 1542 ? 9 :  
                len < 119151 ? 18 : 39);  
runBase = new int[stackLen];  
runLen = new int[stackLen];
```

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant

Size of `runLen`

- 1 `runLen[n-2] > runLen[n-1] + runLen[n]`
- 2 `runLen[n-1] > runLen[n]`

If the above invariant is true for **all** n and `runLen[n] >= 16`, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant - checking last 3 runs is insufficient

If (1) is false and `runLen[n-2] < runLen[n]`: merge at idx $n-2$ and $n-1$, otherwise merge runs at indices $n-1$ and n

`runLen`

| | | | |
|-----|----|----|----|
| 120 | 80 | 25 | 20 |
|-----|----|----|----|

Breaking the invariant

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant - checking last 3 runs is insufficient

If (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$: merge at idx $n-2$ and $n-1$, otherwise merge runs at indices $n-1$ and n

`runLen`

| | | | | |
|-----|----|----|----|----|
| 120 | 80 | 25 | 20 | 30 |
|-----|----|----|----|----|

Breaking the invariant

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant - checking last 3 runs is insufficient

If **(1) is false** and $\text{runLen}[n-2] < \text{runLen}[n]$: **merge at idx $n-2$ and $n-1$** , otherwise merge runs at indices $n-1$ and n

`runLen`

| | | | | |
|-----|----|----|----|----|
| 120 | 80 | 25 | 20 | 30 |
|-----|----|----|----|----|

Breaking the invariant

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant - checking last 3 runs is insufficient

If (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$: merge at idx $n-2$ and $n-1$, otherwise merge runs at indices $n-1$ and n

`runLen`

| | | | |
|-----|----|----|----|
| 120 | 80 | 45 | 30 |
|-----|----|----|----|

Breaking the invariant

Size of `runLen`

- 1 $\text{runLen}[n-2] > \text{runLen}[n-1] + \text{runLen}[n]$
- 2 $\text{runLen}[n-1] > \text{runLen}[n]$

If the above invariant is true for **all** n and $\text{runLen}[n] \geq 16$, then

- ▶ (reversed) runlengths grow exponentially fast (... 87 52 34 17 16)
- ▶ Runs do not overlap: few runs required to cover input array

Breaking the invariant - checking last 3 runs is insufficient

If (1) is false and $\text{runLen}[n-2] < \text{runLen}[n]$: merge at idx $n-2$ and $n-1$, otherwise merge runs at indices $n-1$ and n

`runLen`

| | | | |
|-----|----|----|----|
| 120 | 80 | 45 | 30 |
|-----|----|----|----|

Wrote program that generates testcase

- ▶ that exploits breaking the invariant, by generating too many “short” runs
- ▶ **Triggers exception: insufficient size for `runLen` to store run lengths**

| Language | Smallest array that triggers error |
|----------|------------------------------------|
| Android | 65.536 (2^{16}) |
| Java | 67.108.864 (2^{26}) |
| Python | 562.949.953.421.312 (2^{49}) |

Most powerful supercomputer (Tianhe-2) has $\approx 2^{50}$ bytes of mem.

Our work (I)

Wrote program that generates testcase

- ▶ that exploits breaking the invariant, by generating too many “short” runs
- ▶ **Triggers exception: insufficient size for `runLen` to store run lengths**

| Language | Smallest array that triggers error |
|----------|------------------------------------|
| Android | 65.536 (2^{16}) |
| Java | 67.108.864 (2^{26}) |
| Python | 562.949.953.421.312 (2^{49}) |

Most powerful supercomputer (Tianhe-2) has $\approx 2^{50}$ bytes of mem.

Provided worst-case analysis of broken version

- ▶ Shows the actual minimally required `runLen.length`

Our work (II)

Fixed the algorithm

- ▶ Check that last 4 runs satisfy invariant
- ▶ Executed existing benchmarks (result: same performance) and unit tests (all passed)

```
1  /** ...
2   * merges adjacent runs until the stack invariants are reestablished:
3   *   1. runLen[i - 3] > runLen[i - 2] + runLen[i - 1]
4   *   2. runLen[i - 2] > runLen[i - 1]
5   */
6  private void mergeCollapse() {
7      while (stackSize > 1) {
8          int n = stackSize - 2;
9          if ( (n >= 1 && runLen[n-1] <= runLen[n] + runLen[n+1])
10             || (n >= 2 && runLen[n-2] <= runLen[n-1] + runLen[n])) {
11              if (runLen[n - 1] < runLen[n + 1])
12                  n--;
13              } else if (runLen[n] > runLen[n + 1]) {
14                  break; // Invariant is established
15              }
16              mergeAt (n);
17          }
18      }
```

Analyzing “Real” Software

“because truly understanding it essentially requires doing a formal correctness proof, it’s difficult to maintain”

“Yet another large mass of difficult code can make for a real maintenance burden after I’m dead”

- Tim Peters on Timsort, python-dev mailing list, 2002

Implementation uses features for performance that complicate analysis:
break statements, low-level bitwise ops., arithmetic overflows

Analyzing “Real” Software

“because truly understanding it essentially requires doing a formal correctness proof, it’s difficult to maintain”

“Yet another large mass of difficult code can make for a real maintenance burden after I’m dead”

- Tim Peters on Timsort, python-dev mailing list, 2002

Implementation uses features for performance that complicate analysis:
break statements, low-level bitwise ops., arithmetic overflows

Mechanically proved fixed version with KeY (Java theorem prover)

- ▶ absence of the bug, and all other run-time exceptions
- ▶ termination
- ▶ this requires: formal specifications for all functions

Method contracts

- ▶ precondition (`requires`): condition on the input
- ▶ postcondition (`ensures`): condition on the output / result

```
1  /*@ requires
2     @   stackSize > 0;
3     @   ensures
4     @   (\forall int i; 0<=i && i<stackSize-2;
5     @       elemInv(runLen, i, 16))
6     @   && elemBiggerThanNext(runLen, stackSize-2)
7     @*/
8  private void mergeCollapse()
```

Method contracts

- ▶ precondition (`requires`): condition on the input
- ▶ postcondition (`ensures`): condition on the output / result

```
1  /*@ requires
2     @   stackSize > 0;
3     @   ensures
4     @   (\forall int i; 0<=i && i<stackSize-2;
5     @     elemInv(runLen, i, 16))
6     @   && elemBiggerThanNext(runLen, stackSize-2)
7     @*/
8  private void mergeCollapse()
```

Class Invariant

Property that all instances of a class must satisfy before and after **every** method (call)

- ▶ Can be **assumed** in method precondition
- ▶ Must be **established** at all call sites and method postcondition

Class Invariant (simplified)

```
1  /*@ invariant
2  @    runBase.length == runLen.length
3  @ && (a.length < 120 ==> runLen.length==4)
4  @ && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5  @ && (a.length >= 1542 && a.length<119151 ==> runLen.length==18)
6  @ && (a.length >= 119151 ==> runLen.length==39)
7  @ && (0 <= stackSize && stackSize <= runLen.length)
8  @ && (\forallall int i; 0<=i && i<stackSize-4;
9  @    elemInv(runLen, i, 16))
10 @ && (elemLargerThanBound(runBase, 0, 0))
11 @ && (\forallall int i; 0<=i && i<stackSize-1;
12 @    runBase[i] + runLen[i] == runBase[i+1]);
13 @*/
```

| Name | Definition |
|---|---|
| <code>elemBiggerThanNext2(arr, idx)</code> | $(0 \leq idx \wedge idx + 2 < arr.length) \rightarrow arr[idx] > arr[idx + 1] + arr[idx + 2]$ |
| <code>elemBiggerThanNext(arr, idx)</code> | $0 \leq idx \wedge idx + 1 < arr.length \rightarrow arr[idx] > arr[idx + 1]$ |
| <code>elemLargerThanBound(arr, idx, v)</code> | $0 \leq idx < arr.length \rightarrow arr[idx] \geq v$ |
| <code>elemInv(arr, idx, v)</code> | $elemBiggerThanNext2(arr, idx) \wedge elemBiggerThanNext(arr, idx) \wedge elemLargerThanBound(arr, idx, v)$ |

Class Invariant (simplified)

```
1  /*@ invariant
2  @   runBase.length == runLen.length
3  @   && (a.length < 120 ==> runLen.length==4)
4  @   && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5  @   && (a.length >= 1542 && a.length < 119151 ==> runLen.length==18)
6  @   && (a.length >= 119151 ==> runLen.length==39)
7  @   && (0 <= stackSize && stackSize <= runLen.length)
8  @   && (\forall int i; 0<=i && i<stackSize-4;
9  @     elemInv(runLen, i, 16))
10 @   && (elemLargerThanBound(runBase, 0, 0))
11 @   && (\forall int i; 0<=i && i<stackSize-1;
12 @     runBase[i] + runLen[i] == runBase[i+1]);
13 @*/
```

Length of `runlen` in terms of input length

Class Invariant (simplified)

```
1  /*@ invariant
2  @    runBase.length == runLen.length
3  @ && (a.length < 120 ==> runLen.length==4)
4  @ && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5  @ && (a.length >= 1542 && a.length<119151 ==> runLen.length==18)
6  @ && (a.length >= 119151 ==> runLen.length==39)
7  @ && (0 <= stackSize && stackSize <= runLen.length)
8  @ && (\forallall int i; 0<=i && i<stackSize-4;
9  @    elemInv(runLen, i, 16))
10 @ && (elemLargerThanBound(runBase, 0, 0))
11 @ && (\forallall int i; 0<=i && i<stackSize-1;
12 @    runBase[i] + runLen[i] == runBase[i+1]);
13 @*/
```

Bounds on stackSize (in-use part of runLen)

Class Invariant (simplified)

```
1 /*@ invariant
2 @   runBase.length == runLen.length
3 @ && (a.length < 120 ==> runLen.length==4)
4 @ && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5 @ && (a.length >= 1542 && a.length < 119151 ==> runLen.length==18)
6 @ && (a.length >= 119151 ==> runLen.length==39)
7 @ && (0 <= stackSize && stackSize <= runLen.length)
8 | @ && (\forall int i; 0 <= i && i < stackSize-4;
9 | @   elemInv(runLen, i, 16))
10 | @ && (elemLargerThanBound(runBase, 0, 0))
11 | @ && (\forall int i; 0 <= i && i < stackSize-1;
12 | @   runBase[i] + runLen[i] == runBase[i+1]);
13 | @*/
```

All but the last 4 runs satisfy the invariant while merging

Class Invariant (simplified)

```
1 /*@ invariant
2 @      runBase.length == runLen.length
3 @ && (a.length < 120 ==> runLen.length==4)
4 @ && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5 @ && (a.length >= 1542 && a.length < 119151 ==> runLen.length==18)
6 @ && (a.length >= 119151 ==> runLen.length==39)
7 @ && (0 <= stackSize && stackSize <= runLen.length)
8 @ && (\forall int i; 0<=i && i<stackSize-4;
9 @      elemInv(runLen, i, 16))
10 @ && (elemLargerThanBound(runBase, 0, 0))
11 @ && (\forall int i; 0<=i && i<stackSize-1;
12 @      runBase[i] + runLen[i] == runBase[i+1]);
13 @*/
```

First run starts at non-negative array index

Class Invariant (simplified)

```
1  /*@ invariant
2  @    runBase.length == runLen.length
3  @ && (a.length < 120 ==> runLen.length==4)
4  @ && (a.length >= 120 && a.length < 1542 ==> runLen.length==9)
5  @ && (a.length >= 1542 && a.length<119151 ==> runLen.length==18)
6  @ && (a.length >= 119151 ==> runLen.length==39)
7  @ && (0 <= stackSize && stackSize <= runLen.length)
8  @ && (\forall int i; 0<=i && i<stackSize-4;
9  @    elemInv(runLen, i, 16))
10 @ && (elemLargerThanBound(runBase, 0, 0))
11 @ && (\forall int i; 0<=i && i<stackSize-1;
12 @    runBase[i] + runLen[i] == runBase[i+1]);
13 @*/
```

There are no gaps between consecutive runs

Loop Invariant (simplified)

```
1 /*@ loop_invariant
2   @ \forall int i; 0<=i && i<stackSize-4;
3   @     elemInv(runLen, i, 16);
4   @*/
```

The main verif. condition (simplified)

```
( loop-inv && n==stackSize-2 && n >= 0
  && n>=1 ==> runLen[n-1] > runLen[n] + runLen[n+1]
  && n>=2 ==> runLen[n-2] > runLen[n-1] + runLen[n]
  && runLen[n] > runLen[n+1]
)
```

==> ensures(mergeCollapse)

Recall that `ensures(mergeCollapse)` is (substituting `stackSize-2==n`):

```
(\forall int i; 0<=i && i<n; elemInv(runLen, i, 16))
&& elemBiggerThanNext(runLen, n)
```

pushRun contract (simplified)

```
1  /*@ normal_behavior
2  @ requires
3  @   (runLen > 0 && runBase >= 0)
4  @   && (stackSize > 0 ==> runBase ==
5  @     this.runBase[stackSize-1]+this.runLen[stackSize-1])
6  @   && (runLen + runBase <= a.length)
7  @   && (\forall int i; 0<=i && i<stackSize-2;
8  @     elemInv(this.runLen, i, 16))
9  @   && elemBiggerThanNext(this.runLen, stackSize-2)
10 @   && elemLargerThanBound(this.runLen, stackSize-1, 16)
11 @ ensures
12 @   this.runBase[\old(stackSize)] == runBase
13 @   && this.runLen[\old(stackSize)] == runLen
14 @   && stackSize == \old(stackSize)+1;
15 @*/
16 private void pushRun(int runBase, int runLen) {
17     this.runBase[stackSize] = runBase;
18     this.runLen[stackSize] = runLen;
19     stackSize++;
20 }
```

pushRun contract (simplified)

```
1  /*@ normal_behavior
2  @ requires
3  @ (runLen > 0 && runBase >= 0)
4  @ && (stackSize > 0 ==> runBase ==
5  @     this.runBase[stackSize-1]+this.runLen[stackSize-1])
6  @ && (runLen + runBase <= a.length)
7  @ && (\forall int i; 0<=i && i<stackSize-2;
8  @     elemInv(this.runLen,i,16))
9  @ && elemBiggerThanNext(this.runLen, stackSize-2)
10 @ && elemLargerThanBound(this.runLen, stackSize-1, 16)
11 @ ensures
12 @     this.runBase[\old(stackSize)] == runBase
13 @ && this.runLen[\old(stackSize)] == runLen
14 @ && stackSize == \old(stackSize)+1;
15 @*/
16 private void pushRun(int runBase, int runLen) {
17     this.runBase[stackSize] = runBase;
18     this.runLen[stackSize] = runLen;
19     stackSize++;
20 }
```

The new run has positive length and starts directly after the last run

pushRun contract (simplified)

```
1  /*@ normal_behavior
2  @ requires
3  @   (runLen > 0 && runBase >= 0)
4  @ && (stackSize > 0 ==> runBase ==
5  @   this.runBase[stackSize-1]+this.runLen[stackSize-1])
6  @ && (runLen + runBase <= a.length)
7  @ && (\forall int i; 0<=i && i<stackSize-2;
8  @   elemInv(this.runLen,i,16))
9  @ && elemBiggerThanNext(this.runLen, stackSize-2)
10 @ && elemLargerThanBound(this.runLen, stackSize-1, 16)
11 @ ensures
12 @   this.runBase[\old(stackSize)] == runBase
13 @ && this.runLen[\old(stackSize)] == runLen
14 @ && stackSize == \old(stackSize)+1;
15 @*/
16 private void pushRun(int runBase, int runLen) {
17     this.runBase[stackSize] = runBase;
18     this.runLen[stackSize] = runLen;
19     stackSize++;
20 }
```

The new run cannot extend beyond length of the input array

pushRun contract (simplified)

```
1  /*@ normal_behavior
2  @ requires
3  @   (runLen > 0 && runBase >= 0)
4  @ && (stackSize > 0 ==> runBase ==
5  @   this.runBase[stackSize-1]+this.runLen[stackSize-1])
6  @ && (runLen + runBase <= a.length)
7  @ && (\forall int i; 0<=i && i<stackSize-2;
8  @   elemInv(this.runLen,i,16))
9  @ && elemBiggerThanNext(this.runLen, stackSize-2)
10 @ && elemLargerThanBound(this.runLen, stackSize-1, 16)
11 @ ensures
12 @   this.runBase[\old(stackSize)] == runBase
13 @ && this.runLen[\old(stackSize)] == runLen
14 @ && stackSize == \old(stackSize)+1;
15 @*/
16 private void pushRun(int runBase, int runLen) {
17     this.runBase[stackSize] = runBase;
18     this.runLen[stackSize] = runLen;
19     stackSize++;
20 }
```

The invariant is satisfied by all runs

pushRun contract (simplified)

```
1  /*@ normal_behavior
2  @ requires
3  @   (runLen > 0 && runBase >= 0)
4  @ && (stackSize > 0 ==> runBase ==
5  @   this.runBase[stackSize-1]+this.runLen[stackSize-1])
6  @ && (runLen + runBase <= a.length)
7  @ && (\forall int i; 0<=i && i<stackSize-2;
8  @   elemInv(this.runLen, i, 16))
9  @ && elemBiggerThanNext(this.runLen, stackSize-2)
10 @ && elemLargerThanBound(this.runLen, stackSize-1, 16)
11 @ ensures
12 | @   this.runBase[\old(stackSize)] == runBase
13 | @ && this.runLen[\old(stackSize)] == runLen
14 | @ && stackSize == \old(stackSize)+1;
15 @*/
16 private void pushRun(int runBase, int runLen) {
17     this.runBase[stackSize] = runBase;
18     this.runLen[stackSize] = runLen;
19     stackSize++;
20 }
```

The new run is stored at index stackSize-1

No ArrayIndexOutOfBoundsException if

```
requires (pushRun) && cl. invariant ==> stackSize < len.length
```

pushRun main verification condition

No ArrayIndexOutOfBoundsException if

`requires(pushRun) && cl.invariant ==> stackSize < len.length`

Proof.

Note first: `cl.invariant` \rightarrow `stackSize` \leq `len.length`.

Assume by contradiction that `stackSize` = `len.length` and do a case distinction on `a.length`. We treat `a.length` \leq 119:

- 1 `len.length` = 4 (from `cl.invariant`, ln 3).
- 2 Abbreviate `len[0]+...+len[3]` = SUM, then (`pushRun` ln 7–10)
`len[3]` \geq 16, `len[2]` \geq 17, `len[1]` \geq 34 and `len[0]` \geq 52.
Therefore: SUM \geq 16+17+34+52=119
- 3 `base[3]` + `len[3]` = `base[0]` + SUM (from `cl.invariant`, ln 11–12)
- 4 Previous line, with `pushRun` ln 4–5 implies:
`runBase` + `runLen` = `base[0]` + SUM + `runLen`
- 5 But `base[0]` \geq 0 (`cl.invariant` ln 10) and `runLen` $>$ 0 (`pushRun` ln 3),
contradicting `runBase` + `runLen` \leq 119 (`pushRun` ln 6)

One proof step in KeY

The screenshot displays the KeY IDE interface. On the left, the 'Proof Search Strategy' pane shows a tree of goals and tactics. The selected goal is `CUT: forall int i. (0 <= i & i < self.stackSize & self.runLen[i] <= self.runBase.length) => self.runLen[i] <= self.runBase.length`. The 'Inner Node' pane on the right shows the corresponding Java code for this goal, including a `forall` loop and a `measureByEmpty` tactic. The code is as follows:

```
0 <= i,0,
i,0 < self.stackSize,
i,0 = -1 + self.stackSize,
self.stackSize <= 39,
self.runLen.length <= 39,
self.c.<created> = TRUE,
self.runBase.<created> = TRUE,
self.a.<created> = TRUE,
self.tmp.<created> = TRUE,
wellFormed(heap),
self.<created> = TRUE,
java.util.TimerSort::exactInstance(self) = TRUE,
measureByEmpty,
elemInv(heap, self.runLen, -4 + self.stackSize, javaDivInt(32, 2)),
self.runLen.<created> = TRUE,
elemBiggerThanNext(heap, self.runLen, -3 + self.stackSize),
self.runLen.length >= 1,
self.stackSize >= 1,
self.runBase.length >= 0,
self.runBase.length <= 2147483647,
self.runLen.length = self.runBase.length,
self.runLen.length <= 2147483647,
self.a.length <= 2147483647,
self.a.length >= 0,
bound(int i);(0, self.stackSize, self.runLen[i_1]) <= self.a.length + self.runBase[0] * -1,
self.stackSize <= self.runLen.length,
forall int i;
{
  elemBiggerThanNext(heap, self.runLen, i)
  & elemBiggerThanNext(heap, self.runLen, i)
  & !if (self.runLen = null & 0 <= i & i < self.runLen.length) \then (self.runLen[i] >= javaDivInt(32, 2)) \else (true)
  | i < -1
  | i >= -4 + self.stackSize,
  \if (self.runLen = null & 0 <= -4 + self.stackSize & -4 + self.stackSize < self.runLen.length)
  \then (-4 + self.stackSize + 1 < self.runLen.length >=> self.runLen[-4 + self.stackSize + 1])
  \else (true),
  \if (self.runLen = null & 0 <= -3 + self.stackSize & -3 + self.stackSize < self.runLen.length)
  \then (self.runLen[-3 + self.stackSize] >= javaDivInt(32, 2))
  \else (true),
  \if (self.runLen = null & 0 <= -2 + self.stackSize & -2 + self.stackSize < self.runLen.length)
  \then (self.runLen[-2 + self.stackSize] >= javaDivInt(32, 2))
  \else (true),
  \if (self.runLen = null & 0 <= -1 + self.stackSize & -1 + self.stackSize < self.runLen.length) \then (self.runLen[-1 + self.stackSize] >= 1) \else (true),
  elemLargerThanBound(heap, self.runBase, 0, 0),
  java.lang.Object[]::exactInstance(self.tmp) = TRUE,
  self.mingGallo <= 2147483647,
  self.mingGallo >= -2147483648
}
==>
self.runLen[i_0] >= 0,
self.tmp = null,
self.a = null,
self.tmp = self.a,
self.runBase = null,
self.c = null,
self.runLen = self.runBase,
```

At the bottom of the IDE, a red banner displays the message: **KeY Hint: Running out of memory? Use the command line option -j-Xmx next time you start KeY.**

Proof Stats - summary

| | # Rule Apps | # Interactive | LoSpec | LoC |
|---------------|-------------|---------------|--------|-----|
| total | 2.211.263 | 5.029 | 334 | 333 |
| pushRun | 26.248 | 94 | 17 | 5 |
| mergeCollapse | 415.133 | 1.529 | 47 | 13 |

Proof Stats - summary

| | # Rule Apps | # Interactive | LoSpec | LoC |
|---------------|-------------|---------------|--------|-----|
| total | 2.211.263 | 5.029 | 334 | 333 |
| pushRun | 26.248 | 94 | 17 | 5 |
| mergeCollapse | 415.133 | 1.529 | 47 | 13 |

Evaluation of the problem

- ▶ Bug unlikely to be encountered by accident
- ▶ Possible security hazard: bug may be exploitable in DoS attack
- ▶ Extensive testing unable to expose bug:
input size too large, structure too complex
- ▶ Manual code reviews (Google) unable to expose bug
- ▶ Core libraries in widely used languages can contain subtle bugs undetected for years

- ▶ Scientific paper (CAV 2015), articles (ERCIM, Bits & Chips)
- ▶ Published blog post `viewed 361274 times`

Responses: general public

- ▶ Scientific paper (CAV 2015), articles (ERCIM, Bits & Chips)
- ▶ Published blog post **viewed 361274 times**


 **Joshua Bloch**
@joshbloch [Follow](#)

Congratulations to Stijn de Gouw et al. for finding and fixing a bug in TimSort using formal methods!
envisage-project.eu/proving-android...

12:33 AM - 25 Feb 2015

126 RETWEETS 65 FAVORITES

JetBrains IntelliJ IDEA Blog (March 10, 2015):
"The first of its kind, this result is both an important development for the Java community and a proof of concept for the feasibility of formal verification and automated theorem proving. Perhaps more importantly, the tool used to detect and identify this bug is completely open source and **available to try yourself!**"

 **Damian Gryski**
@dgryski [Follow](#)

These #golang timsorts
go-search.org/search?q=timso... probably need to be fixed envisage-project.eu/proving-android...

5:17 PM - 24 Feb 2015

7 RETWEETS 9 FAVORITES

Messages (8)

[msg236533](#) - (view)

Author: [Tim Peters \(tim.peters\)](#)


Date: 2015-02-24 19:32

Some researchers found an error in the logic of `merge_collapse`, explained here, and with corrected code shown in section 3.2:

<http://envisage-project.eu/proving-android-java-and-python-sorting-algorithm-is-broken-and-how-to-fix-it/>

This affects all current versions of Python. However, I marked the priority "low" because, as the article also notes, there's currently no machine in existence with enough memory to hold an array large enough for a contrived input to trigger an overflow of the pending-runs stack.

It should be fixed anyway, and their suggested fix looks good to me.

 **Peter O'Hearn**
February 24 · [@](#)

An impressive use of formal methods. Congrats to the team behind this for their dedicated and valuable work.
http://envisage-project.eu/proving-android-java-and-python...

```
mergeCollapse && n = sta
mLen[n-1] > runLen[n] +
mLen[n-2] > runLen[n-1]
runLen[n] > runLen[n+1]

==> \ensur
```


Proving that Android's, Java's and Python's sorting algorithm is broken (and showing how to fix...

10.03.2015 08:17

[@baso Developer](#) · [Vorlage](#) | [Nächste](#) >

Fehler in Standardsortieralgorithmus mit formalen Methoden aufgedeckt

 [verfassen](#) / 193 Downloads



Android, Java und Groovy nutzen alle den TimSort-Algorithmus. Informatiker eines Verbundunternehmens konnten mit Hilfe eines von ihnen entwickelten Tools

 **Pascal Hartig** @passy [Follow](#)

The bug found in the Python and Java implementations of Timsort was also present in this Haskell version: github.com/tmjb/haskell-...

 **Pascal Hartig** @passy [Follow](#)

I wonder if the invariant could have been encoded in the type system, or if LiquidHaskell, Iris, Agda impls would have exposed this.

4:23 PM - 26 Feb 2015

Java

- ▶ Submitted bug report to Java issue tracker

Java

- ▶ Submitted bug report to Java issue tracker
- ▶ Bug was previously found and “fixed” by increasing `runLen.length`

```
int stackLen = (len < 120 ? 5 :  
                len < 1542 ? 10 :  
                len < 119151 ? 19 24  
                : 40);  
runBase = new int[stackLen];  
runLen = new int[stackLen];
```

Responses: developer communities

Java

- ▶ Submitted bug report to Java issue tracker
- ▶ Bug was previously found and “fixed” by increasing `runLen.length`
- ▶ Bug now fixed by further increasing `runLen.length` based on worst-case analysis

Discussion on OpenJDK mailing list

Stack length increased previously by JDK-8011944 was insufficient for some cases. Please review and push - Lev Priima, 11 Feb 2015

```
int stackLen = (len < 120 ? 5 :  
                len < 1542 ? 10 :  
                len < 119151 ? 24 :  
                40 49 );  
runBase = new int[stackLen];  
runLen = new int[stackLen];
```

Java

- ▶ Submitted bug report to Java issue tracker
- ▶ Bug was previously found and “fixed” by increasing `runLen.length`
- ▶ Bug now fixed by further increasing `runLen.length` based on worst-case analysis

Discussion on OpenJDK mailing list

Stack length increased previously by JDK-8011944 was insufficient for some cases. Please review and push
- Lev Priima, 11 Feb 2015

Hi Lev, The fix looks fine. Did you consider the improvements suggested in the paper to reestablish the invariant?
- Roger Riggs, Feb 11, 2015

Just briefly looked at it, w/o evaluating formal proof ...
- Lev Priima, Feb 11, 2015

Responses: developer communities

Java

- ▶ Submitted bug report to Java issue tracker
- ▶ Bug was previously found and “fixed” by increasing `runLen.length`
- ▶ Bug now fixed by further increasing `runLen.length` based on worst-case analysis
- ▶ Purported class invariant still broken
- ▶ Not amenable to mechanic verification

Python

- ▶ Bug report filed by Tim Peters
- ▶ Bug fixed by checking last 4 runs (verified version)

Android

- ▶ No bug report or fix so far

Formal methods work!

Useful links

Blog post

<http://tinyurl.com/timsort-bug>

Website with full paper, test programs and proofs

<http://www.envisage-project.eu/timsort-specification-and-verification>

KeY (Java theorem prover)

<http://www.key-project.org>

Timsort description

<http://bugs.python.org/file4451/timsort.txt>

OpenJDK dev discussion

<http://mail.openjdk.java.net/pipermail/core-libs-dev/2015-February/thread.html#31405>