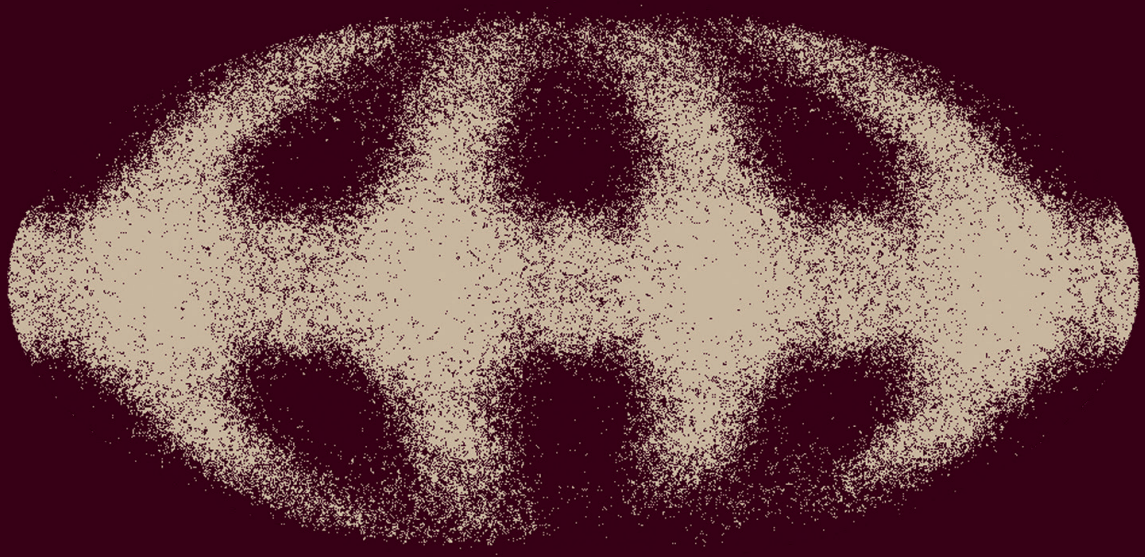
Universität Münster

IceCube Gen2

IceCube Upgrade



Development of an optical module for IceCube-Gen2 and sensitivity studies on NGC 1068 with IceCube Upgrade

Universität Münster
Institut für Kernphysik

Dissertation Markus Dittmer 2026







Experimentelle Physik

Development of an optical module for IceCube-Gen2 and sensitivity studies on NGC 1068 with IceCube Upgrade

Inaugural-Dissertation
zur Erlangung des Doktorgrades Dr. rer. nat.
der Naturwissenschaften im Fachbereich Physik
der Mathematisch-Naturwissenschaftlichen Fakultät
der Universität Münster

vorgelegt von

Markus Dittmer

aus Münster

–2026–

Dekan:	Prof. Dr. Uwe Thiele
Erster Gutachter:	Prof. Dr. Alexander Kappes
Zweiter Gutachter:	Prof. Dr. Christian Weinheimer
Tag der mündlichen Prüfung:	30.09.2026
Tag der Promotion:	30.09.2026

Abstract

The IceCube Neutrino Observatory's recent identification of the active galactic nucleus NGC 1068 as a steady neutrino source highlights the potential of high-energy neutrino astronomy. To improve the detector's sensitivity at lower energies and enhance the reconstruction of such sources, the recent detector addition IceCube Upgrade deployed densely instrumented strings featuring novel optical modules, specifically the multi-PMT digital optical module (mDOM) along with a precursor module, the long optical module (LOM). Looking beyond, the planned IceCube-Gen2 extension will substantially expand the detector's optical array with these next-generation optical modules, requiring extensive prototyping and scalable production methods.

This thesis details specific hardware and analytical advancements driving these future extensions. For the IceCube-Gen2 optical modules, it presents a scalable mechanical prototyping approach based on pre-cast silicone gel pads to simplify production. This hardware development is underpinned by comprehensive characterizations of photomultiplier tube (PMT) models specifically developed for these modules, also verifying compliance with the collaboration's specifications. To quantify the expected detector improvements, Geant4 simulations are used to benchmark the performance of the new module geometries against legacy designs, alongside detailed modeling of the intrinsic scintillation background originating from the pressure vessel glass. On the analysis side, this work presents a time-integrated unbinned maximum likelihood point-source analysis of NGC 1068. Using IceCube Upgrade simulation datasets, the study quantifies the projected sensitivity gains and the potential of IceCube Upgrade to improve the resolution of obscured active galactic nuclei.

Zusammenfassung

Die kürzliche Identifizierung des aktiven Galaxienkerns NGC 1068 als stetige Neutrinoquelle durch das IceCube Neutrino Observatorium unterstreicht das Potenzial der Hochenergie-Neutrinoastronomie. Um solche Quellen besser aufzulösen und die Energieschwelle des Detektors zu senken, verwendet das kürzlich installierte IceCube Upgrade dicht instrumentierte Stränge mit verbesserten optischen Modulen: das multi-PMT digitale optische Modul (mDOM) und, in Form eines Vorläufermoduls, das lange optische Modul (LOM). Die geplante IceCube-Gen2-Erweiterung wird das optische Detektorarray erheblich vergrößern, was eine sorgfältige Entwicklung und Erprobung von Prototypen sowie skalierbare Produktionsverfahren für diese optischen Module der nächsten Generation erfordert.

Diese Dissertation beschreibt Fortschritte in der Hardwareentwicklung und Datenanalyse, die diese zukünftigen Erweiterungen vorantreiben. Für die optischen Module von IceCube-Gen2 wird ein skalierbares Konzept für die mechanische Entwicklung von Prototypen vorgestellt, das vorgefertigte Silikongelpads zur Vereinfachung der Produktion nutzt. Diese Hardwareentwicklung wird durch umfassende Charakterisierungen speziell entwickelter Photomultiplier-Modelle untermauert, wodurch gleichzeitig die Einhaltung der Spezifikationen der Kollaboration nachgewiesen wird. Zur Quantifizierung der erwarteten Detektorverbesserungen werden Geant4-Simulationen eingesetzt, um die optischen Sensitivitäten der neuen Modulgeometrien mit bisherigen Designs zu vergleichen. Ergänzend erfolgt eine detaillierte Modellierung des intrinsischen Szintillationsuntergrunds, der vom Glas des Druckbehälters ausgeht. Analytisch präsentiert diese Arbeit eine zeitintegrierte, ungebinnte Maximum-Likelihood-Punktquellenanalyse von NGC 1068. Unter Verwendung von Simulationsdatensätzen von IceCube Upgrade quantifiziert die Studie die prognostizierten Sensitivitätsverbesserungen und demonstriert die Fähigkeit von IceCube Upgrade, staubverdunkelte aktive Galaxienkerne präziser zu untersuchen.

Contents

1 Introduction **1**

I Foundation **3**

2 The High-Energy Universe and Neutrino Astronomy **5**
2.1 High-Energy Cosmic Rays 5
2.1.1 Acceleration Mechanisms and Propagation 7
2.1.2 Atmospheric Neutrinos and Muons 9
2.2 High-Energy Neutrino Sources 10
2.2.1 Neutrino Production Mechanisms 11
2.2.2 Active Galactic Nuclei 13
2.2.3 Other Source Candidates 14
2.3 Propagation of Astrophysical Neutrinos to Earth 16
2.3.1 Neutrino Oscillations 17
2.3.2 Flavor Composition at Earth 19

3 Neutrino Detection with the IceCube Neutrino Observatory **22**
3.1 Neutrino Interactions and Cross Sections 22
3.1.1 Interaction Channels 22
3.1.2 Cross Section and Energy Dependence 23
3.2 Neutrino Propagation through the Earth 24
3.2.1 The GeV Regime: Oscillations in Matter 24
3.2.2 The TeV–PeV Regime: Earth Absorption 25
3.3 Cherenkov Detector Operation Principle 26
3.4 The IceCube Neutrino Observatory 27
3.5 Event Morphology 29
3.6 IceCube Upgrade 32
3.7 IceCube-Gen2 35

4 Photomultiplier Tubes **38**
4.1 Operation Principle 38
4.2 Photocathode Properties 41
4.3 Gain and Charge Spectrum 42
4.4 Background 45
4.4.1 In Darkness (uncorrelated Background) 45
4.4.2 Signal-induced (correlated Background) 46

II Developing the optical Sensor for the next Generation Neutrino Telescope at the South Pole 47

5 Towards the IceCube-Gen2 optical Module 49

6 Photomultiplier Tube Investigation 53
6.1 Data Acquisition . . . 54
6.2 Temperature-dependent Characterization . . . 55
6.2.1 Gain Calibration and Pulse Shape Parameters . . . 55
6.2.2 Timing Resolution and Pulse Types Probabilities . . . 58
6.2.3 Dark Rates . . . 60
6.3 Quantum Efficiency . . . 62
6.4 Performance Metrics for IceTray . . . 64
6.5 Scans across the photosensitive Area . . . 71
6.6 Summary . . . 76

7 Mechanical Prototyping and Component Validation 78
7.1 Prototyping Efforts . . . 78
7.1.1 Gel Pad Casting . . . 78
7.1.2 Mechanical Prototype for Gen2-OM . . . 81
7.2 Component Testing . . . 85
7.2.1 Pressure Vessel . . . 85
7.2.2 Gel Pad Efficiency . . . 87
7.2.3 Influence of metallic Holding Structure on PMT Noise . . . 89
7.3 Summary and current Status . . . 90

8 Optical Module Performance Simulations 92
8.1 Geant4 Simulation Framework . . . 92
8.2 Effective Area . . . 95
8.2.1 Angular Acceptance . . . 96
8.2.2 Shadowing Losses . . . 99
8.2.3 In-ice spectral Performance . . . 100
8.3 Summary and Outlook . . . 103

9 Optical Module Background Estimation 104
9.1 Underlying Principle . . . 104
9.2 Isotope Concentrations in the Vessel Glass . . . 105
9.2.1 Concentration in mDOM Vessels . . . 106
9.2.2 Components of borosilicate Glass . . . 110
9.3 Scintillation Spectrum . . . 113
9.3.1 Experimental Setup . . . 113
9.3.2 Corrections . . . 115
9.3.3 Corrected Spectrum . . . 118
9.4 Scintillation Yield . . . 121
9.5 Simulation for expected in-ice Background Rates . . . 123
9.6 Summary and Outlook . . . 133

III IceCube Upgrade Sensitivity Study on Neutrino Emission from active Galaxy NGC 1068 135

10 NGC 1068: A Key to understanding high-energy Astrophysics 137
10.1 Astrophysical Insights and high-energy Particle Acceleration . . . 137
10.2 The strategic Value of NGC 1068 in IceCube and multi-Messenger Astrophysics . . . 139

11 Point Source Analysis Method 141
11.1 The unbinned Likelihood Function . . . 141
11.2 Hypothesis Testing and Test Statistic . . . 144
11.3 Source Injection, Sensitivity and Discovery Potential . . . 145

12 Analysis Framework 148
12.1 Event Samples . . . 148
12.1.1 IceCube DeepCore-based Baseline Dataset (GRECO) . . . 149
12.1.2 IceCube Upgrade Simulation (IC93) . . . 150
12.2 Detector Performance Metrics . . . 151
12.3 Analysis Setup . . . 159
12.3.1 Overview . . . 159
12.3.2 Dataset Generation and statistical Framework . . . 160
12.3.3 Analysis Cuts . . . 161

13 Results 166
13.1 Projected Upgrade Performance across the Sky . . . 166
13.2 Sensitivity Improvements for NGC 1068 . . . 168
13.2.1 Temporal Evolution and accumulated Gain . . . 168
13.2.2 Sensitive Energy Range . . . 170
13.2.3 Physics Reach: spectral Constraints . . . 171
13.3 Applicability to transient Searches . . . 172
13.4 Additional Validations . . . 174
13.4.1 Fit Bias . . . 174
13.4.2 Demonstration of spatial Resolution . . . 176
13.5 Deployment Constraints and the IC91 Approximation . . . 177
13.6 Summary . . . 179

14 Summary and Outlook 181

Bibliography 183

A Scaling towards Mass Production 200

B Supplementary optical Module Performance Comparison 203

C PMT Manufacturer Data 209

D Supplementary Information on Module Background 212

Acknowledgements 215

Glossary

ADC Analog-to-Digital Converter.
AGN Active Galactic Nucleus.
AP1 Afterpulsing type 1.
AP2 Afterpulsing type 2.
BDT Boosted Decision Tree.
CAD Computer-Aided Design.
CC Charged Current, weak interaction mediated by the W-boson.
CDF Cumulative Distribution Function.
CE Collection Efficiency.
CL Confidence Level.
CP Charge-Parity (as in δ_{CP} violation).
CR Cosmic Ray.
DAQ Data Acquisition.
DC Direct Current.
DE Detection Efficiency.
DIS Deep Inelastic Scattering.
DOM Digital Optical Module.
DVT Design Verification Testing.
D-Egg Dual optical sensors in an Ellipsoid Glass for Gen2.
FAT Final Acceptance Testing.
FEP Full-Energy Peak.
FT Fall Time.
FWHM Full Width at Half Maximum.
GNN Graph Neural Network.
GRECO GeV Reconstructed Events with Containment for Oscillations.
GRL Good Run List.
Gen2-OM Optical module proposed for IceCube-Gen2.
GRB Gamma-Ray Burst.
GZK Greisen-Zatsepin-Kuzmin.
HA coating Hamamatsu Photonics insulating coating surrounding a PMT.
HECR High-Energy Cosmic Ray.
HESE High Energy Starting Event.
HV High Voltage.
IC Inverse Compton (scattering).
IC91 IceCube Upgrade simulation with two strings removed from the IC93 geometry.
IC93 IceCube Upgrade simulation based on the initially planned seven-string geometry.
KDE Kernel Density Estimation.
LC Local Coincidence.
LED Light-Emitting Diode.
LMP Linear Motion Potentiometer.
LOM Long Optical Module (precursor to Gen2-OM).
MC Monte Carlo.
mDOM Multi-PMT Digital Optical Module.
MJD Modified Julian Date.
MPE Multiple Photoelectrons.
NC Neutral Current.
NGC 1068 Entry in the New General Catalogue (Messier 77), also known as the Squid Galaxy.
OM Optical Module.
PCA Penetrator Cable Assembly.
PDF Probability Density Function.
PE Photoelectron.
PMNS Pontecorvo-Maki-Nakagawa-Sakata (matrix).
PMT Photomultiplier Tube.
PSF Point Spread Function.
QE Quantum Efficiency.
QUESO Quick Upgrade Event Selection for Oscillations (internal name for IC93).
R15458-02/20 Hamamatsu PMT type used in mDOMs.
R16293-01 Hamamatsu PMT type used in LOMs/Gen2-OMs.
R5912-100-70 Hamamatsu PMT type used in D-Eggs.
R7081-02 Hamamatsu PMT type used in DOMs.
RA Right Ascension.
ROI Region Of Interest.
RT Rise Time.
SMBH Supermassive Black Hole.
SN Serial Number.
SNR Signal-to-Noise Ratio.
SPE Single Photoelectron.
TIR Total Internal Reflection.
TS Test Statistic.
TT Transit Time.
TTS Transit Time Spread.
UHE Ultra-High-Energy.
UHECR Ultra-High-Energy Cosmic Ray.
UV Ultraviolet.
Vitrovex Brand of borosilicate glass by Nautilus.

1 Introduction

Astroparticle physics is an intriguing field driven by curiosity to understand the cosmos, although explaining it to friends and family often makes it seem abstract and far removed from everyday life. From humanity's place in the solar system to the investigation of violent phenomena in the universe using increasingly complex particle detectors, the field seeks to answer some of the most fundamental questions about nature. In recent decades, it has entered the era of multi-messenger astronomy, in which cosmic accelerators, including active galactic nuclei that produce a diverse array of high-energy particles, are studied collaboratively. This approach leverages a wide range of specialized detectors spanning the entire electromagnetic spectrum, alongside neutrinos, cosmic rays and gravitational waves, collectively providing a comprehensive view of the harshest physical environments in the universe.

Among these cosmic messengers, neutrinos occupy a unique position. Owing to their extremely small interaction cross section, neutrinos traverse vast distances through matter almost unperturbed, as they participate predominantly in weak interactions. This characteristic preserves the directional information of neutrinos, allowing their arrival directions to be traced back to their sources and providing an unobscured view of the physical processes that produced them. Their elusive nature, however, necessitates detectors with enormous instrumented volumes. The field of neutrino astronomy was pioneered by the Homestake experiment in 1968 [1], followed by the deployment of the first water-based Cherenkov detector, the Baikal neutrino telescope, in 1993 [2]. Today, neutrinos play a pivotal role in multi-messenger astronomy, complementing electromagnetic, gravitational-wave and cosmic ray observations to provide a more complete understanding of astrophysical sources.

At the forefront of neutrino astronomy is the IceCube Neutrino Observatory [3], the world's largest detector by volume, instrumenting approximately $1\,\mathrm{km}^3$ of glacial ice at the South Pole as its detection medium. Neutrino interactions within this volume produce charged secondary leptons that emit Cherenkov radiation, which is recorded by photomultiplier tubes housed inside Digital Optical Modules (DOMs). The detector consists of 86 strings, each equipped with 60 DOMs, as well as a more densely instrumented sub-array, *DeepCore*, optimized for neutrinos with energies of $\mathcal{O}(10\,\mathrm{GeV})$. Since its completion, IceCube has achieved numerous scientific milestones, including the first evidence for an astrophysical high-energy neutrino flux [4] and the identification of neutrino sources such as the blazar TXS 0506+056 [5], the active galactic nucleus NGC 1068 [6] and diffuse neutrino emission from the Galactic Plane of the Milky Way [7].

Building upon these successes, a recent addition to the detector, IceCube Upgrade [8], augments the existing optical array with five additional strings equipped with advanced optical modules designed to enhance detector calibration and low-energy performance. Concurrently, plans for a substantial expansion to $8\,\mathrm{km}^3$ under IceCube-Gen2 [9, 10] are underway, aiming to extend the detector's sensitivity to astrophysical neutrinos over a broad energy range. IceCube Upgrade also serves as a testbed for the IceCube-Gen2 optical modules that are being evaluated for future deployment.

The work described in this thesis has contributed to the development and characterization of next-generation optical modules and assesses potential sensitivity improvements enabled by IceCube Upgrade and the multi-PMT optical modules. It is structured into three main parts: first, the theoretical foundations of neutrino astronomy, along with the principles of detector technology, are introduced. The second part presents the hardware development, characterization and simulation studies for the IceCube-Gen2 optical modules. The final part develops a point-source analysis of the active galactic nucleus NGC 1068, evaluating the scientific potential of the low-energy detector extension.

Part I

Foundation

2 The High-Energy Universe and Neutrino Astronomy

This chapter establishes the astrophysical foundation for IceCube and the analysis presented in this thesis. Specifically, it traces the physical journey from the acceleration of primary cosmic rays to the production of secondary neutrinos, to their propagation across cosmic distances and eventual detection on Earth. This progression provides the physical context needed to motivate both the search for specific astrophysical sources, such as active galactic nuclei (AGNs), and the need for measurements with ground-based instruments like IceCube.

2.1 High-Energy Cosmic Rays

High-energy cosmic rays (HECRs) are primarily fully ionized atomic nuclei originating from extraterrestrial sources, with energies spanning more than ten orders of magnitude, from a few GeV ($\sim 10^9\,\mathrm{eV}$) to hundreds of EeV ($\sim 10^{20}\,\mathrm{eV}$) [11]. They are the principal messengers of the non-thermal universe[1], providing insights into some of the harshest astrophysical environments in the cosmos, from supernova remnants to the relativistic jets of supermassive black holes. In terms of composition, the cosmic ray flux is dominated by protons and helium nuclei, with heavier nuclei and electromagnetic components contributing only a minor fraction of the total flux [11]. The vast energy range reflects a hierarchy of astrophysical accelerators: While solar and stellar processes produce lower-energy particles, only the most powerful known astrophysical environments are capable of accelerating particles to ultra-high energies. These accelerators are rare, which results in a steeply falling flux of particles reaching the Earth, empirically described by a power law:

$$\frac{dN}{dE} \propto E^{-\gamma}, \tag{2.1}$$

where E is the particle energy and γ is the spectral index. The value of γ is a key diagnostic of the underlying acceleration and propagation physics [11].

The measured cosmic ray spectrum, as shown in Figure 2.1, is continuous over many orders of magnitude, implying a common acceleration mechanism. However, it exhibits distinct deviations from a single power law, most notably the *knee*, the *second knee* and the *ankle*. These features mark transitions between different source populations and elemental compositions. At energies around 3 PeV (the knee), the spectrum steepens ($\gamma \approx 2.7 \rightarrow 3.1$) [13]. This transition is generally interpreted as the point where protons either reach the maximum acceleration limit of Galactic sources (such as supernova remnants) or can no longer be magnetically confined, causing them to escape the Milky Way [11]. Because the maximum energy in both magnetic acceleration and confinement scales with

[1] The "non-thermal universe" refers to astrophysical emission originating from particles accelerated to relativistic energies out of thermal equilibrium, rather than from standard thermal blackbody radiation (e.g., starlight). It is characterized by processes such as synchrotron emission, inverse Compton scattering and hadronic particle cascades.

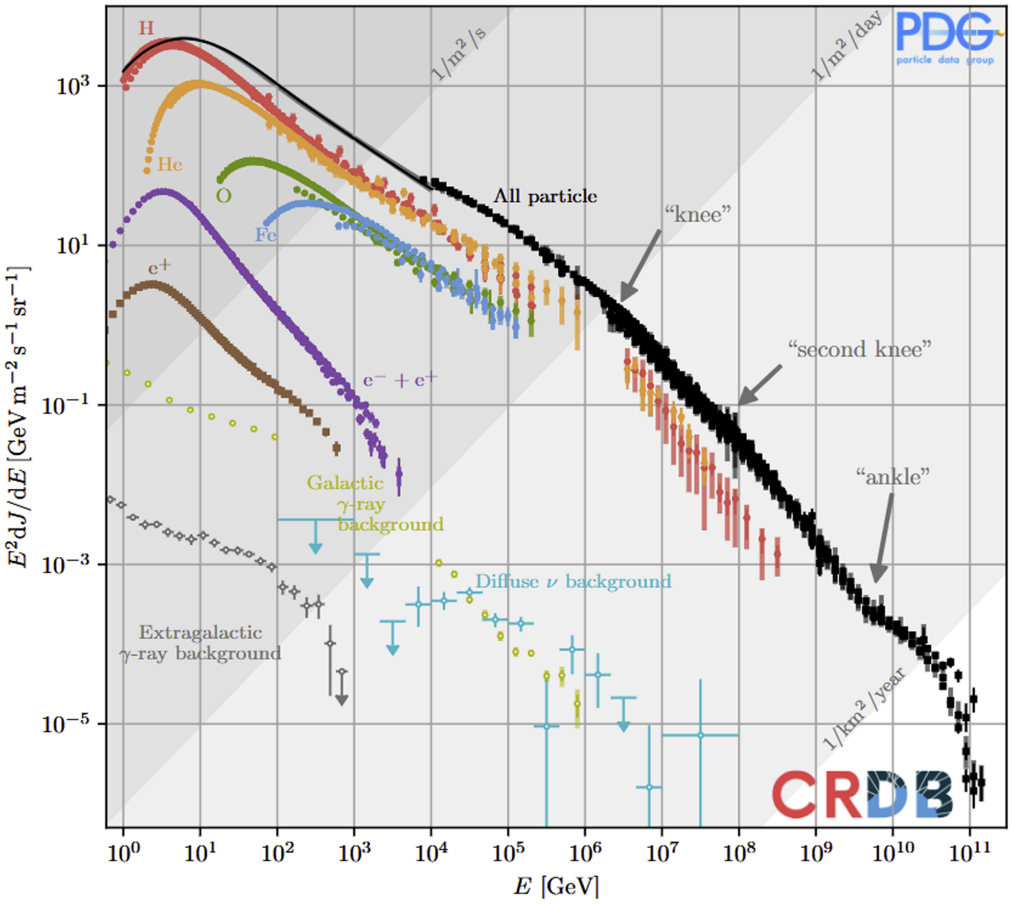


Figure 2.1: The all-particle cosmic ray energy spectrum (black) alongside the contributions of respective components (colored). While the measured spectrum extends down to MeV energies (dominated by solar emission), the high-energy regime relevant to neutrino astronomy begins at energies of several tens of GeV. The main spectral features, the knee (~ 3 PeV) and the ankle (~ 5 EeV), alongside the second knee, are highlighted. The flux is multiplied by E^2 to emphasize changes in the spectral shape. Taken from [12].

the nuclear charge Z [11], heavier Galactic nuclei can reach proportionally higher energies. As a result, beyond the first knee, the cosmic ray composition becomes increasingly heavy. This trend is expected to continue toward the second knee ($\sim 100\,\text{PeV}$ to $1\,\text{EeV}$), which is believed to mark the analogous limit for the heaviest nuclei from Galactic sources, such as iron ($Z = 26$), causing the overall Galactic flux to drop off.

At approximately 5 EeV (the ankle), the spectrum flattens again. This feature marks the energy scale where the hard spectrum of the extragalactic sources (e.g., AGNs) overtakes the fading tail of the heavy Galactic component [13]. Finally, at the end of the spectrum ($E > 5 \times 10^{19}\,\text{eV}$), the observed flux exhibits sharp suppression. Historically, this was entirely attributed to the *Greisen-Zatsepin-Kuzmin* (GZK) cutoff [14, 15], a propagation effect rather than a limit of the accelerators themselves. To a proton traveling at such Lorentz factors ($\Gamma > 10^{10}$), photons of the cosmic microwave background (CMB) appear as high-energy γ-rays. The protons interact with these photons via the Δ-resonance:

$$p + \gamma_{\text{CMB}} \to \Delta^+ \to n + \pi^+ \quad (\text{or } p + \pi^0). \tag{2.2}$$

While this specific interaction drains energy from primary protons, heavier cosmic ray nuclei are similarly attenuated through photo-disintegration upon scattering with the CMB and the cosmic infrared background. Together, these propagation interactions effectively limit the "visible" horizon for all UHECRs to the local universe ($\sim 50\,\mathrm{Mpc}$). However, recent precision measurements by the Pierre Auger Observatory indicate that this flux suppression is actually dominated by **source exhaustion** [16, 17]. In this modern paradigm, extragalactic accelerators merely reach their maximum acceleration capability, leading to an increasingly heavy composition at the highest energies. Crucially, for the fraction of cosmic rays that do undergo GZK-like interactions, where the primary particle is lost (or shifted to lower energies), the energy is redistributed among secondary particles. The decay of the secondary pion ($\pi^+ \to \mu^+ + \nu_\mu$) produces a flux of secondary *cosmogenic* or *GZK neutrinos* at EeV energies [18].

Unlike charged cosmic rays, these neutrinos are unaffected by magnetic fields and experience negligible attenuation from photon backgrounds such as the CMB, theoretically allowing them to carry spectral and spatial information from the distant universe. However, the expected magnitude of this flux is tied to the source exhaustion paradigm: because the extragalactic flux becomes dominated by heavy nuclei at these energies, the resulting photo-disintegration yields fewer high-energy neutrinos than the original pure-proton GZK models [19, 20].

2.1.1 Acceleration Mechanisms and Propagation

The non-thermal power law spectrum of cosmic rays implies that they are not produced by thermal emission (which would yield an exponential Boltzmann distribution) but rather by stochastic acceleration processes. The standard paradigm for this acceleration is the Fermi mechanism, which occurs in astrophysical plasmas carrying magnetic fields.

In his original 1949 proposal [21], Fermi postulated that particles gain energy through repeated scattering encounters with moving magnetic clouds in the interstellar medium. Because the clouds move randomly, head-on collisions, which increase particle energy, occur only slightly more frequently than trailing collisions, which decrease it. The resulting net fractional energy gain per interaction is proportional to the square of the cloud velocity ($\Delta E/E \propto \beta^2$). This quadratic dependence gives the process its name, *second-order Fermi acceleration* and makes it inefficient and too slow to account for the observed high-energy flux.

The currently favored model is *first-order Fermi acceleration*, commonly known as **diffusive shock acceleration** (DSA) [11]. This occurs in environments such as supernova shock fronts or relativistic AGN jets. Here, particles cross a shock boundary repeatedly. Due to the organized bulk motion of the shock, every crossing results in an energy gain proportional to the first power of the velocity ($\Delta E/E \propto \beta$), leading to its "first-order" designation. In its simplest form, this mechanism predicts a power law spectrum with a spectral index of $\gamma \approx 2.0$ at the source. The steepening to the observed $\gamma \approx 2.7$ (see Figure 2.1) is attributed to propagation effects: Higher-energy cosmic rays escape Galactic magnetic fields more readily than lower-energy ones, modifying the spectrum by a diffusion factor $\delta \approx 0.7$ (where the diffusion coefficient is proportional to E^δ) [11].

Regardless of the acceleration efficiency, a charged particle can only be accelerated as long as it remains confined within the source. If the particle's Larmor radius exceeds the physical size of the accelerator R, it escapes and stops gaining further energy. This

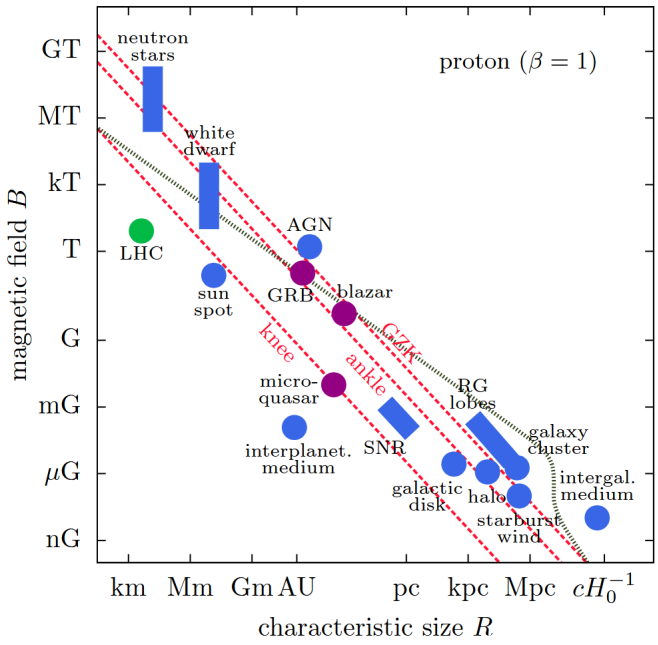


Figure 2.2: The Hillas plot showing the magnetic field strength B against the size R of various astrophysical objects. The diagonal lines represent the minimum requirements to accelerate protons to specific energies; below these lines, the sources cannot confine them. Red lines correspond to the confinement regions for protons at the knee, ankle and GZK cutoff, while the gray line indicates a second-order limit from synchrotron losses. Taken from [23].

geometric constraint, known as the **Hillas criterion** [22], provides an approximate upper limit on the achievable energy as per

$$E_{\max} \approx ZeBR\beta c, \tag{2.3}$$

where Z is the charge number of the nucleus, e is the elementary charge, B is the magnetic field strength and βc is the shock velocity of the accelerator. This relationship is visualized in the *Hillas plot* in Figure 2.2. To accelerate charged particles to ultra-high energies ($> 10^{19}$ eV), a source must possess a sufficient combination of vast dimensions (such as galaxy clusters) and strong magnetic fields (such as neutron stars). AGNs sit in the "high B, large R" center of this phase space, identifying them as prime candidates for cosmic ray acceleration and, by extension, high-energy neutrino production.

Once charged particles are accelerated, they dissipate energy via radiation. A central challenge in high-energy astrophysics is determining the nature of the accelerated charged particles based on this radiation. The high-energy γ-ray emission observed from sources such as AGNs is typically interpreted through two primary models, which may also act in tandem:

1. **Leptonic Models:** The emission is dominated by accelerated electrons. Low-energy photons (such as those from the CMB or accretion disk) are up-scattered to γ-ray energies via inverse Compton (IC) scattering.
2. **Hadronic Models:** The emission originates from accelerated protons or nuclei. These interact with ambient matter (pp) or photon fields ($p\gamma$) to produce neutral pions (π^0) and charged pions ($\pi^\pm$).

Crucially, neutral pions from hadronic interactions decay into γ-rays ($\pi^0 \to \gamma\gamma$), producing a signal that can be indistinguishable from leptonic IC emission. However, charged pions decay into neutrinos ($\pi^\pm \to \mu^\pm \nu_\mu(\bar{\nu}_\mu)$). Therefore, neutrinos are the "smoking gun" for hadronic acceleration. The detection of high-energy neutrinos from a source provides compelling evidence for the presence of accelerated protons, addressing the century-old mystery of cosmic ray origins [24, 25].

2.1.2 Atmospheric Neutrinos and Muons

The direct detection of primary cosmic rays is limited to small-scale satellite experiments such as Fermi-LAT [26] and AMS-02 [27]. For indirect detection, ground-based (or underground) detectors such as HAWC [28] and the Pierre Auger Observatory [29] observe *extensive air showers* produced when a HECR nucleus collides with an atmospheric nucleus, such as nitrogen or oxygen.

This interaction initiates a hadronic cascade, producing a stream of secondary particles dominated by pions and kaons [11]. While neutral pions decay rapidly into γ-rays, the charged mesons decay into muons and neutrinos, with the muons subsequently decaying into electrons and additional neutrinos:

$$\begin{aligned}\pi^\pm \to & \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\ & \hookrightarrow e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu).\end{aligned}$$

Thus, ground-based detectors do not view the cosmos directly, but rather through the screen of the Earth's atmosphere. This atmospheric flux constitutes the primary background for astrophysical neutrino searches and is categorized into two distinct components based on the lifetime of the parent meson.

The **conventional neutrino flux** arises from the decay of light mesons ($\pi^\pm, K^\pm$). At high energies, these particles have relatively long lifetimes ($\tau_\pi \approx 26\,\text{ns}$) and relativistic decay lengths that can exceed the scale height of the atmosphere. Hence, there is competition between decay (producing a neutrino) and further interaction (losing energy in a collision with air). As energy increases, the interaction probability dominates, causing the atmospheric neutrino spectrum to steepen relative to the primary cosmic ray spectrum by one power of energy ($E^{-2.7} \to E^{-3.7}$) [30]. This component exhibits a zenith-angle dependence at sufficiently high energies ($\propto \sec\theta$), as mesons entering near the horizon travel through a thinner atmosphere for longer, favoring decay [18].

The rate of atmospheric muons and neutrinos observed at ground level is not constant but exhibits **seasonal variations** due to fluctuations in the thermodynamic state of the atmosphere. This effect is driven by the ideal gas law: in the summer, the stratospheric temperature rises, causing the atmosphere to expand and the local density to decrease. In this less dense medium, the mean free path for meson interaction increases, giving pions and kaons more time to decay before colliding. Thus, for neutrino detectors, the observed atmospheric background rate is correlated with the stratospheric temperature [31].

While technically present also at lower energies, the **prompt neutrino flux** begins to dominate the atmospheric background at high energies ($\gtrsim 100\,\text{TeV}$). This flux originates from the decay of heavy hadrons containing charm quarks, such as D-mesons and Λ_c baryons. These particles have short lifetimes ($\tau \sim 10^{-12}\,\text{s}$) and decay almost immediately after production, regardless of the local air density, resulting in an approximately

isotropic flux. Since they decay before they can interact, the prompt atmospheric neutrino spectrum follows the primary cosmic ray spectrum ($\sim E^{-2.7}$ [30]), making it *harder* than the conventional flux. While atmospheric neutrinos as a whole constitute the general background for IceCube, this prompt component represents an irreducible high-energy floor for astrophysical searches [18]. Because its harder spectrum mimics the expected astrophysical signal, it cannot be easily filtered out using simple energy thresholds.

Beyond their distinct energy spectra, these atmospheric components also exhibit characteristic flavor compositions that differ from astrophysical expectations. In the standard hadronic cascade, electron neutrinos are produced almost exclusively through the subsequent decay of secondary muons. However, because high-energy muons produced in conventional air showers experience significant relativistic time dilation, they typically reach the ground before this decay can occur. With this electron-neutrino production channel heavily suppressed, the conventional atmospheric neutrino flux is dominated by muon neutrinos ($\nu_e \ll \nu_\mu$) [32]. In contrast, the semileptonic decays driving the prompt flux couple equally to electrons and muons, resulting in a symmetric flavor ratio of approximately $\nu_e : \nu_\mu \approx 1 : 1$ [33]. Crucially, the atmospheric tau neutrino background is negligible ($\nu_\tau \approx 0$). In the conventional flux, the heavy τ lepton cannot be produced by the decay of lighter pions and kaons. In the prompt flux, its production is heavily suppressed because it relies on the relatively infrequent production of D_s mesons in the air shower [33]. Therefore, detecting high-energy tau neutrinos provides a comparatively low-background signature of astrophysical origin [18].

2.2 High-Energy Neutrino Sources

The search for cosmic ray accelerators is complicated by the inherent physical properties of the available astronomical messengers. Because primary cosmic rays are electrically charged, they are deflected by Galactic and intergalactic magnetic fields. As a result, except at ultra-high energies, they lose their directional information and cannot be traced back to their sources. While the highest-energy cosmic rays ($> 10^{19}\,\mathrm{eV}$) experience minimal magnetic deflection and roughly preserve their pointing information, their observable horizon is restricted to the local universe due to the GZK suppression discussed previously. Similarly, while electrically neutral γ-rays travel in straight lines and are excellent messengers at lower energies, the universe becomes increasingly opaque to them at very high energies. Above the TeV scale, extragalactic photons are attenuated via pair production on background photon fields such as the cosmic microwave background (CMB) and extragalactic background light (EBL) [34], reducing their mean free path to the order of 100 Mpc.

Identifying the most energetic accelerators, therefore, requires a messenger that both points back to its source and survives cosmological baselines. As electrically neutral leptons with sub-eV masses ($< 0.45\,\mathrm{eV}$ [35]), neutrinos do not participate in electromagnetic or strong interactions. While they are subject to gravitational lensing like photons, they are unaffected by intergalactic magnetic fields that scramble the trajectories of charged cosmic rays. Furthermore, their exceedingly small weak cross section with matter ($\sigma_\nu \sim 10^{-38}\,\mathrm{cm}^2 \times E_\nu/\mathrm{GeV}$ [36]) allows them to escape radiation-choked astrophysical environments and traverse the observable universe without attenuation. By providing the only unobstructed view of the high-energy universe, this unique capability motivates the construction of large-scale neutrino observatories [37].

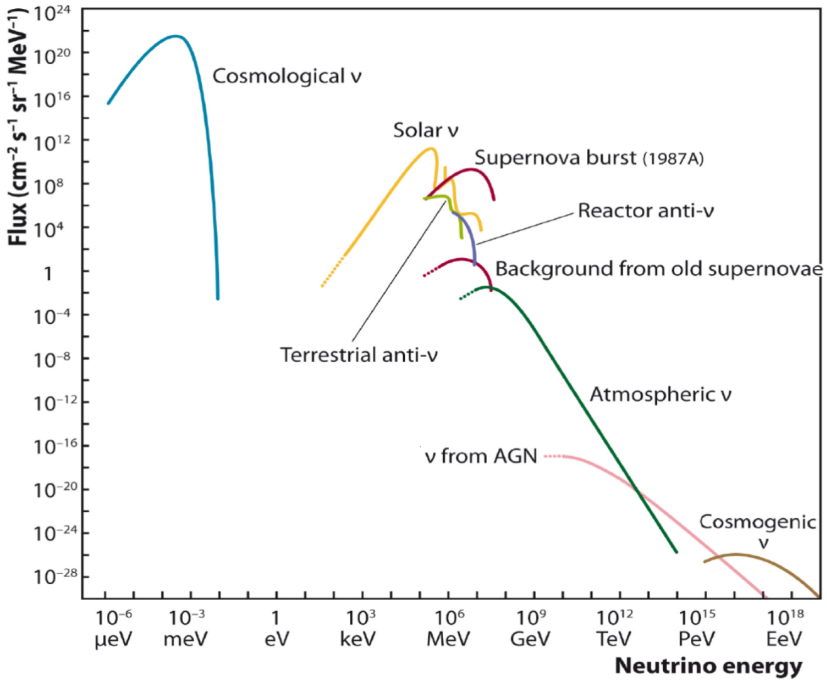


Figure 2.3: The neutrino energy spectrum spanning approximately 24 orders of magnitude. Different colors indicate the dominant neutrino production sites across the energy range. Taken from [36].

Since astrophysical neutrinos are secondary particles from interactions of cosmic rays, their energy spectra are intimately linked. In high-energy hadronic interactions, the primary proton transfers a fraction of its energy to secondary mesons, with the leading pion typically acquiring approximately 20 % of the initial energy. Subsequently, the decay chain $\pi \to \mu \to e$ distributes the energy among four leptons. Therefore, an individual neutrino carries approximately 5 % of the primary proton's energy ($E_\nu \approx E_p/20$) [18]. This kinematic relationship implies that the power law behavior of the primary cosmic rays is imprinted onto the neutrino flux, but shifted to lower energies, as seen in Figure 2.3.

2.2.1 Neutrino Production Mechanisms

Astrophysical neutrinos are produced in the interactions of HECRs with surrounding matter or radiation fields. Unlike the thermal neutrino flux from stellar fusion (which peaks at the MeV scale), the high-energy astrophysical neutrinos targeted by observatories like IceCube (typically spanning the TeV to PeV regime) originate predominantly from hadronic processes. Depending on the environment of the accelerator, these interactions are classified into two primary channels: **hadron-nuclear** (pp) and **photohadronic** ($p\gamma$) interactions.

In environments with low gas density but high photon density (such as the accretion disks and relativistic jets of AGNs), cosmic ray protons interact with ambient photons (UV, X-ray or synchrotron). The dominant channel for this interaction is the formation of the $\Delta^+(1232)$ resonance, identical to the GZK process described in Section 2.1 but occurring within the source itself rather than in the intergalactic medium:

$$p + \gamma \to \Delta^+ \to \begin{cases} p + \pi^0 & \text{(fraction 2/3)} \\ n + \pi^+ & \text{(fraction 1/3)} \end{cases}$$

The branching ratios are determined by isospin conservation, favoring neutral pion production by a factor of two. As mentioned before, the neutral pions decay into γ-rays, while the charged pions initiate the neutrino production chain [11].

In environments with high matter density (such as starburst galaxies), cosmic rays interact inelastically with cold gas protons. Unlike the resonant $p\gamma$ process, pp interactions produce charged and neutral pions with roughly equal probability ($\pi^+ \approx \pi^- \approx \pi^0$) due to isospin symmetry [11]. Thus, pp sources act as reservoirs for both neutrino and γ-ray production, with a slightly softer spectral index compared with $p\gamma$ sources.

Regardless of the production channel (pp or $p\gamma$), the subsequent neutrino emission is governed by the weak decay of the charged pions. As introduced in Section 2.1.2, the complete decay chain proceeds as follows:

$$\begin{aligned} \pi^+ &\to \mu^+ + \nu_\mu \\ &\hookrightarrow e^+ + \nu_e + \bar{\nu}_\mu \end{aligned} \qquad \text{and} \qquad \begin{aligned} \pi^- &\to \mu^- + \bar{\nu}_\mu \\ &\hookrightarrow e^- + \bar{\nu}_e + \nu_\mu. \end{aligned}$$

Counting the neutrino flavors in the full chain yields the canonical ratio of

$$(\nu_e : \nu_\mu : \nu_\tau)_S \approx (1 : 2 : 0) \tag{2.4}$$

at the source. Note that intrinsic ν_τ production within the sources is negligible. Because the large mass of the τ lepton ($m_\tau \approx 1.77\,\text{GeV}$) makes it kinematically inaccessible for pion decay ($m_\pi \approx 0.14\,\text{GeV}$), ν_τ production would require the decay of much heavier, rarely produced charmed mesons (such as D_s). Consequently, the initial source flavor ratio is effectively devoid of tau neutrinos and any significant ν_τ flux observed at Earth is predominantly the result of neutrino oscillations during cosmic propagation (see Section 2.3).

The standard (1:2:0) ratio assumes that the secondary muon decays freely. However, in sources with strong magnetic fields or high matter density (such as the cores of AGNs), the muon may lose energy via synchrotron radiation or collision before it can decay. If the synchrotron cooling time is shorter than the muon lifetime, the muon is "damped" [38]. It eventually decays at low energies, contributing only to the thermal background rather than the high-energy flux. In this scenario, only the initial neutrino from the pion decay survives at high energies: $\pi^+ \to \mu^+(\text{damped}) + \nu_\mu$. This creates a **muon-damped** source with a flavor ratio of purely muon neutrinos:

$$(\nu_e : \nu_\mu : \nu_\tau)_S \approx (0 : 1 : 0).$$

Identifying such a deviation is a key science goal for neutrino detectors, as it would directly probe the magnetic field strength of the accelerator.

2.2.2 Active Galactic Nuclei

Among the most powerful particle accelerators in the universe are active galactic nuclei (AGNs). These are the compact central regions of galaxies, powered by the accretion of matter onto a supermassive black hole (SMBH). While the host galaxies can be massive (up to $10^{12}\,M_\odot$), the central SMBHs themselves typically range from 10^6 to $10^{10}\,M_\odot$ [39].

The energy source of an AGN is gravitational potential energy. As gas and dust fall into the gravitational well of the SMBH, they possess angular momentum, preventing direct radial infall. Instead, the matter forms a flattened *accretion disk* orbiting the black hole. Within this disk, differential rotation creates friction, which heats the plasma to millions of degrees, releasing intense thermal radiation in the UV and X-ray bands. Simultaneously, the ionized plasma drags magnetic field lines along with it. The rotation of the disk winds these field lines into a tight helix. This twisted magnetic field collimates a fraction of the infalling matter and launches it outward along the rotation axis at relativistic speeds [40]. These **relativistic jets** can extend for kiloparsecs, carrying charged particles that are accelerated via shock waves to produce HECRs. The extreme environments in the vicinity of SMBHs are capable of accelerating protons well beyond PeV energies, making AGNs prime candidates for the universe's *PeVatrons* [18, 41].

While AGNs exhibit a diverse range of observational properties (classifying them as blazars, quasars, radio galaxies or Seyfert galaxies), the **unified model** suggests that these classes represent the same physical object viewed from different angles [42]. As illustrated in Figure 2.4, the central engine is surrounded by a dusty toroidal structure. When viewed edge-on, the line of sight to the central engine is obscured by the dusty torus which is classified as **type II Seyfert galaxies**. NGC 1068, the focus of the analysis in the third part of this thesis, is a prototypical type II Seyfert galaxy. From Earth, it is viewed through the Compton-thick environment of the torus, which blocks optical and soft X-ray emission from the core.

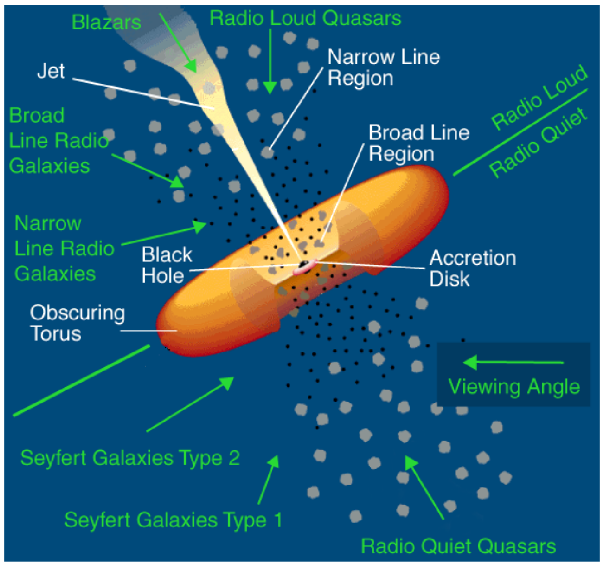


Figure 2.4: Schematic illustration of the AGN unified model. The central supermassive black hole and its accretion disk are surrounded by a dusty torus. The classification of the AGN depends on the viewing angle. Taken from [43].

Historically, neutrino searches focused on blazars, assuming neutrinos would be produced in the jets alongside the observed γ-rays. However, the recent association of high-energy neutrinos with NGC 1068 [6, 44] presents observational tension: the source is "neutrino bright but gamma-ray dim" in the GeV–TeV range probed by Fermi-LAT [26] and IceCube [3]. Standard jet models fail to explain this, as the expected γ-ray flux accompanying the neutrinos would violate observational upper limits. A compelling solution is provided by the **Corona Model** (or Hidden Core Model) [45]. In this scenario, cosmic ray acceleration does not occur in a remote jet, but in the hot, turbulent *corona* situated immediately above the accretion disk:

1. **Acceleration:** Magnetic turbulence in the corona accelerates protons to energies up to $\mathcal{O}$(PeV).
2. **Production:** These protons interact with the intense X-ray field radiated by the accretion disk via the $p\gamma$ mechanism (Section 2.2.1), producing pions.
3. **Opacity:** The target photon density is so high (often reaching the Eddington luminosity limit[2]) that the environment becomes opaque to γ-rays. The produced neutral pions decay into very-high-energy (VHE) γ-rays, which rapidly interact with the ambient X-ray background via pair production:

$$\gamma_{\text{VHE}} + \gamma_{\text{X-ray}} \rightarrow e^+ + e^-.$$

   The resulting electromagnetic cascade thermalizes the energy down to MeV scales, extinguishing the high-energy γ-ray signal.

Crucially, the neutrinos produced by the charged pions escape this dense, obscured environment unperturbed. This creates a distinctive signature: a source that is invisible in TeV γ-rays but luminous in neutrinos. Hence, for "hidden cores" like NGC 1068, neutrinos are not merely complementary messengers; they are the only probe capable of penetrating the torus and the corona to reveal the physics of the central engine. Regardless of whether the emission originates from the compact corona (a region only a few Schwarzschild radii across) or an extended kiloparsec-scale jet, the cosmological distance to the host galaxy causes the entire AGN to appear as an unresolved neutrino point source on the sky when observed from Earth. Detailed observation of these faint sources requires precise pointing to distinguish their localized high-energy excess from the terrestrial atmospheric background.

2.2.3 Other Source Candidates

In addition to type II Seyfert galaxies such as NGC 1068, several source classes have been proposed as potential high-energy cosmic neutrino production sites because they satisfy the Hillas criterion and possess the necessary baryonic loading to generate detectable neutrino fluxes.

[2]The Eddington limit is the maximum luminosity an accreting object can achieve while remaining in hydrostatic equilibrium. At this limit, the outward radiation pressure balances the inward gravitational pull on the surrounding plasma. If an AGN were to significantly exceed this limit, the radiation pressure would expel the accreting matter, effectively depleting the central engine of accretion material [46].

Blazars

Blazars are a subclass of AGNs where the relativistic jet is aligned within a few degrees of the observer's line of sight. Unlike in the hidden core scenario, the emission is dominated by the jet itself. In 2017, IceCube detected a high-energy neutrino coincident with a flaring state of the blazar TXS 0506+056 [5]. Follow-up observations by other detectors such as Fermi-LAT and MAGIC [47] confirmed enhanced γ-ray emission, providing the first multi-messenger evidence for a specific neutrino source. However, analyses of the diffuse neutrino flux [48] suggest that blazars alone cannot account for the entirety of the observed cosmic signal, implying a significant contribution from other populations.

Starburst Galaxies

While AGNs are characterized by their central engines, these phenomena are not mutually exclusive with other astrophysical environments. Many (composite) galaxies host both an active nucleus and regions of high star formation. A prominent example is NGC 1068 itself, which features a heavily obscured Seyfert core surrounded by a circumnuclear starburst ring [49]. In such hybrid objects, both environments can theoretically accelerate cosmic rays and produce high-energy neutrinos.

However, the underlying physics of the starburst component itself differs. Unlike the accretion-driven central engine of an AGN, this emission is powered by a high rate of star formation. The frequent supernovae in these regions accelerate cosmic rays, which then interact with interstellar gas via *pp* interactions. These galaxies act as reservoirs, accelerating cosmic rays and converting a considerable fraction of their energy into neutrinos and γ-rays before they can escape [50]. Because the star formation rate in these distant sources is higher than in our own galaxy, integrating their aggregate neutrino emission over cosmological distances is expected to contribute significantly to the observed isotropic diffuse neutrino background.

Diffuse Galactic Emission

To describe the distribution of localized cosmic sources as opposed to extended structures, two primary coordinate systems are employed [39]. *Equatorial coordinates* represent the standard projection of Earth's latitude and longitude onto the celestial sphere (declination δ, right ascension α). Conversely, *Galactic coordinates* are a system centered on the Sun, where the fundamental plane ($b = 0°$) is aligned with the disk of the Milky Way. The projection of this dense disk onto the sky is known as the **Galactic Plane** (typically $|b| < 15°$). In equatorial coordinates, this band appears as a characteristic "U" or "S" shape arching across the sky.
In 2023, IceCube reported the first observation of high-energy neutrinos from the Galactic Plane with a significance of 4.5σ [7], as shown in Figure 2.5. Unlike point sources, this emission is diffuse, spanning the entire Galactic disk. It arises from the interaction of cosmic rays trapped by the Galactic magnetic field with interstellar gas and dust (*pp* interactions). This observation supports the reservoir mechanism within our own galaxy, demonstrating that the Milky Way acts as a vast cosmic ray beam dump. It establishes the Galactic Plane as the most prominent extended source of high-energy neutrinos in our immediate cosmic neighborhood.

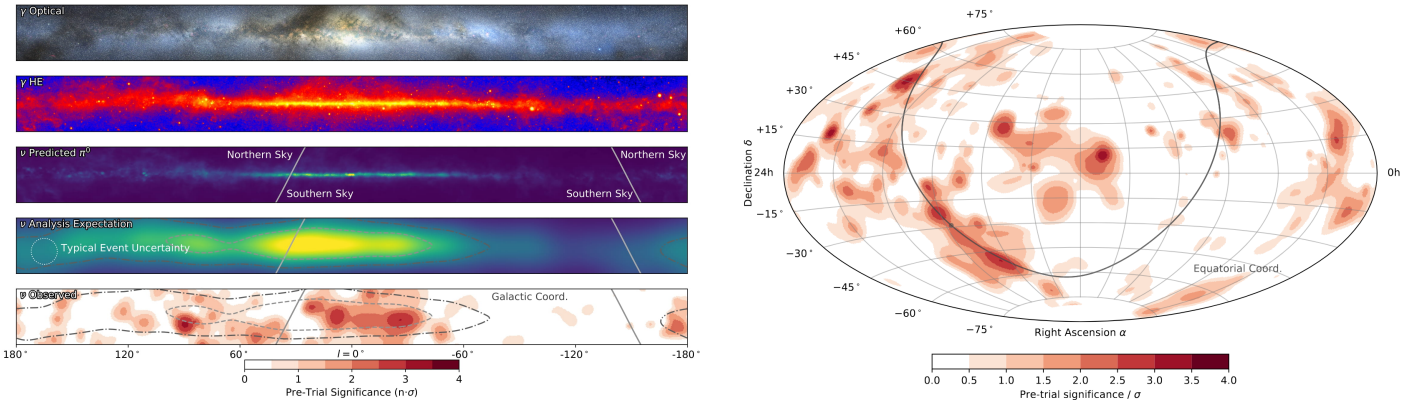


Figure 2.5: (left) Multi-messenger view of the Galactic Plane in Galactic coordinates. From top to bottom: 1) Optical view (obscured by interstellar dust); 2) Gamma-ray flux (Fermi-LAT), tracing π^0 decay; 3) Predicted neutrino emission based on hadronic models; 4) IceCube model expectations; 5) Inset from the right panel of measured significance. **(right)** All-sky neutrino significance map in equatorial coordinates. The Galactic Plane is drawn as a gray curve with the Galactic Center marked by the dot at $\delta = -29°$. Taken from [7].

Transient Sources

While AGNs and starburst galaxies are generally considered steady emitters, the high-energy universe is also populated by *transients*. These sources release energy in short, violent bursts, for which targeted analyses benefit from a reduced background by restricting the search window to the specific time of the event. Among many types, transients include:

- **Gamma-Ray Bursts (GRBs):** Energetic electromagnetic explosions occurring when massive stars collapse or neutron stars merge, releasing energy as focused jets of gamma radiation [51, 52]. The *fireball model* [53] predicts that protons accelerated in relativistic shock waves should produce a transient flux of neutrinos. Despite theoretical promise, IceCube searches constrain the contribution of GRBs to the overall observed diffuse neutrino flux [54].
- **Tidal Disruption Events (TDEs):** These occur when a star passes sufficiently close to a SMBH and is torn apart by tidal forces. The resulting accretion flare can power a transient jet capable of high-energy neutrino production [55].
- **Supernovae:** The collapse of massive stars. Specifically, *choked jet* supernovae are hypothesized to produce neutrinos without a visible γ-ray counterpart, as the jet fails to break out of the stellar envelope [56].

2.3 Propagation of Astrophysical Neutrinos to Earth

Produced in astrophysical environments, neutrinos travel distances ranging from Mega- to Gigaparsecs before reaching Earth. Unlike charged cosmic rays, which diffuse through magnetic fields, or γ-rays, which cascade on background radiation or interstellar matter, neutrinos interact only via gravity and the weak force. As a consequence of this minimal interaction, they preserve the original directional and spectral information of their source. However, their flavor composition is not constant; it evolves during propagation due to the quantum mechanical phenomenon of neutrino oscillation.

2.3.1 Neutrino Oscillations

Neutrinos are produced and detected in weak-interaction flavor eigenstates associated with the charged leptons: electron (ν_e), muon (ν_μ) and tau (ν_τ). These are not eigenstates of the free Hamiltonian. Instead, a neutrino of flavor α is created as a coherent superposition of three mass eigenstates ν_i ($i = 1, 2, 3$) with definite masses m_i:

$$|\nu_\alpha\rangle = \sum_{k=1}^{3} U^*_{\alpha k} |\nu_k\rangle.$$

The mixing between these bases is described by the **Pontecorvo-Maki-Nakagawa-Sakata (PMNS)** matrix, U_{PMNS}, as conceptualized in [57, 58, 59]. It is generally parameterized as the product of three rotation matrices, each characterized by a mixing angle θ_{ij} (where $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$) and a complex CP-violating phase δ_{CP}:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{\text{PMNS}} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \qquad \text{with} \tag{2.5}$$

$$U_{\text{PMNS}} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix}}_{\text{Atmospheric}} \underbrace{\begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{bmatrix}}_{\text{Reactor}} \underbrace{\begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\text{Solar}}. \tag{2.6}$$

For Majorana neutrinos, the full mixing matrix contains an additional diagonal phase matrix, $\text{diag}(1, e^{i\alpha_{21}/2}, e^{i\alpha_{31}/2})$ [60]. However, because these phases cancel out in the derivation of the oscillation probability, they have no observable effect on standard neutrino oscillations and are conventionally omitted here. Each angle in the standard parameterization governs the mixing in a specific sector of neutrino physics. Current global fits from NuFIT 6.0 [61] provide the following best-fit values:

- $\theta_{12} = 33.68^{+0.73}_{-0.70}\,^\circ$ (Solar Sector): Dictates the flavor transitions of electron neutrinos produced by nuclear fusion in the Sun. The matter effect within the dense solar interior resonantly enhances this mixing, ultimately leading to the observed deficit of solar electron neutrinos at Earth.
- $\theta_{23} = 48.6^{+0.7}_{-0.9}\,^\circ$ (Atmospheric Sector): Drives the disappearance of muon neutrinos produced by cosmic ray interactions in the Earth's atmosphere. Because current fits still suffer from an octant ambiguity, a major goal of current experiments is to determine the true octant of θ_{23} and whether the mixing is maximal (45°).
- $\theta_{13} = (8.58 \pm 0.11)\,^\circ$ (Reactor Sector): Represents the cross-mixing between the solar and atmospheric sectors. The non-zero value measured by reactor experiments acts as the necessary gateway that makes the observation of leptonic CP violation physically possible.
- $\delta_{\text{CP}} = 285^{+25}_{-28}\,^\circ$ (CP-violating phase): A value deviating from 0° or 180° introduces an asymmetry between neutrino and antineutrino oscillation probabilities, potentially offering insight into the universe's matter-antimatter imbalance. Current global fits mildly prefer values near 270°, although CP-conserving values remain allowed and leptonic CP violation has not yet been established experimentally.

The probability of a neutrino (or antineutrino) changing from flavor α to β in vacuum is given by:

$$P_{\substack{\nu_\alpha \to \nu_\beta \\ \bar{\nu}_\alpha \to \bar{\nu}_\beta}}(L, E) = \delta_{\alpha\beta} - 4 \sum_{k>j} \mathrm{Re}(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin^2\left(1.27 \frac{\Delta m^2_{kj} L}{E}\right) \pm 2 \sum_{k>j} \mathrm{Im}(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin\left(2.54 \frac{\Delta m^2_{kj} L}{E}\right), \tag{2.7}$$

where the positive sign of the imaginary term applies to neutrinos and the negative sign to antineutrinos. In this parameterization, the baseline L is expressed in km, the neutrino energy E in GeV and the mass-squared splittings Δm^2_{kj} in eV^2. Oscillations are driven by the interference between mass eigenstates, which accumulate phase differences relative to one another as they propagate. The frequency of this oscillation depends on the difference in the squares of the masses $\Delta m^2_{ij} \equiv m^2_i - m^2_j$. Experiments have identified two distinct mass scales: *solar splitting* with $\Delta m^2_{21} = (7.49 \pm 0.19) \times 10^{-5}\,\mathrm{eV}^2$ and *atmospheric splitting* with $\Delta m^2_{31} = (2.513^{+0.021}_{-0.019}) \times 10^{-3}\,\mathrm{eV}^2$ [61, 62]. The sign of the solar splitting is measured to be positive, whereas the sign of the atmospheric splitting is currently undetermined, though global fits favor a positive sign (as quoted above). This ongoing ambiguity results in two viable scenarios for the **neutrino mass ordering (NMO)**: *normal ordering* ($\Delta m^2_{31} > 0$) with $m_1 < m_2 < m_3$ and *inverted ordering* ($\Delta m^2_{31} < 0$) with $m_3 < m_1 < m_2$, as illustrated in Figure 2.6. Determining the true mass ordering is a fundamental goal in its own right to complete our understanding of the cosmic neutrino sector, whose relic background constitutes the second most abundant known particle population in the universe [63]. Furthermore, it dictates the theoretical expectations for other major frontiers: assuming neutrinos are Majorana particles, an inverted ordering generally implies a non-zero lower limit on the effective Majorana mass within the standard light-neutrino exchange mechanism, which is within reach of next-generation neutrinoless double beta decay experiments [64]. Simultaneously, establishing the mass ordering is crucial for disentangling matter effects from true leptonic CP violation, thereby alleviating parameter degeneracies in long-baseline accelerator experiments [65].

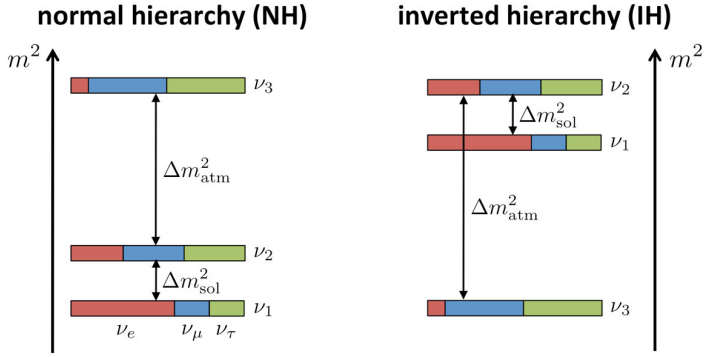


Figure 2.6: The two possible neutrino mass orderings. In normal ordering **(left)**, ν_1 is the lightest mass eigenstate, while in the inverted ordering **(right)**, ν_3 is the lightest. The color composition of each bar indicates the flavor content of the corresponding mass eigenstate ($|U_{\alpha i}|^2$). Courtesy of the JUNO Collaboration.

For terrestrial baselines, the $\sin^2$ terms in Equation (2.7) drive the observable oscillation pattern. However, for astrophysical sources, the distance L is astronomical (Mpc to Gpc). Over these immense scales, two distinct physical effects act to suppress the observable interference pattern: phase **averaging** and wave packet **decoherence**. First, due to the finite energy resolution of the detector, the energy spread of the emitted neutrinos and the vast physical size of the emission region, the rapid oscillation phases cannot be experimentally resolved. Integrating over these macroscopic uncertainties causes the oscillatory terms to smear out. Second, on a quantum mechanical level, neutrinos are described by a wave packet with an intrinsic spatial width [66]. Because the distinct mass eigenstates travel at slightly different group velocities, their respective wave packets gradually drift apart. For many astrophysical production scenarios, the propagation distance exceeds the *coherence length*, which depends on the neutrino energy, the mass splitting and the initial spatial width of the wave packet, causing the wave packets to separate. This permanent loss of wave packet overlap renders the interference terms observationally inaccessible, a regime known as wave packet decoherence in astrophysical neutrino propagation [11].

Whether through phase averaging or wave packet decoherence, the interference terms become unobservable over astrophysical baselines, reducing the oscillation probability to an incoherent sum of mass-eigenstate contributions:

$$\langle P_{\alpha\to\beta}\rangle = \sum_{k=1}^{3} |U_{\alpha k}|^2 |U_{\beta k}|^2. \tag{2.8}$$

This averaged probability governs the flavor composition expected at Earth, independent of distance and neutrino energy for vacuum oscillations (apart from energy-dependent mixing effects) [67].

2.3.2 Flavor Composition at Earth

The final observable in neutrino astronomy is the flavor ratio measured at the detector $(\nu_e : \nu_\mu : \nu_\tau)_\oplus$, which is commonly visualized using the **flavor triangle** (or ternary plot). As discussed in Section 2.2.1, the environment of the accelerator determines the initial flavor ratio $(\nu_e : \nu_\mu : \nu_\tau)_S$ at the source. Commonly discussed flavor ratio scenarios are:

1. **Pion Decay (1:2:0):** The standard scenario for AGNs and starbursts, where neutrinos originate from the full decay chain $\pi^\pm \to \mu^\pm \to e^\pm$ as detailed in Section 2.2.1.
2. **Muon-damped (0:1:0):** Expected in dense or highly magnetized environments, where the muon loses energy before decaying. These conditions may occur in AGN cores such as NGC 1068. The flux is initially composed only of muon-flavor neutrinos (see Section 2.2.1).
3. **Neutron Decay (1:0:0):** In sources where high-energy neutrons escape the production region before further interactions and the conventional pion-decay channel is suppressed (for example by synchrotron cooling in strongly magnetized environments), the neutrino flux can be dominated by neutron beta decay, $n \to p + e^- + \bar{\nu}_e$, yielding an approximately pure electron-antineutrino flavor composition at the source [70].

While treated as distinct scenarios here, sources may transition between these regimes (e.g., from pion-decay to muon-damped) as the energy of the parent particles increases.

Using the time-independent transition probability derived above (Equation (2.8)), the arriving flavor flux at Earth is calculated by projecting the initial source composition through the squared PMNS matrix elements. The large values of the mixing angles (in particular, the atmospheric mixing angle $\theta_{23} \approx 48.6°$ and the solar angle $\theta_{12} \approx 33.7°$) redistribute the flavor composition. Rather than merging into a single indistinguishable ratio, this deterministic mapping translates distinct source models into separate, well-defined points clustered within the central region of the flavor triangle. For the standard pion decay scenario, this redistribution results in approximate equipartition. The initial ratio $(1:2:0)_S$ evolves into approximately equal parts of all flavors at Earth:

$$(1:2:0)_S \xrightarrow{\text{Decoherence}} (\approx 1:1:1)_\oplus. \tag{2.9}$$

This result has direct implications for observational neutrino astronomy: for all standard astrophysical production scenarios, propagation over cosmic baselines predicts a flux of all three flavors at Earth. This expectation necessitates a detector architecture capable

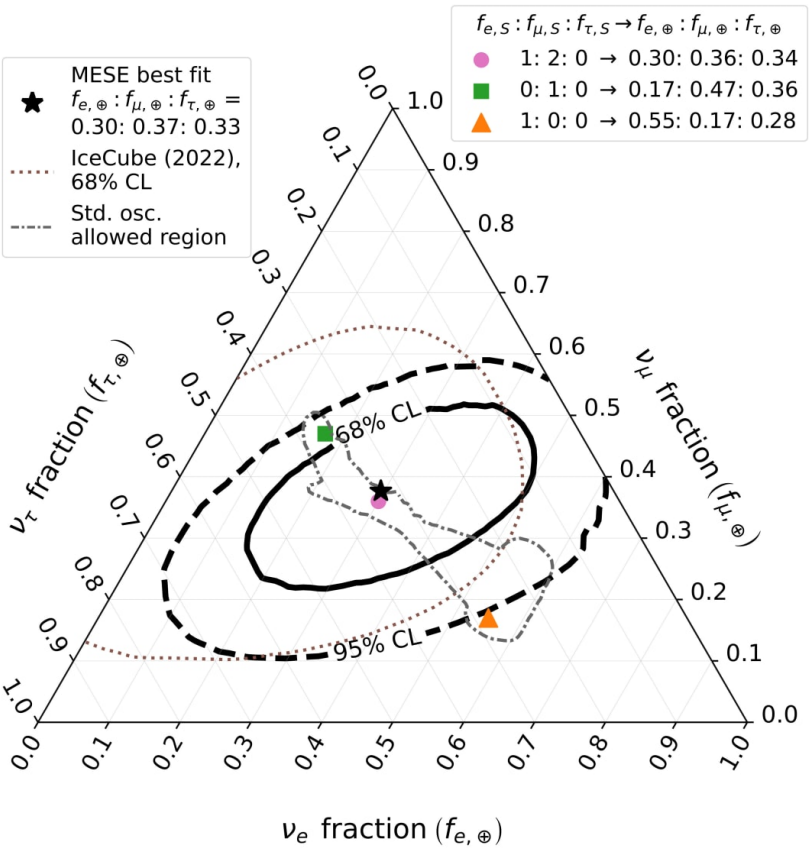


Figure 2.7: Flavor composition of high-energy cosmic neutrinos arriving at Earth. **Experimental fits:** The solid black and dotted brown contours show the 68 % and 95 % confidence level (CL) allowed regions from a recent analysis [68] and a previous IceCube measurement [69], respectively. **Theoretical models:** The specific markers represent the theoretically expected flavor ratios at Earth assuming idealized source production mechanisms: standard pion decay (pink circle), muon-damped pion decay (green square) and neutron decay (orange triangle). The dash-dotted gray region indicates the total kinematically allowed parameter space for standard astrophysical production scenarios after mixing (including uncertainties on oscillation parameters). Taken from [68].

of identifying and distinguishing all three neutrino flavors.

While the redistribution alters the composition, the theoretical arrival ratios of different production models land at fundamentally distinct, albeit close, locations on the flavor triangle, as shown in Figure 2.7. The flavor triangle represents a normalized three-component parameter space in which the sum of the flavor fractions is unity [71]. To aid interpretation, the corners represent a pure state (e.g., the top corner is $100\,\%\,\nu_\mu$), the edges represent a mix of two flavors with zero contribution from the third and the center represents an equipartition of all three flavors $(1:1:1)$. For instance, a muon-damped source $(0:1:0)_S$ evolves to a ratio of approximately $(0.17:0.47:0.36)_\oplus$ at Earth, shifting away from the center toward the ν_μ–ν_τ edge.

Testing these deviations offers a unique way to probe the internal physics of the source and the fundamental properties of neutrinos. For instance, confirming a muon-damped signature would provide evidence of strong magnetic fields, whereas finding a ratio far outside the standard region (e.g., near the corners of the triangle) would point to exotic physics such as neutrino decay or Lorentz invariance violation. However, while these signatures are theoretically well-separated, identifying the underlying source mechanism experimentally remains challenging. This difficulty stems from detector limitations, specifically finite energy and angular resolutions, low event statistics and the topological degeneracy between electron and low-energy tau neutrino cascades (see Section 3.5). These limitations currently result in large observational uncertainties that span multiple theoretical scenarios, as seen in the broad contours of Figure 2.7. Despite these large uncertainties, the flavor composition of high-energy cosmic neutrinos is consistent with standard particle physics (standard oscillations) and standard astrophysics (sources powered by pion decay). Furthermore, the measurement is now precise enough to statistically disfavor neutron decay as the dominant mechanism for producing cosmic neutrinos, suggesting that sources are likely dominated by pp or $p\gamma$ interactions where charged pions and muons decay efficiently [68].

IceCube Upgrade (see Section 3.6), with its enhanced particle identification and angular resolution, combined with IceCube-Gen2 (see Section 3.7), which will provide higher statistics and sufficient instrumented volume to constrain large event signatures, aims to shrink these contour regions and potentially resolve the specific production mechanisms of cosmic accelerators.

3 Neutrino Detection with the IceCube Neutrino Observatory

The property of neutrinos to travel unperturbed through space is a double-edged sword: while it preserves directional and energy information from the astrophysical production site, it necessitates enormous detection volumes to capture these neutrinos, especially because their flux decreases rapidly toward the highest astrophysical energies (see Figure 2.3). A comprehensive understanding of high-energy neutrino detection requires not only knowledge of their interaction mechanisms within the detector medium, but also of their propagation through the Earth. Governed by the neutrino interaction cross section, the probability of a neutrino reaching the detector depends on both its energy and the distance it travels through the Earth. This chapter first outlines the fundamental interaction physics (Section 3.1), then applies these principles to neutrino propagation through the Earth (Section 3.2), discusses the subsequent secondary detection principle (Section 3.3) and finally describes the IceCube detector in the following sections.

3.1 Neutrino Interactions and Cross Sections

At the energy scales relevant for the IceCube Neutrino Observatory, which primarily targets neutrino energies from the TeV to PeV range (with lower-energy sensitivity provided by DeepCore and IceCube Upgrade), the interaction landscape differs from low-energy solar or reactor neutrino physics. While quasi-elastic scattering and resonance production occur at the lowest end of this spectrum, **deep inelastic scattering** (DIS) off nucleons dominates the interaction cross section for high-energy atmospheric and astrophysical neutrinos [72]. In the DIS regime, the incoming neutrino scatters off individual quarks that constitute the nucleons of the detector medium.

3.1.1 Interaction Channels

Neutrino-nucleon DIS interactions occur via the exchange of weak vector bosons, proceeding through two primary channels that govern the resulting event topology in the detector, as illustrated by the Feynman diagrams in Figure 3.1.

Charged current (CC): Mediated by the exchange of a charged $W^{\pm}$ boson, the neutrino ν_ℓ (or antineutrino $\bar{\nu}_\ell$) of flavor $\ell \in \{e, \mu, \tau\}$ converts into its corresponding charged lepton counterpart $\ell^{\pm}$, while the target nucleon N fragments into a hadronic system X:

$$\nu_\ell + N \xrightarrow{W^{\pm}} \ell^- + X \quad \text{and} \quad \bar{\nu}_\ell + N \xrightarrow{W^{\pm}} \ell^+ + X. \tag{3.1}$$

The outgoing charged lepton preserves the flavor identity of the incident neutrino, which provides the basis for flavor-dependent event signatures such as track-like or cascade-like morphologies (see Section 3.5).

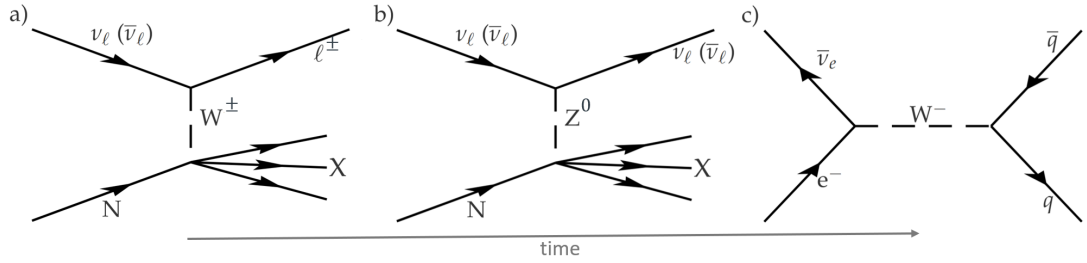


Figure 3.1: Leading-order Feynman diagrams for (a) CC and (b) NC neutrino-nucleon deep inelastic scattering, alongside (c) the Glashow resonance. Modified from [73].

Neutral current (NC): Mediated by the exchange of a neutral Z^0 boson, the neutrino transfers a portion of its four-momentum to the target nucleon and exits the interaction vertex as a neutrino of the same flavor, leaving behind the disrupted hadronic system:

$$\nu_\ell + N \xrightarrow{Z^0} \nu_\ell + X \quad \text{and} \quad \bar{\nu}_\ell + N \xrightarrow{Z^0} \bar{\nu}_\ell + X. \tag{3.2}$$

In both interaction channels, the kinematics are characterized by the inelasticity parameter (Bjorken-y), defined as the fraction of the incident neutrino's energy transferred to the hadronic system in the rest frame of the target nucleon: $y = E_{\text{had}}/E_\nu$, where E_ν is the energy of the incident neutrino and E_{had} is the total energy of the hadronic final state [74]. The average inelasticity $\langle y \rangle$ depends on the weak interaction channel and the parton structure of the nucleon. At energies below approximately 1 PeV, valence quarks contribute significantly to the scattering process [75]. In this regime, the helicity structure of the weak interaction leads to different inelasticity distributions for neutrinos and antineutrinos, with neutrinos transferring about half of their energy to the hadronic system on average. As the neutrino energy increases, interactions probe progressively smaller parton momentum fractions (Bjorken-x), where sea quark and antiquark contributions become increasingly important in the nucleon parton distributions. Consequently, the neutrino and antineutrino inelasticity distributions converge, while the average inelasticity gradually decreases [75].

3.1.2 Cross Section and Energy Dependence

The probability of a neutrino interacting in the detector medium is governed by the total neutrino-nucleon cross section, which exhibits a strong energy dependence. Deep inelastic scattering remains the dominant interaction mechanism throughout the high-energy regime extending toward the EeV scale, however, the energy dependence of its cross section changes across this range. From tens of GeV up to roughly 10 TeV, the DIS cross section grows linearly with the incident neutrino energy ($\sigma \propto E_\nu$). As energies reach the high-TeV and PeV scales, the momentum transfer Q^2 of the interaction becomes comparable to the squared mass of the exchanged weak bosons (M_W^2 or M_Z^2). The presence of the squared boson propagator term, proportional to $1/(Q^2 + M_{W,Z}^2)^2$, dampens the interaction probability [76]. As a result, the cross section growth softens, transitioning from a linear dependence to a weaker power law dependence of approximately $\sigma \propto E_\nu^{0.36}$ [77], as illustrated in Figure 3.2.

A notable exception to this scaling occurs for electron antineutrinos at $E_\nu \approx 6.3\,\text{PeV}$. At this specific energy, an incident $\bar{\nu}_e$ can interact with an atomic electron at rest to

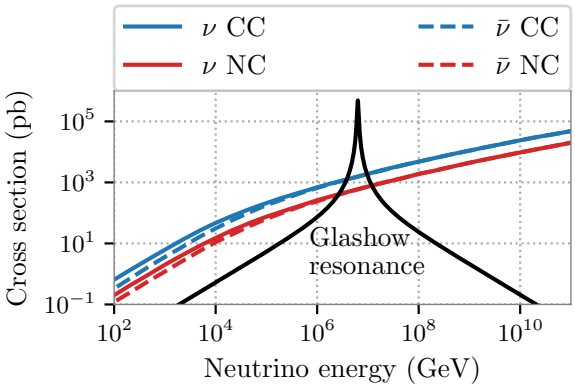


Figure 3.2: Total neutrino and antineutrino cross sections as a function of neutrino energy. CC and NC nucleon DIS data are taken from [77]. The Glashow resonance cross section for $\bar{\nu}_e e^- \to W^-$ is calculated following [78] (accounting for atomic Doppler broadening produces no visible effect on this logarithmic scale).

resonantly produce an on-shell W^- boson ($\bar{\nu}_e + e^- \to W^-$). Known as the Glashow resonance, this process causes the cross section to spike by orders of magnitude over the standard DIS baseline, providing a unique Standard Model signature accessible to IceCube observations [79].

3.2 Neutrino Propagation through the Earth

In the context of IceCube, neutrinos arriving from the Northern sky must traverse a substantial portion of the Earth before reaching the detector. Because the Earth effectively acts as a dense, energy-dependent filter, understanding neutrino propagation through the Earth is crucial for accurate simulation, modeling and analysis. In this regard, two dominant physical effects become relevant depending on the neutrino energy: flavor oscillations governed by matter effects at lower energies and attenuation due to neutrino interactions at higher energies. To accurately model this propagation, analyses utilize the preliminary reference Earth model (PREM) [80], which maps the Earth's internal density profile, including the major density discontinuities between the crust, mantle and core.

3.2.1 The GeV Regime: Oscillations in Matter

For neutrinos with energies from a few GeV to approximately 100 GeV, the interaction cross section is sufficiently small that physical absorption in the Earth is negligible [81]. However, the dense matter alters the neutrino flavor evolution. While propagating through the Earth, electron neutrinos undergo coherent forward scattering via CC weak interactions with the abundant atomic electrons, introducing an additional effective matter potential known as the Wolfenstein potential $V_{CC} = \sqrt{2} G_F N_e$, where G_F is the Fermi coupling constant and N_e is the local electron number density [82]. This matter effect induces the Mikheyev-Smirnov-Wolfenstein (MSW) resonance [83]. At specific energies and mantle densities, the effective mixing angle is maximized, leading to enhancements in flavor transitions compared with vacuum oscillations. Furthermore, neutrinos passing

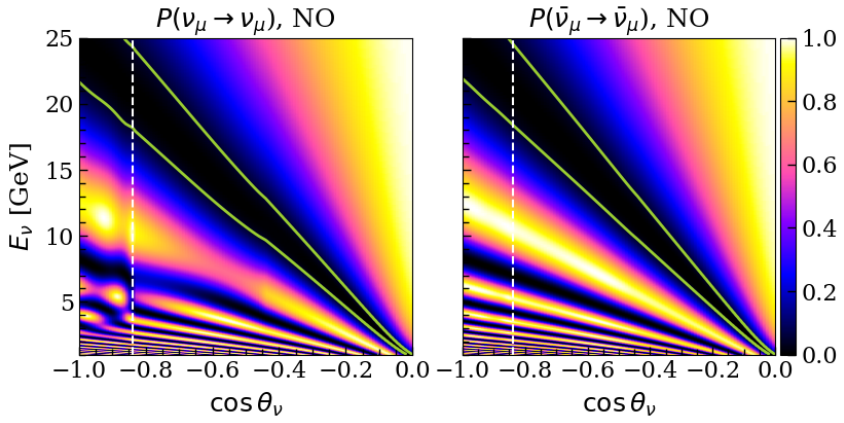


Figure 3.3: The $\nu_\mu \to \nu_\mu$ **(left)** and $\bar{\nu}_\mu \to \bar{\nu}_\mu$ **(right)** survival probability oscillograms for atmospheric neutrinos propagating through the Earth. The vertical axis shows neutrino energy (E_ν) and the horizontal axis shows the cosine of the zenith angle ($\cos\theta_\nu$), representing the path length through the PREM density profile. The dark regions indicate enhanced ν_μ disappearance arising from vacuum oscillations and matter effects, including MSW resonance in mantle-crossing trajectories and parametric resonance for core-crossing trajectories. Taken from [85].

through the Earth's core experience an abrupt increase in density. This layered density profile can induce parametric resonance [84], which further increases the transition probabilities.

As shown in the survival oscillogram in Figure 3.3, these combined Earth-matter effects manifest as distinct "valleys" of ν_μ disappearance. For trajectories crossing the core ($\cos\theta_z \approx -1$), this maximal disappearance occurs at energies near 25 GeV, providing physical motivation for the dense sensor geometry of IceCube Upgrade (Section 3.6).

3.2.2 The TeV–PeV Regime: Earth Absorption

As discussed in Section 2.3.1, for high-energy astrophysical neutrinos, the oscillation length scales astronomically, effectively freezing the flavor ratio over the diameter of the Earth. Instead of oscillating, these neutrinos are subject to the rising DIS cross section shown in Figure 3.2.

The characteristic distance a neutrino can travel before interacting is given by its interaction length $L_{\text{int}} = 1/(\sigma\rho N_A)$, where ρ is the density of the medium and N_A is Avogadro's number. At energies above approximately 100 TeV, the interaction cross section becomes so large that L_{int} drops below the diameter of the Earth. Thus, the Earth becomes largely opaque to up-going high-energy neutrinos [86]. This absorption effect restricts IceCube's field of view for the highest-energy astrophysical sources, suppressing the flux from the Northern Hemisphere and making the observatory primarily sensitive to the horizon and the downward-going Southern sky. However, this same attenuation provides a macroscopic method to directly measure the neutrino-nucleon cross section via Earth absorption at energies inaccessible to terrestrial particle accelerators [81].

Tau Regeneration

An exception to Earth absorption exists for tau neutrinos. When a high-energy ν_e or ν_μ undergoes a CC interaction in the bedrock beneath the detector, it is effectively absorbed,

as the resulting outgoing leptons (the electron or muon) will eventually deposit their energy and stop far below the instrumented volume. However, when a high-energy ν_τ undergoes a CC interaction, it produces a massive τ lepton that possesses a short lifetime (approximately 2.9×10^{-13} s). At TeV to PeV energies, the τ typically decays back into a ν_τ before substantial radiative energy losses occur. This weak decay occurs via both leptonic and hadronic channels [87]:

$$\tau^- \to \nu_\tau + \begin{cases} e^- + \bar{\nu}_e & (\sim 17.8\,\%) \\ \mu^- + \bar{\nu}_\mu & (\sim 17.4\,\%) \\ \text{hadrons (e.g., } \pi^-, \rho^-) & (\sim 64.8\,\%) \end{cases} \tag{3.3}$$

Therefore, regardless of the specific decay branch, rather than being completely absorbed, the original tau neutrino is "regenerated" at a lower energy. This successive cascading process allows ν_τ fluxes to penetrate through the Earth, creating a pile-up of regenerated tau neutrinos at lower energies and making the Earth intrinsically more transparent to tau neutrinos than to any other flavor [88].

3.3 Cherenkov Detector Operation Principle

Neutrino detection in large-volume observatories relies on an indirect measurement principle: the observation of charged secondary leptons and hadronic showers produced in neutrino interactions within the detector medium (Section 3.1). As these relativistic secondary charged particles propagate through a dielectric medium, they induce *Cherenkov radiation* [89], providing a detectable optical signature. Cherenkov radiation is emitted when a charged particle travels through a dielectric medium faster than the phase velocity of light in that medium. The threshold energy for a particle of rest mass m_0 to emit Cherenkov radiation is given by

$$E_{\text{th}} = \frac{m_0 c_0^2}{\sqrt{1 - 1/n^2}}, \tag{3.4}$$

where c_0 is the speed of light in vacuum and n is the refractive index of the medium [90]. Assuming $n \approx 1.33$ for ice, this threshold is approximately 0.8 MeV for electrons, 160 MeV for muons and 2.7 GeV for taus. In CC interactions, the outgoing lepton carries a significant fraction of the incident neutrino's energy. Hence, even neutrino interactions at MeV energies can produce secondary electrons above the Cherenkov threshold.

Microscopically, as the charged particle traverses the dielectric medium, its electromagnetic field temporarily polarizes the surrounding atoms. When the particle's velocity v exceeds the phase velocity of light in the medium (c_0/n), the subsequent depolarization of these atoms emits photons that constructively interfere. Based on Huygens' principle, this results in a coherent, cone-shaped macroscopic wavefront along the particle's trajectory, as illustrated in Figure 3.4. The emission angle θ_C relative to the particle track is defined by

$$\cos(\theta_C) = \frac{c_0}{nv} = \frac{1}{\beta n}. \tag{3.5}$$

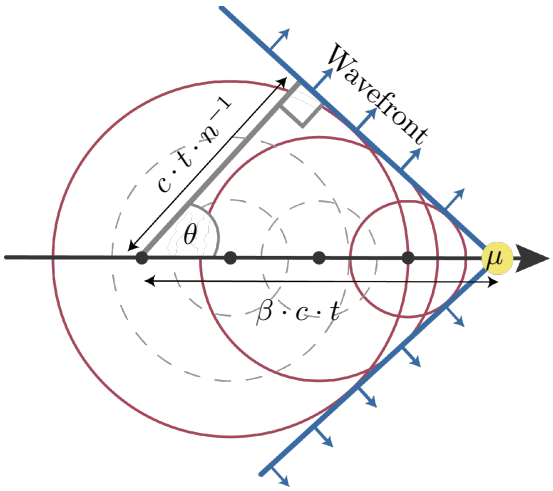


Figure 3.4: Schematic of Cherenkov radiation. A charged muon (black line) traverses a dielectric medium with $v > c_0/n$. The emission of electromagnetic radiation is shown as red circles, forming a coherent conical wavefront (blue). Taken from [91].

For relativistic particles in ice, the Cherenkov angle is approximately 41°. The number of Cherenkov photons N emitted per wavelength λ over a track length L is described by

$$\frac{dN}{d\lambda} = \frac{2\pi z^2 \alpha}{\lambda^2} L \sin^2(\theta_C), \tag{3.6}$$

derived from the Frank-Tamm formula [92], where z is the charge number of the incident particle and α is the fine-structure constant [93]. The $1/\lambda^2$ dependence dictates that the emission spectrum favors shorter wavelengths, particularly in the ultraviolet and blue regions of the spectrum. Because the flux of astrophysical neutrinos decreases rapidly towards high energies, observing the resulting Cherenkov photons at a statistically significant rate requires instrumenting immense volumes of a transparent dielectric medium. This naturally limits viable choices on Earth to the atmosphere (e.g., MAGIC [47]) and large bodies of water or ice, with this thesis focusing on ice-based detection. While deep glacial ice exhibits long absorption lengths (high transparency), its scattering length is comparatively short. Because both scattering (such as Rayleigh or Mie scattering) and absorption are wavelength-dependent, the initial $1/\lambda^2$ spectrum is effectively filtered as photons propagate, meaning the spectrum of photons that successfully reach the optical modules is skewed toward longer optical wavelengths.

3.4 The IceCube Neutrino Observatory

The IceCube Neutrino Observatory [3] is currently the largest detector in the world by instrumented volume. It encompasses an instrumented volume of one cubic kilometer within the transparent glacial ice at the geographic South Pole, using the glacier as a Cherenkov medium for indirect neutrino detection. As illustrated in Figure 3.5, the main optical array consists of 86 vertical strings deployed in holes drilled into the ice, with each string instrumenting 60 Digital Optical Modules (DOMs). These DOMs, which house photomultiplier tubes (PMTs) and onboard digitizing electronics in a transparent pressure vessel, are positioned at depths between 1450 m and 2450 m below the surface.

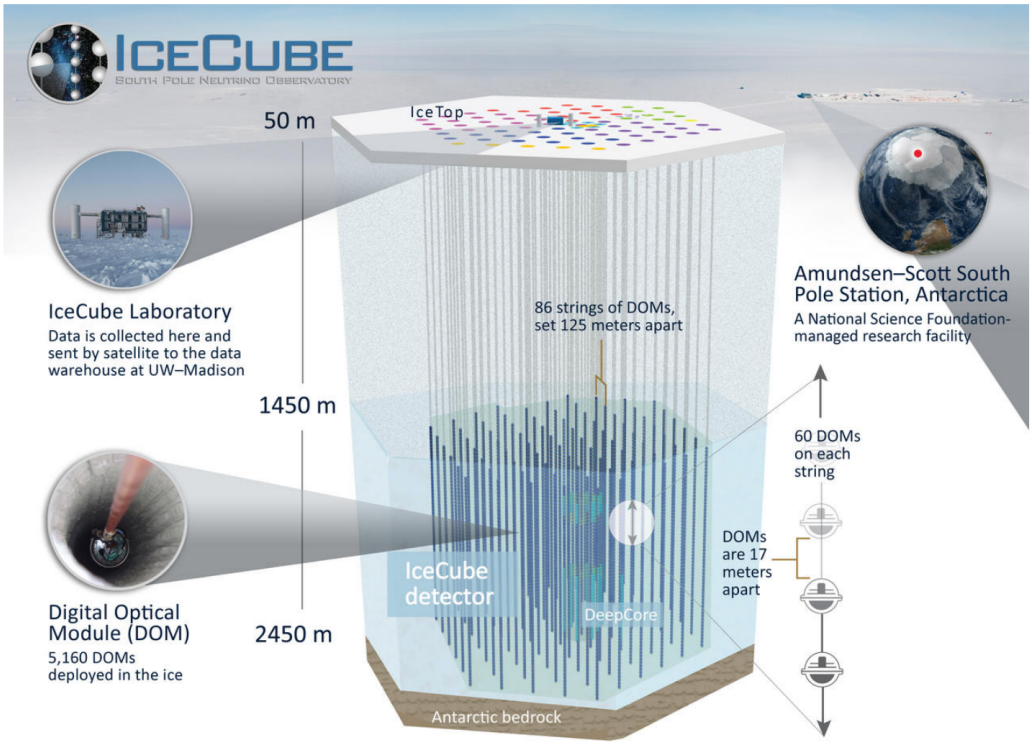


Figure 3.5: Schematic of the IceCube Neutrino Observatory. The deployment depth ends approximately 400 m above bedrock. Courtesy of the IceCube collaboration.

Strings are arranged in a hexagonal grid with a horizontal spacing of 125 m and a vertical inter-DOM spacing of 17 m [3]. This geometry optimizes the total instrumented volume for detecting high-energy astrophysical neutrinos, resulting in an energy threshold for the main array of approximately 100 GeV, with optimal sensitivity for neutrino energies $E_\nu \gtrsim 1\,\text{TeV}$ [94]. The deep ice above the detector provides a natural 1450 m thick shield that reduces the flux of downward-going atmospheric muons. Moreover, the scale of the main optical array allows its outermost strings and top layers to serve as an active veto region. By identifying penetrating atmospheric muons, this outer layer effectively shields the inner array, allowing for the isolation of neutrino interactions that occur within the fiducial volume.

Unlike natural open-water Cherenkov detectors, the deep glacial ice provides a static environment without dynamic currents. In addition, it is inherently free of bioluminescence and has much lower concentrations of radioactive isotopes, such as ^{40}K, that commonly contaminate seawater arrays. Thus, the intrinsic detector noise is introduced solely by the hardware itself, specifically the dark noise of the PMTs (Section 4.4) and scintillation occurring within the optical modules (Section 9.1). Distinct from this intrinsic instrumental noise, the dominant physics background for most astrophysical analyses consists of penetrating atmospheric muons and atmospheric neutrinos.

Systematic calibration studies [95] have revealed that the optical properties of the ice vary with depth. These variations closely track mineral dust concentrations deposited over millennia, which correlate with the Earth's climatic history. The ice with the lowest scattering and absorption in the detector, referred to as the *physics region*, is located at depths between 2100 m and 2450 m with an absorption length of ~ 200 m and scattering length of ~ 50 m at 400 nm [96]. Above this lies a prominent *dust layer* between 2000 m

and 2100 m, characterized by high scattering and absorption. These dust layers are not perfectly horizontal and exhibit a macroscopic spatial tilt across the detector footprint, driven by the glacial flow over the uneven underlying bedrock [96]. Lastly, the *bulk ice* displays optical anisotropy, where the scattering and absorption of Cherenkov photons exhibit a distinct directional dependence that is closely aligned with the local glacial flow axis [97 (pp.17–20)].

To increase IceCube's sensitivity to lower energies, a densely instrumented infill array called *DeepCore* [94] was deployed. DeepCore comprises eight dedicated strings, alongside surrounding standard IceCube strings that serve as a boundary. Six of these dedicated strings are equipped with high quantum efficiency PMTs (see Section 4.2). To extend sensitivities into the GeV regime, 50 of the 60 DOMs on each DeepCore string are deployed in the highly transparent physics region below the dust layer, with a vertical spacing of just 7 m. The remaining 10 DOMs of each string are positioned above the dust layer (between 1750 m and 1850 m) with a 10 m spacing, acting as an overhead veto against background tracks. The horizontal distance between DeepCore strings ranges from 42 m to 72 m, which corresponds to about 1.5 times the effective scattering length of the ice surrounding most of DeepCore. This dense configuration lowers the detector's energy threshold, enabling neutrino reconstruction down to energies of approximately 10 GeV.

Complementing the deep-ice array is *IceTop* [98], a surface air shower array designed to detect the electromagnetic and hadronic components of extensive air showers induced by cosmic rays. IceTop consists of 81 stations located near the surface positions of the IceCube strings. Each station comprises two light-tight ice Cherenkov tanks separated by 10 m, with two DOMs deployed in each tank. By requiring a trigger across three or more stations, IceTop can reconstruct cosmic ray air showers with primary energies extending from approximately 300 TeV up to the EeV range, probing the knee region of the cosmic ray spectrum (see Section 2.1). In the denser infill region located directly above DeepCore, this threshold is reduced to roughly 100 TeV. Crucially, IceTop also operates in coincidence with the optical array, providing a partial veto against atmospheric muons originating from cosmic ray showers.

All signals and power lines from both the in-ice strings and IceTop stations are routed to the IceCube Laboratory (ICL). The ICL is a surface building centrally located on the array where the data acquisition systems collect, timestamp, and pre-process the event data before it is filtered and transmitted via satellite for offline analysis (see Section 12.1 for details).

3.5 Event Morphology

When a neutrino interacts within the instrumented volume, the spatial and temporal distribution of the emitted Cherenkov light allows IceCube to distinguish between neutrino flavors and interaction types. Depending on the kinematics and the flavor of the involved neutrino, the observed Cherenkov light predominantly forms one of three distinct event morphologies: tracks, cascades or double bangs as shown in Figure 3.6. Because a fraction of the incident neutrino's energy is transferred to the target nucleon during deep inelastic scattering (see Section 3.1), every interaction inherently produces a localized shower of secondary hadrons at the interaction vertex. As a result, the distinct detector topologies described below are often a superposition of this initial hadronic cascade and the subsequent behavior of the outgoing lepton.

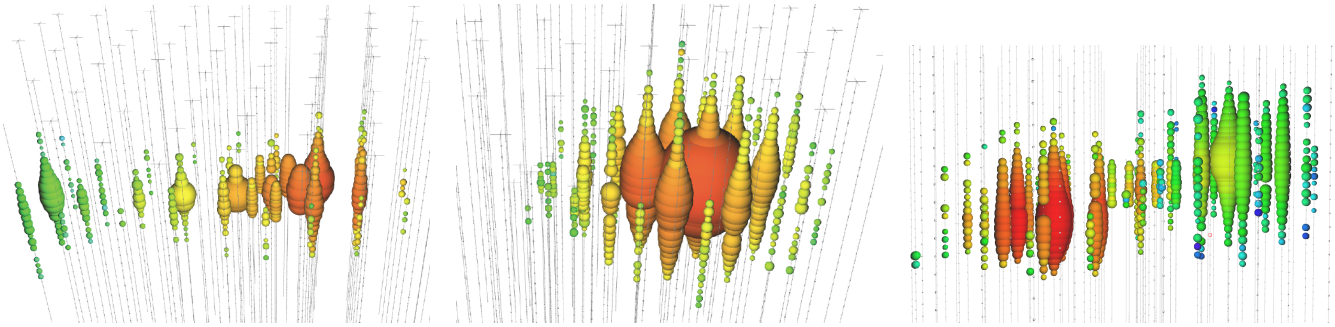


Figure 3.6: Cherenkov light morphologies in IceCube. Colors indicate photon arrival times, ranging from early (red) to late (blue), while bubble sizes represent the logarithmic light intensity. **(left)** A track-like event typical for ν_μ CC interactions. **(center)** A cascade event typical for ν_e CC, low-energy ν_τ CC and all NC interactions. **(right)** A double-bang topology characteristic of a high-energy ν_τ CC interaction, showing the interaction vertex and subsequent tau decay. Courtesy of the IceCube collaboration.

Tracks are produced by relativistic muons propagating through the detector. While often the result of ν_μ CC interactions, these muons can also be penetrating atmospheric muons or the decay products of tau leptons originating from ν_τ CC interactions. Because muons suffer relatively low radiative energy losses compared with electrons due to their much larger mass, they can propagate for kilometers through the ice at TeV energies. Their passage creates an extended trail of Cherenkov light, which becomes increasingly stochastic at higher energies due to discrete radiative losses such as bremsstrahlung and pair production. The extended lever arm of this track yields excellent angular resolution, typically 0.5° for events above 10 TeV [99]. However, the energy resolution for tracks depends heavily on event containment.

For typical through-going tracks, the muon originates outside the array and escapes the instrumented volume before stopping. In a ν_μ CC interaction, this results in the outgoing muon carrying away the lepton energy fraction $(1 - y)E_\nu$. Because the full track is uncontained, the total calorimetric energy cannot be measured. Instead, the specific energy loss dE/dx along the visible track is utilized to reconstruct the energy of the incident muon, providing a lower limit for the parent neutrino energy. This dE/dx-based reconstruction is more challenging and less precise than measuring total energy deposition. If the interaction vertex is contained within the volume (a starting track), the initial hadronic cascade energy (yE_ν) is calorimetrically measured alongside the track, providing a closer proxy to the true neutrino energy. Similarly, if a muon fully stops within the array, its total energy can be inferred calorimetrically.

While high-energy muon tracks are readily distinguishable, the reconstruction of low-energy starting tracks (such as those targeted by the GRECO sample introduced in Section 12.1.1) is impacted by event-by-event fluctuations in the inelasticity parameter y. In a high-y interaction, the majority of the neutrino's energy is channeled into the vertex hadronic cascade, leaving a low-energy, short-lived muon. The resulting faint muon track is obscured by the bright, localized vertex cascade, making the overall topology appear cascade-like and complicating flavor tagging.

Cascades (also known as showers) originate from electron neutrino CC interactions, most tau neutrino CC interactions and all NC interactions that produce localized particle showers. For example, an outgoing electron from a ν_e CC interaction will immediately initiate an electromagnetic cascade, while the disrupted nucleon at the interaction vertex

creates a hadronic cascade. Because Cherenkov photons undergo multiple scattering in the ice, the Cherenkov emission appears as a localized, nearly spherical "blob" of light. Since the entire cascade is often fully contained within the detector, these events offer more precise calorimetric energy measurements (often achieving an energy resolution better than 10 % for $> 10\,\text{TeV}$ [100, 101]). Crucially, the total visible energy depends on the interaction channel. While a contained ν_e CC interaction deposits nearly the full incident neutrino energy into visible secondaries within the instrumented volume, an NC interaction only deposits the hadronic fraction yE_ν, as the outgoing neutrino carries away the remaining $(1-y)E_\nu$ completely undetected. Conversely, the near-spherical pattern of the light pool washes out directional information, leading to a relatively poor angular resolution, although for contained events above 50 TeV it is typically better than 5° [102].

Double bangs are the characteristic signature of high-energy ν_τ CC interactions. The initial neutrino interaction creates a hadronic cascade (the first bang). The resulting tau lepton then travels a finite distance through the ice before decaying, initiating a second cascade (the second bang). The distance that the tau travels is proportional to its energy, given by $L \approx 50\,\text{m} \cdot (E_\tau/1\,\text{PeV})$ [103]. Because IceCube's optical modules in the bulk ice are spaced 17 m apart vertically and 125 m apart horizontally, this distinct separation is most clearly resolvable at PeV energies and above. Below this energy threshold, the tau decays too rapidly and the two cascades merge into a topology indistinguishable from a standard single cascade. Despite these geometric threshold challenges, IceCube recently isolated its first seven astrophysical ν_τ candidate events, marking a major milestone in high-energy flavor physics [103].

The exact topology of the second signature is dictated by the decay channel of the tau lepton (see Equation (3.3)). In roughly 82.6 % of cases, the tau decays either hadronically ($\sim 64.8\,\%$) or into an electron ($\sim 17.8\,\%$). Both of these channels produce a secondary localized shower, completing the classic double cascade or double-bang signature. However, in the remaining $\sim 17.4\,\%$ of cases, the tau decays into a muon ($\tau^- \to \mu^- + \bar{\nu}_\mu + \nu_\tau$). In this scenario, the second bang is replaced by an extended muon track emerging from the tau decay vertex, creating a distinct cascade-to-track topology.

Stochastic Radiative Losses

While low-energy muons traversing the ice behave primarily as minimum-ionizing particles with continuous and uniform energy losses, high-energy muons behave differently, as shown in Figure 3.7. Above the critical energy of approximately 1 TeV in ice, radiative energy loss mechanisms (specifically bremsstrahlung, electron-positron pair production and photonuclear interactions) begin to dominate over continuous ionization. These radiative losses are discrete and stochastic, producing localized electromagnetic and hadronic subcascades along the muon's trajectory. Thus, rather than a continuous line of Cherenkov emission, a high-energy track manifests as a series of bright, energetic bursts. By measuring the amplitude and frequency of these stochastic energy depositions, the detector samples the muon's specific energy loss, dE/dx. Because these discrete radiative processes dominate the TeV regime, this specific energy loss becomes approximately proportional to the muon energy ($dE/dx \propto E_\mu$), a behavior reflected in Figure 3.7. Exploiting this direct proportionality enables the reconstruction of the incident muon's energy, which is crucial for establishing a lower limit on the parent neutrino's energy even when the track is uncontained.

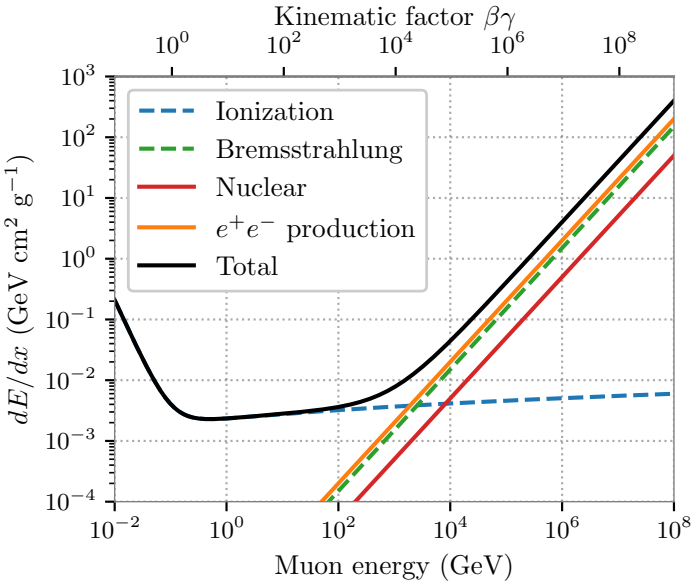


Figure 3.7: Average muon energy loss in water as a function of the muon energy. The total energy loss is shown together with the contributions from ionization, bremsstrahlung, photonuclear interactions and e^+e^- pair production. Adapted from [104 (p. 46)], using the parameterizations of [105].

3.6 IceCube Upgrade

The final drilling for the original IceCube array concluded with its 86th hole on December 18, 2010. Following over a decade of successful operation, the observatory underwent its next major extension under IceCube Upgrade. Fifteen years after the original array's completion, on December 20, 2025, the first new strings were deployed. While initially planned to comprise seven strings, the 2025/2026 deployment season successfully concluded with five newly instrumented strings integrated into the detector. As illustrated in Figure 3.8, this extension introduces an even denser infill compared with the existing DeepCore sub-array. The vertical spacing between individual sensors is reduced to approximately 3 m, with an inter-string distance of 20 m to 30 m relative to the surrounding strings, which is optimized to provide suitable baselines for calibration. Together, these five strings house approximately 500 newly developed optical modules.

The main science goals of IceCube Upgrade target improvements in the GeV energy regime and reduction of systematic uncertainty through advanced ice calibration [8]. A primary physics driver is the precise measurement of atmospheric neutrino oscillations. For up-going neutrinos with a path length comparable to the Earth's diameter, the first oscillation maximum for ν_μ disappearance and the corresponding peak for ν_τ appearance occur at approximately 25 GeV (see Section 3.2.1). While the existing DeepCore sub-array has probed this region, it lacks the necessary density to fully capture the kinematics of these low-energy interactions. IceCube Upgrade's densely instrumented optical modules are expected to yield a significant enhancement in event rates and reconstruction accuracy down to $\mathcal{O}(1\,\text{GeV})$. By enabling precision measurements of ν_τ appearance alongside ν_μ disappearance, IceCube Upgrade will serve as a test for the unitarity of the PMNS matrix (see Equation (2.7)) [106]. Any observed deviations from unitarity would provide compelling evidence for new physics beyond the Standard Model.

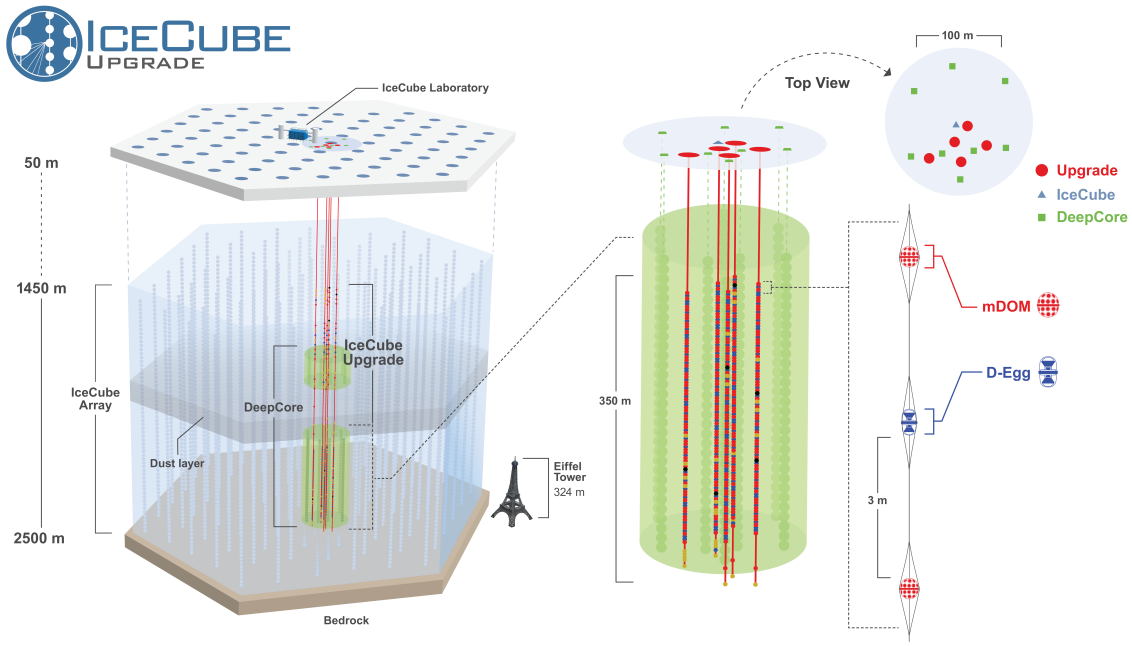


Figure 3.8: Schematic of IceCube Upgrade (red) embedded within DeepCore (green) and the main optical array (blue). Courtesy of the IceCube collaboration.

Additionally, the improved sensitivity and lowered energy threshold in the GeV regime are relevant for indirect dark matter searches. The new strings will enhance the detector's sensitivity to weakly interacting massive particle (WIMP) candidates postulated at these energies, which accumulate in the gravitational wells of the Sun or the Galactic Center [107].

In current neutrino analyses, the precise modeling of the bulk glacial ice optical properties constitutes a major source of systematic uncertainty. To better constrain the complex depth-dependent scattering, absorption and the aforementioned optical anisotropy, IceCube Upgrade deploys a comprehensive suite of novel in-module (acoustic emitter [108] and camera system [109]) as well as standalone (Pencil Beam [110], Sweden Camera [111], Acoustic Module [112] and POCAM [113]) calibration devices. A source of systematic uncertainty in the GeV energy region arises from the refrozen *hole ice*. This ice, which refreezes after the deployment of the detector strings, forms a bubble column in the central region of the deployment holes that scatters photons significantly more than the surrounding bulk ice and is a major systematic for oscillation analyses. New calibration devices and strategies have been developed to specifically calibrate this refrozen ice [114]. Crucially, the improved calibration constants and ice models derived from IceCube Upgrade will also be applied to previously recorded IceCube data, improving the angular and spatial reconstruction of previously published astrophysical neutrino events. Finally, IceCube Upgrade operates as a technological testbed for IceCube-Gen2 (see Section 3.7), a future large-scale detection extension. It provides a platform to test instrumentation, verify large-scale calibration techniques and conduct in situ performance studies of new sensor designs.

Optical Modules

All of IceCube's optical modules inherently consist of photomultiplier tubes (PMTs) housed alongside readout electronics in a transparent, watertight pressure vessel. The extreme environmental conditions dictate this baseline design, as the water in the deployment boreholes refreezes from the surface down to the bedrock, experiencing surface temperatures below $-40\,°\mathrm{C}$, with a temperature gradient reaching up to $-8\,°\mathrm{C}$ at depth [115]. Optical modules are subjected to high physical stress, with pressures reaching up to 690 bar [3]. Thus, the internal components are enclosed within a pressure vessel. To withstand this pressure while maintaining transparency for the surviving Cherenkov photons (which shift toward longer optical wavelengths due to scattering and absorption in the ice), these pressure vessels are constructed from borosilicate glass.

Inside an optical module, a holding structure (often mechanically mounted to the pressure vessel glass) sets the position of the PMTs. While the internal volume is flushed with dry nitrogen to protect the electronics, creating no direct optical interface between the PMT glass and the pressure vessel glass would result in internal reflection and photon loss. To seamlessly bridge this interface, a transparent optical gel with a refractive index matched to the pressure vessel is utilized. The PMTs are operated via an active base and the readout electronics are routed through a penetrator cable assembly (PCA) to the main data cable. During deployment, optical modules are secured with a harness to the string, which is lowered from the surface into the deep ice.

For IceCube Upgrade, the primary optical modules installed in the physics region are the multi-PMT digital optical module (mDOM) [116, 117] and the dual optical sensors in an ellipsoid glass intended for Gen2 (D-Egg) [118, 119]. Both designs are engineered to improve photon detection efficiencies as well as enhance the calibration capability of the detector and are structurally compared with the legacy DOM in Figure 3.9. The standard DOM consists of a single, downward-facing 10-inch PMT. In contrast, D-Eggs utilize an elongated pressure vessel containing two 8-inch PMTs directed towards both module poles. The mDOM features 24 smaller 3-inch PMTs distributed nearly uniformly over the pressure vessel.

Regarding specific design nuances, standard DOMs and pDOMs (which are refurbished DOMs) include a mu-metal cage surrounding the PMT to provide magnetic shielding, as the response of large-aperture PMTs is sensitive to their orientation within the Earth's magnetic field. During D-Egg assembly, the optical gel is poured to the edge

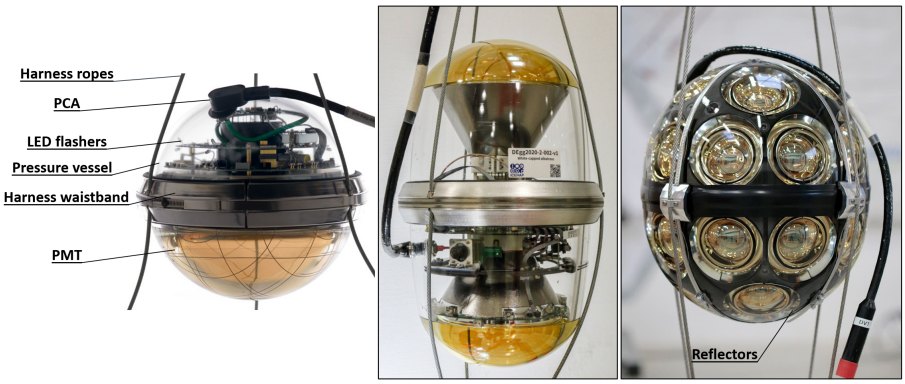


Figure 3.9: Comparison of the structural designs of the DOM (left), D-Egg (center) and mDOM (right). Courtesy of the IceCube collaboration.

of the PMT photocathode, similar to the DOMs, but the design intentionally exchanges the mu-metal shielding with FineMet cones along the PMT's backside. Furthermore, both mDOMs and D-Eggs house camera modules (including a blue LED to illuminate the surroundings) [120] to document ice formation during and after deployment, as well as LED flashers [121] to measure optical scattering and absorption. The multi-PMT design philosophy of the mDOM draws inspiration from the optical modules developed for KM3NeT [122] utilizing multiple smaller PMTs with reflectors to maximize photon collection. Transitioning to a multi-PMT architecture significantly enhances module performance by distributing an expanded effective photocathode area across numerous smaller tubes, achieving a nearly homogeneous 4π angular acceptance and intrinsic directional sensitivity. Additionally, this cathode segmentation naturally extends the dynamic range, simplifies discrete photon counting and enables the effective suppression of uncorrelated noise (see Section 4.4.1 and Chapter 9) through local coincidence triggers between individual PMTs within the same module.

Alongside mDOMs and D-Eggs, IceCube Upgrade incorporates multiple additional sparsely deployed optical modules. These include WOMs [123], pDOMs [124] and the LOM16/LOM18 sensor candidates, which act as prototypes for the future large-scale extension, IceCube-Gen2 (see Chapter 5). These new types of optical sensors are specifically developed to fit into the narrower boreholes envisioned for IceCube-Gen2, aiming to reduce drilling costs compared with the current 60 cm diameter holes [125]. While common design philosophies are shared across optical modules, each prototype introduces specific design considerations, necessitating independent research, development cycles and engineering solutions.

3.7 IceCube-Gen2

The future planned detector extension is called IceCube-Gen2 [9, 10, 126] and represents a major scale expansion designed to push the detector's capabilities across an unprecedented energy spectrum. As illustrated in Figure 3.10, this extension is not merely a volumetric increase of the deep-ice optical array from $1\,\mathrm{km}^3$ to approximately $8\,\mathrm{km}^3$. It broadens the observatory's multi-messenger reach by extending the cosmic ray surface array and deploying a new radio array. With these integrated components, combined with the low-energy extensions of IceCube Upgrade, the complete detector system will span an energy regime ranging from $\mathcal{O}(1\,\mathrm{GeV})$ to beyond 100 EeV. The core IceCube-Gen2 optical array is designed to increase the detection rate of cosmic TeV–PeV neutrinos by roughly a factor of ten, providing the sensitivity to detect astrophysical sources that are five times fainter than current limits [10]. To achieve this, 120 new strings will be deployed in an optimized "Sunflower" configuration, departing from the previous hexagonal grid. The hexagonal geometry of the legacy array inherently contains symmetric axes that create "corridors" through the instrumented volume, which can allow atmospheric muons to penetrate deeply with fewer optical module hits, resulting in a pronounced azimuth-dependent detector response. The Sunflower geometry breaks these long lines of symmetry, effectively eliminating these muon corridors and ensuring a more uniform, isotropic azimuthal acceptance. The strings are spaced 240 m apart to maximize the effective volume for high-energy events. Each string encompasses 80 optical modules and covers a larger vertical depth than the existing detector; the instrumentation spans from 1369 m (a boundary determined by ice scattering properties) down to 2689 m, approaching

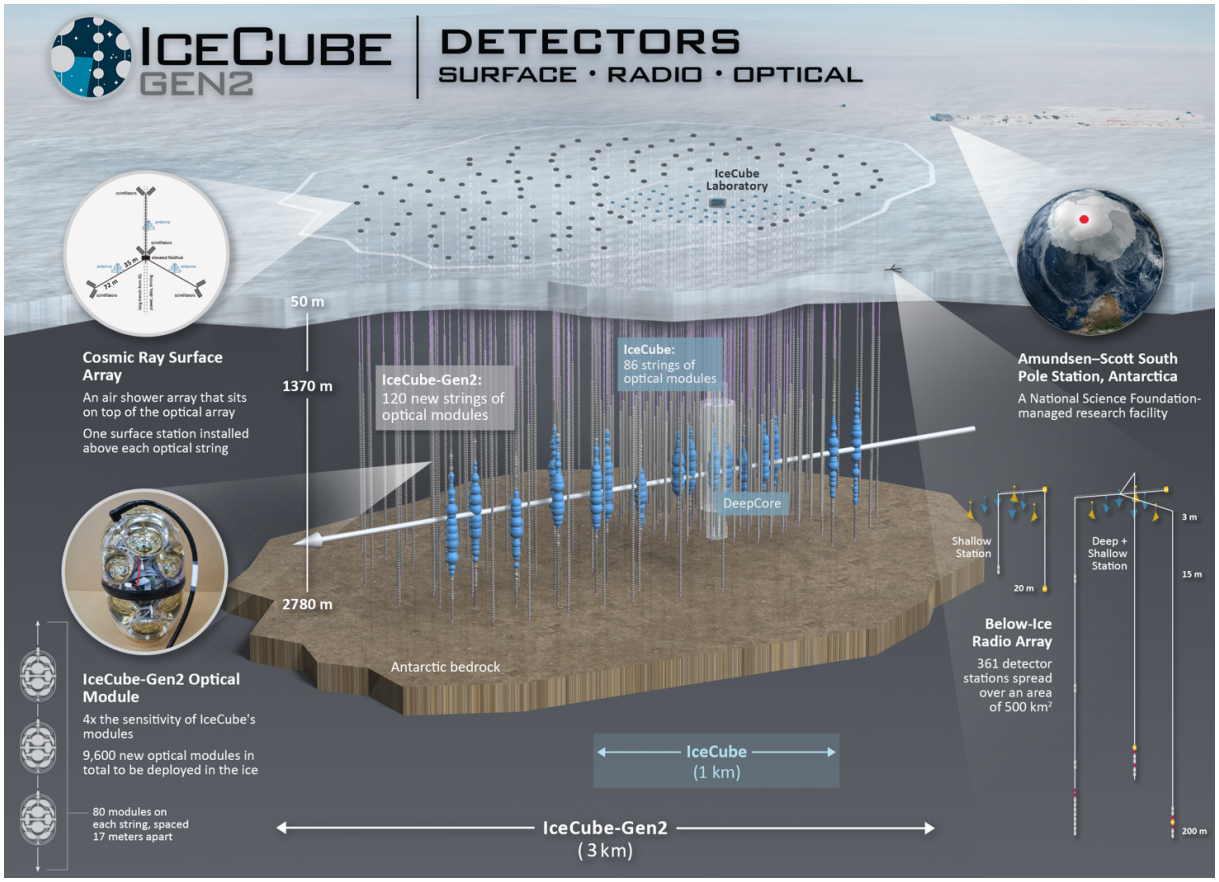


Figure 3.10: Overview of the IceCube-Gen2 layout, illustrating the extended deep-ice optical array, the surface array and the shallow radio array in the vicinity of the main array. Courtesy of the IceCube-Gen2 collaboration.

the bedrock estimated at 2780 m within the IceCube-Gen2 region. By capturing longer muon tracks over this expanded volume, the array will improve the angular resolution for horizontally passing events by roughly a factor of three [10].

The proposed optical modules for the future optical array play a central role in the second part of this thesis, with a detailed conceptual introduction in Chapter 5. Building on the development of IceCube Upgrade's mDOM and D-Egg modules, two new IceCube-Gen2 optical module design candidates (LOM16 and LOM18) were developed. These modules integrate successful features of their predecessors alongside a novel gel pad design. Five LOM16s and four LOM18s were successfully deployed within IceCube Upgrade for in situ testing. Following a development cycle that began in late 2020, the project is now entering a new phase of design iterations to scale up and ultimately mass-produce nearly 10 000 optical modules over a seven-year production phase. A technical advancement in these modules is the PMT readout mechanism. With the ability to extract PMT signals from the second-to-last dynode rather than just the anode, the optical modules achieve a vastly improved dynamic range for bright events, increasing the saturation threshold up to $\sim 5000\,\mathrm{PE}/25\,\mathrm{ns}$ per PMT [127].

Above the deep optical strings sits the extended cosmic ray surface array, spanning a footprint of approximately $8\,\mathrm{km}^2$. The array comprises 120 individual stations, with one positioned directly above each string. Each station is equipped with eight scintillation

panels and three radio antennas. These antennas are elevated to mitigate snow accumulation and are designed to be periodically raised to adapt to the changing surface level. Compared with the existing IceTop, this configuration yields a 30-fold increase in the aperture for cosmic ray air showers measured in coincidence with the deep-ice array. This serves a dual purpose. First, the surface array acts as a veto against cosmic ray-induced atmospheric muons, reducing the background for high-energy astrophysical neutrinos in the optical array. Second, it operates as a standalone observatory for PeV–EeV cosmic rays. By analyzing surface air showers alongside deep-ice muon tracks, the mass composition of primary cosmic rays can be probed, enabling a more precise mapping of the transition from Galactic to extragalactic sources [128].

While the optical array probes neutrinos up to the PeV scale, the IceCube-Gen2 radio array pushes the detector's reach into the ultra-high-energy (UHE) regime to detect exceedingly rare cosmogenic neutrinos (see Section 2.1). It achieves this by relying on the Askaryan effect, a macroscopic counterpart to optical Cherenkov emission [129]. When a UHE neutrino interacts in the ice, it initiates an energetic cascade. As the secondary electromagnetic shower develops, it sweeps up atomic electrons from the surrounding medium via Compton scattering, while positrons generated within the shower rapidly annihilate. This creates a compact, net negative charge excess, typically $\sim 20\,\%$ at shower maximum [130]. Because the longitudinal thickness of this charge bunch is smaller than radio wavelengths (spanning roughly 200 MHz to 1200 MHz [130]), the resulting Cherenkov radiation from the shower particles is macroscopically coherent. The emitted radio power scales quadratically with shower energy ($P \propto E^2$), producing a nanosecond-long pulse that can travel for over a kilometer in the radio-transparent ice [131]. To capture these transient pulses, the array exploits this kilometer-scale attenuation length by deploying a sparse $500\,\mathrm{km}^2$ network of shallow and hybrid stations. Shallow stations utilize near-surface log-periodic dipole antennas (LPDAs), including upward-facing LPDAs tasked specifically with vetoing background radio emissions from cosmic ray air showers. Hybrid stations augment these surface antennas with deep, omni-directional cylindrical antennas lowered 150 m into the firn [132].

4 Photomultiplier Tubes

Photomultiplier tubes (PMTs) are highly sensitive sensors capable of detecting individual photons with ns time resolution at low intrinsic background, making them indispensable across a wide range of particle physics experiments. Their applications span direct dark matter detection (XENON, PandaX), neutrino physics (Daya Bay, JUNO), gamma-ray astronomy (MAGIC, CTA) and cosmic ray physics (Pierre Auger, HAWC), as well as Cherenkov detectors (Super-Kamiokande, IceCube). In the context of IceCube, PMTs serve as the central sensor technology of the detector. As the primary components responsible for signal capture, a precise understanding of their characteristics is essential for accurate event reconstruction and the control of systematic uncertainties. Accordingly, the detailed characterization and simulation of these devices play a central role in the second part of this thesis. The following sections outline the operating principles and physical characteristics of these devices. Unless indicated otherwise, the general theory presented here is derived from standard literature [133, 134, 135].

4.1 Operation Principle

Fundamentally, a PMT operates by converting an incident photon into a photoelectron via the photoelectric effect. This single photoelectron is then multiplied by a factor of $\mathcal{O}(10^5 - 10^8)$ through a chain of secondary electron emissions to produce a measurable electrical signal at the anode. A visualization of this process is shown in Figure 4.1. The functional structure of a PMT consists of several key components:

Glass Envelope: This transparent enclosure defines the physical geometry of the PMT and maintains the vacuum required for free electron propagation. The specific glass type determines the transmission cutoff and thus limits the spectral range. PMTs are typically categorized by the outer diameter of the glass bulb at the widest point near the photocathode edge (e.g., a "4-inch PMT").

Photocathode: A thin, semi-transparent photosensitive layer (typically K_2CsSb or Rb_2CsSb) evaporated onto the inner surface of the entry window. It is responsible for the initial conversion of photons into photoelectrons.

Focusing Electrodes: Located between the photocathode and the first dynode, these components shape the electric field to improve photoelectron pathing (collection efficiency). A conductive coating on the PMT "neck" provides a connection to supply the photocathode potential while electrostatically focusing electrons towards the central axis. A front plate defines the entrance aperture to the dynode system, often supplemented by a focusing mesh (or "spider") to further homogenize the potential at the entrance of the first dynode.

Dynode Chain: This system is responsible for electron multiplication. It consists of a series of electrodes (dynodes) whose surfaces are designed for secondary electron emission. Common emissive materials include alkali antimonides (e.g., SbCs), beryllium oxide (BeO) or magnesium oxide (MgO) formed on a CuBe substrate. Because alkali vapors (e.g., Cs, K) migrate and condense on nearby internal structures during photocathode deposition, the first dynode inevitably becomes photosensitive. The

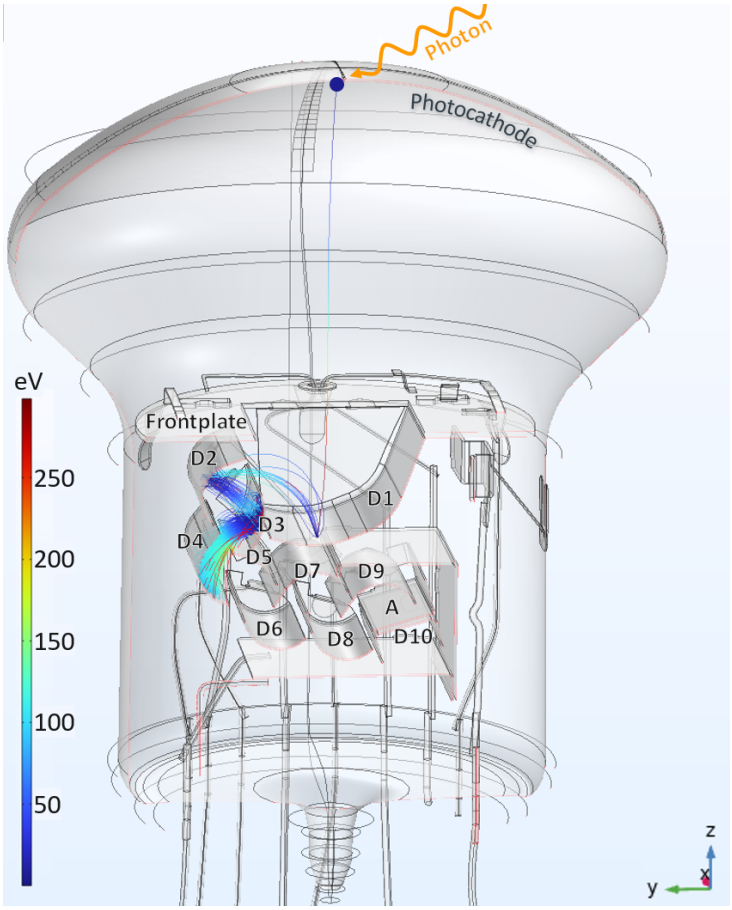


Figure 4.1: Visualization of the PMT multiplication process based on a simulation developed in [136, 137, 138] under the author's supervision. An incident photon (orange) strikes the photocathode, releasing a photoelectron (blue dot). This photoelectron is accelerated towards the first dynode (D1), where kinetic energy is indicated by color. Upon impact, multiple secondary electrons are released and guided through the dynode chain, undergoing successive multiplication steps. Note that the anode geometry (A) is simplified in this model, effectively occluding the tenth dynode, and the mesh in the dynode entry window is missing. A longitudinal cut plane is applied and electron propagation is shown only up to the fourth dynode (D4).

curvature of the dynode surfaces is optimized to guide secondary electrons efficiently to the subsequent stage. While dynode side walls provide mechanical support and limit stray electrons, active multiplication mainly occurs on the coated curved surfaces. Uncoated regions (such as bare copper-beryllium alloy or stainless steel) can also emit secondary electrons, but their secondary emission coefficient is lower [139]. The PMTs investigated in this work utilize a 10-stage dynode structure.

Anode: Unlike the dynodes, the final electrode in the chain is not coated for secondary emission but serves to collect the aggregate charge of the secondary electron cloud. This readout signal has a magnitude proportional to the total number of secondary electrons (gain).

Feedthrough Pins: Electrical connections penetrate the glass base to supply voltage to the internal components. Each internal component is held at a specific electric potential to shape the internal fields that govern electron trajectories. An external voltage divider circuit (base) is attached to these pins to distribute potentials across

components, where the distribution is defined by a divider ratio. The PMTs in this work utilize a ratio of $3 : 1 : \cdots : 1$, strengthening the field between the photocathode and the first dynode by a factor of 3 relative to the inter-dynode gaps.

Process Components: Additional structures, such as alkali sources used for photocathode deposition, are present on the front plate and near the inner wall of the glass envelope. Deposition occurs after evacuation and sealing to prevent evaporation of the photocathode layer.

The PMT response is sensitive to numerous external factors:

- **Light Intensity:** At high photon fluxes, pulse pile-up occurs. In extreme cases, space-charge effects can lead to saturation, deviating from a linear response.
- **Wavelength:** The incident photon energy determines both the conversion probability (quantum efficiency) and the initial kinetic energy of the photoelectron, influencing (photo)electron trajectories.
- **Position and Angle of incident Light:** The photoelectron path and, thus, collection efficiency vary depending on where and at which angle the photon strikes the photocathode, altering the electric field experienced by the photoelectron.
- **External Fields:** Even the Earth's magnetic field can deflect electron trajectories, affecting gain and collection efficiency depending on the orientation of the PMT.
- **Temperature:** Thermal energy influences the secondary emission coefficient of the dynodes as well as the background (namely thermionic emission).

Readout is typically performed in one of two modes. In **pulse mode**, the voltage drop across a load resistor is digitized, typically using an oscilloscope or flash ADC. As illustrated in Figure 4.2 (left), the anode signal from a single photoelectron is not an instantaneous delta function but a broadened pulse. The arrival of the pulse is delayed relative to the photon interaction by the **transit time (TT)**. The shape, characterized by parameters such as rise time (RT), fall time (FT) and pulse width (in terms of FWHM),

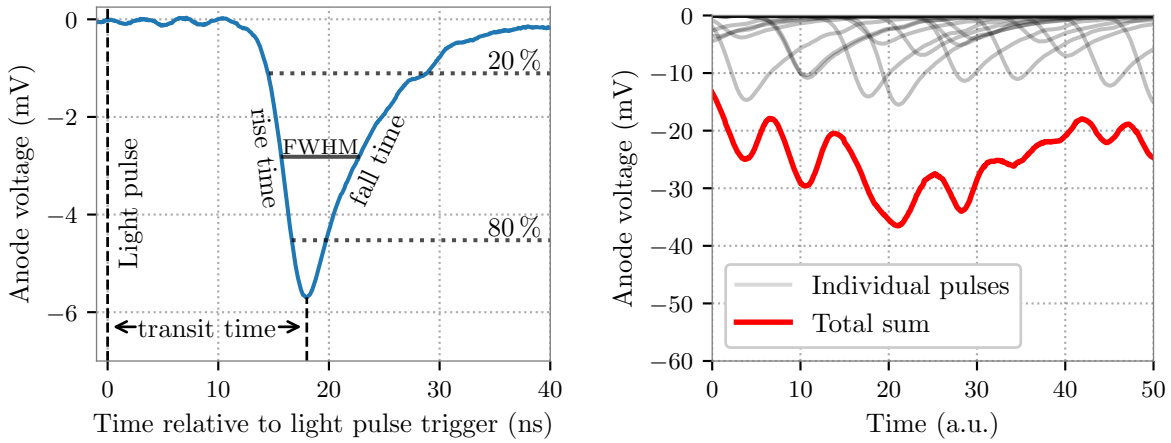


Figure 4.2: (left) Average single-photoelectron (SPE) pulse from a PMT used in this work, measured with a passive base. The horizontal lines indicate the time intervals used to determine the pulse parameters. The transit time is not to scale for illustrative purposes. **(right)** Illustration of pulse pile-up. Individual pulses (gray), randomly sampled from the SPE pulse shown in the left panel, overlap to form the total continuous signal (red).

is determined by the geometry of the multiplication system due to the spread in transit times of the secondary electrons as they move through the dynodes and is generally non-Gaussian [140].

Even for a steady, monochromatic light source emitting single photons, the peak transit time varies from photon to photon due to stochastic differences in photoelectron emission angle, position and initial energy, resulting in different electron trajectories. This jitter is quantified as the **transit time spread (TTS)** and defines the timing resolution of the PMT. While TTS is commonly approximated as the FWHM of a Gaussian fitted to an arrival time histogram, the true distribution is often asymmetric, especially when illuminating the entire photocathode. At higher light intensities, the apparent timing resolution improves due to the statistical averaging of multiple photoelectrons arriving within the resolution window.

In **current mode**, typically used for high-intensity or continuous light sources where individual pulses may overlap (see Figure 4.2, right), the anode current is integrated over a specific time window. The signal is treated as a continuous direct current (DC) proportional to the incident photon flux, typically measured with a picoammeter.

4.2 Photocathode Properties

The photocathode is a thin semiconductor layer, typically $\mathcal{O}(20\,\text{nm})$ thick, deposited on the inner surface of the PMT entry window. Its function is to convert incident photons into free electrons via the external photoelectric effect. For IceCube, which requires high sensitivity in the 300 nm to 600 nm spectral range[1], bialkali photocathodes composed of K_2CsSb are the standard choice due to their low thermionic noise and favorable band structure. The conversion process is described by the *Spicer three-step model* [141, 142]:

1. **Absorption:** An incident photon with energy $h\nu$ is absorbed within the lattice, exciting an electron from the valence band to the conduction band. This effectively corresponds to the internal photoelectric effect.
2. **Electron Transport:** The excited electron diffuses towards the vacuum interface. During this transport, it undergoes scattering (primarily electron-phonon and electron-electron), losing kinetic energy. The probability of reaching the surface depends on the absorption depth relative to the electron's inelastic mean free path.
3. **Surface Escape:** To exit into the vacuum, the electron must have sufficient residual energy to overcome the surface potential barrier (electron affinity E_A). The kinetic energy E_{pe} of the emitted photoelectron is given by:

$$E_{\text{pe}} = h\nu - E_g - E_A - E_{\text{loss}}, \tag{4.1}$$

where E_g is the band gap energy and E_{loss} represents energy dissipated during transport. Consequently, even monochromatic light produces a spread of initial photoelectron energies, which contributes to the transit time spread discussed in Section 4.1.

[1] Although Cherenkov photons are generally produced in the UV range, as described by Equation (3.6), the effective wavelengths reaching optical modules are shifted to longer wavelengths due to scattering in ice and depend on the photon propagation distance with respect to the interaction vertex. 400 nm is often quoted as the peak value of the distance-dependent scattered wavelength distributions.

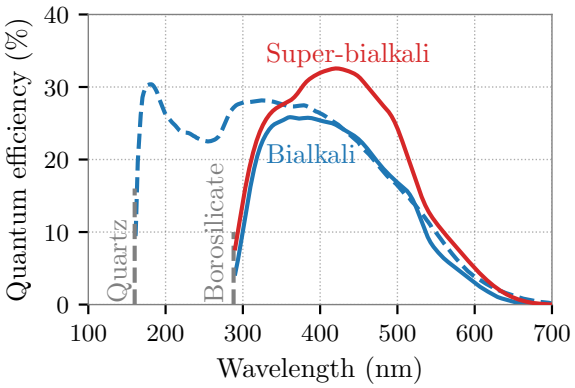


Figure 4.3: Spectral response of different photocathode types (colored lines) and the transmission limits of common window materials (gray dashed lines) of PMTs measured in this thesis.

The efficiency of this process is quantified by the **quantum efficiency (QE)**, defined as the ratio of emitted photoelectrons N_{pe} to incident photons N_γ:

$$\text{QE}(\lambda) = \frac{N_{\text{pe}}}{N_\gamma} = S_k \cdot \frac{hc}{e\lambda}, \tag{4.2}$$

where manufacturers often specify the radiant sensitivity S_k (in A/W). As shown in Figure 4.3, the spectral response is bounded by two physical limits:

- **Cutoff:** Defined by the transmission of the glass window material. Standard borosilicate glass, used in IceCube PMTs for its mechanical durability and low dark rates, becomes opaque to UV light ($\lesssim 280\,\text{nm}$).
- **Threshold:** Determined by the band gap and electron affinity of the semiconductor ($h\nu < E_g + E_A$). For bialkali materials, this threshold occurs in the near-IR region ($\gtrsim 700\,\text{nm}$).

Standard bialkali photocathodes achieve peak QEs of roughly 25 %. Improving the photocathode crystallinity and surface purity leads to *high-QE* or *super-bialkali* variants [143] that can reach peak efficiencies exceeding 35 % (as used in high-QE DeepCore DOM and D-Egg PMTs).

4.3 Gain and Charge Spectrum

The amplification process within the dynode chain is governed by the **secondary emission factor** δ, defined as the average number of secondary electrons emitted per incident primary electron. This factor depends on the kinetic energy of the incident electron and the surface properties of the dynode material. It follows a power law relationship with the inter-stage voltage v [135 (p. 205)]:

$$\delta = a \cdot v^k, \tag{4.3}$$

where a is a material-specific constant and the exponent $k \approx 0.7$–0.8 characterizes the surface emission properties. For the PMTs investigated in this work, the higher potential

difference applied between the photocathode and the first dynode results in a higher first-stage gain of $\delta_1 \approx 12$–14, while subsequent stages typically exhibit $\delta_i \approx 3$–7 [139].

The total number of secondary electrons reaching the anode per photoelectron is called the **gain** G. It is the product of the secondary emission factors of all N dynodes and the **collection efficiency** (CE) at each stage:

$$G = \prod_{i=1}^{N} \mathrm{CE}_i \cdot \delta_i. \tag{4.4}$$

The CE is defined as the ratio of electrons emitted from stage $i-1$ that successfully arrive on stage i. The total CE is dominated by the first dynode (CE_1), which is typically $> 90\,\%$, and subsequent stages generally achieve CE $> 99\,\%$. The gain exhibits a power law dependence on the total supply voltage V, making it highly sensitive to voltage variations. In practice, this relationship is fitted as:

$$G(V) = a \cdot V^b. \tag{4.5}$$

In typical characterization setups, the PMT is illuminated by a pulsed light source (e.g., an LED or laser) whose photon emission follows Poisson statistics. Thus, the number of initial photoelectrons generated per trigger is Poisson-distributed. Coupled with the statistical nature of the secondary emission process, the total number of electrons reaching the anode varies with each event. The resulting distribution of integrated charge pulses is the *charge spectrum*, measured in pulse mode, as illustrated in Figure 4.4 (left). For a Poissonian light source, this spectrum consists of three distinct components:

- **The Pedestal (0 PE):** In the absence of a photon conversion (the $n = 0$ case of the Poisson distribution), the integration of the electronic baseline determines the zero-charge reference. This forms a Gaussian peak dominated by electronic noise, referred to as the pedestal.
- **Single Photoelectron (SPE) Peak:** The primary signal resulting from exactly one photon detection ($n = 1$).
- **Multi-Photoelectron (MPE) Peaks:** Due to the Poisson statistics of the light source, there is a distinct probability of generating multiple photoelectrons ($n \geq 2$) within a single time window, even at a steady, low-intensity illumination. As the mean light intensity increases, the relative frequency of these MPE events grows. The resulting secondary electron avalanches overlap at the anode, forming a summed signal. Since PMTs are linear until saturation (typically a few hundred PE), the mean charge of an n-PE event corresponds to n times the SPE charge.

Assuming this Poisson-distributed illumination, the total charge response is modeled using an analytical function derived in [144], with modifications for the under-amplified tail described in [73 (Appendix A)]. The probability density function $S(q)$ for the observed charge q is given by:

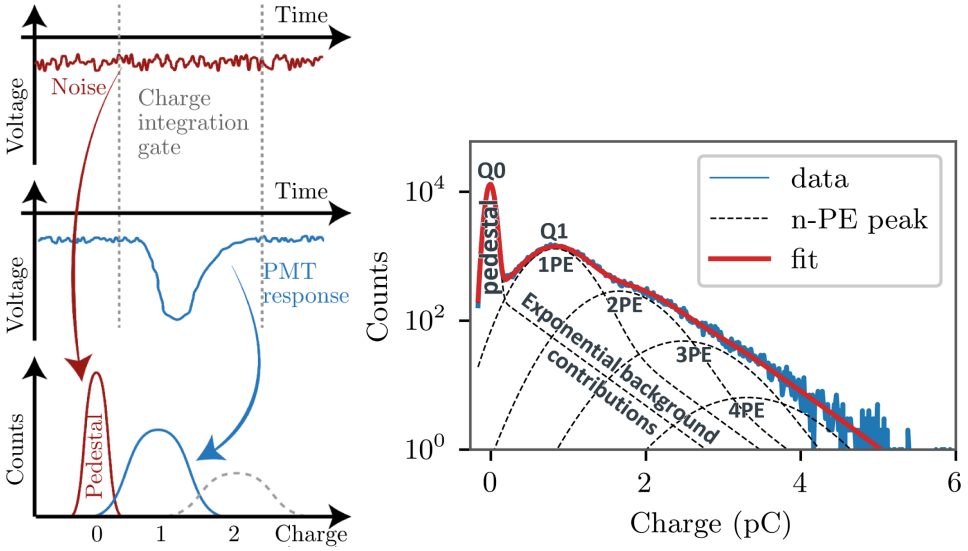


Figure 4.4: (left) Conceptual sketch of the charge spectrum acquisition. Taken from [91]. **(right)** Charge spectrum measured in this work and fitted with the model described by Equation (4.6).

$$S(q) = \sum_{n=0}^{\infty} \underbrace{\frac{\mu^n \cdot e^{-\mu}}{n!}}_{\text{Poisson emission}} \cdot \Bigg[\underbrace{(1-P_u) \cdot \mathcal{G}(q, Q_0 + nQ_1, \sqrt{\sigma_0^2 + n\sigma_1^2})}_{n\text{-PE peak}} + \underbrace{\frac{P_u \lambda}{2} \cdot e^{-\lambda\left(q - Q_0 + nQ_1 - \lambda n \sigma_1^2/2\right)} \cdot \left[1 - \operatorname{erf}\left(\frac{Q_0 + nQ_1 - q + \lambda n \sigma_1}{\sigma_1 \sqrt{2n}}\right)\right]}_{\text{Under-amplified Tail}} \Bigg]. \tag{4.6}$$

Here, the response is a superposition of Poisson-weighted contributions from n photoelectrons (where μ is the average Poissonian light intensity in PE). The deposited charge of n photoelectrons is modeled by a Gaussian $\mathcal{G}$ with mean $Q_n = Q_0 + n \cdot Q_1$ and width $\sigma_n = \sqrt{\sigma_0^2 + n\sigma_1^2}$, where Q_0 and σ_0 represent the pedestal mean and width. With probability P_u, a (photo)electron may backscatter or miss a dynode, resulting in an under-amplified signal. This is modeled by the convolution of an exponential $\lambda e^{-\lambda q}$ with the Gaussian resolution, represented by the term involving the complementary error function.

The gain is determined from the fitted parameter Q_1 (the mean charge of a single photoelectron) via $G = Q_1/e$. PMTs are typically operated at a **nominal gain**, achieved by adjusting the supply voltage until the gain matches the target value. For the PMTs in this work, the target gain is set to 5×10^6 by the manufacturer. Additionally, the fit yields the *SPE resolution* (Q_1/σ_1) to quantify the energy resolution of the dynode chain, while the *peak-to-valley ratio* (extracted directly from the charge spectrum) serves as a figure of merit for discriminating single-photon signals from electronic noise.

4.4 Background

As PMTs are highly sensitive detectors capable of resolving single photons, they are inherently susceptible to background noise. Even in a completely light-tight environment, pulses are measurable at the anode. These background sources are generally categorized into two classes: intrinsic uncorrelated background, which occurs randomly in the absence of light, and signal-induced correlated background, which is prompted by photon detection.

4.4.1 In Darkness (uncorrelated Background)

The dark count rate (or dark current) is the superposition of several physical mechanisms, distinguishable by their amplitude and voltage or temperature dependence:

- **Thermionic Emission:** The photocathode material is chosen for its low work function to maximize quantum efficiency, which inevitably allows thermal electrons to escape the lattice energy barrier according to the Richardson-Dushman equation [145]:
$$j_s = AT^2 \cdot e^{-\frac{W}{k_B T}}, \tag{4.7}$$
where j_s is the current density of the emission ($\mathrm{mA/mm^2}$), $A \approx 1202\,\mathrm{mA/mm^2/K^2}$ is Richardson's constant (multiplied by a material-specific correction factor), W is the work function of the photocathode, k_B is the Boltzmann constant and T is the temperature. These electrons undergo the full amplification chain, resulting in SPE-equivalent pulses. This process is temperature-dependent, reducing significantly at the environmental temperatures of IceCube from $-40\,^\circ\mathrm{C}$ to $-10\,^\circ\mathrm{C}$.
Notably, thermionic emission also occurs at the dynodes, as they are coated with the same photosensitive material. However, the photocathode has a larger surface area and electrons emitted from dynodes further down the chain experience less total amplification, resulting in sub-SPE pulses that are often below the trigger threshold, akin to prepulses (see Section 4.4.2). Therefore, the dark current caused by thermionic emission exhibits a dependency on the supply voltage that is nearly identical to the exponential slope of the gain.
- **Leakage Currents:** Ohmic leakage currents can form distinct conductive paths across the insulating base, the glass stem, or between the anode and the last dynode due to imperfect insulation, humidity or surface contamination. While this primarily contributes a DC offset relevant for current-mode measurements, it can manifest as noise pulses if breakdown occurs. Unlike thermionic emission, the leakage current increases linearly with the applied voltage.
- **Field Emission:** In the presence of excessive electric fields, electrons can tunnel through the potential barrier of electrode surfaces. This typically occurs at sharp geometries or imperfections within the dynode structure at the upper end of the rated supply voltage range. Since these electrons may originate from dynodes further down the chain, they often appear as under-amplified pulses. This contribution is independent of temperature but scales exponentially with the applied voltage. For typical nominal supply voltages used in this work, this effect does not dominate.
- **Radioactive Decays:** Trace radioactive isotopes within the glass envelope (predominantly ^{40}K, Th and U) decay, emitting gamma rays or alpha particles. These particles can produce Cherenkov radiation and scintillation light within the glass,

which can be detected by the photocathode. Additionally, stray electrons from the dynode system can deviate from the intended trajectory and strike the glass envelope, resulting in equivalent photon emission. The glass acts as a scintillator, releasing photons via luminescence (see Chapter 9 for details). To suppress stray electron interactions, the PMTs used in this work feature a conductive coating around the glass stem ("HA-coating" for Hamamatsu PMTs). This coating is held at the cathode potential to electrostatically repel stray electrons away from the glass and back into the internal volume.

- **Cosmic Rays:** Atmospheric muons passing through the PMT glass generate Cherenkov light, resulting in large-amplitude background events.

4.4.2 Signal-induced (correlated Background)

Correlated noise manifests as spurious pulses that supersede or appear in conjunction with an ordinary pulse and are categorized by their arrival time relative to the main pulse:

- **Prepulses** arrive $\mathcal{O}(10\,\text{ns})$ before the expected main pulse time. They occur when an incident photon passes through the photocathode without interacting and strikes the first dynode directly, generating a photoelectron that enters the multiplication system at the second stage. Since the photon's travel time is shorter than the electron transit time from the cathode to the dynode, the signal arrives early. Because prepulses lack the first amplification stage, they are significantly under-amplified. A prepulse (sometimes called prepulsing or early pulse) replaces the main pulse.
- **Late pulses** arrive $\mathcal{O}(10\,\text{ns})$ to $\mathcal{O}(100\,\text{ns})$ after the main pulse. When a primary electron strikes a dynode, it may undergo elastic or inelastic backscattering. The backscattered primary electron travels against the electric field (e.g., towards the photocathode), decelerates, reverses direction, and strikes the dynode a second time. The delay time depends on the backscattered electron flight path and varies with the field strength (scaling with $1/\sqrt{V}$). Two regimes are distinguished [146]:
  1. Inelastic Scattering: The primary electron loses energy and scatters back. It returns to the dynode to generate a signal, effectively replacing the main pulse with a delayed, slightly under-amplified one.
  2. Elastic Scattering **(Afterpulsing type I)**: The primary electron creates secondary electrons upon the first impact (generating a main pulse) but also reflects elastically. The backscattered primary electron returns later to generate an additional signal. Thus, the noise pulse follows a main pulse.

  Both types can also occur at later dynode stages, forming a continuum in the delayed arrival times, but they become increasingly under-amplified.
- **Afterpulsing (type II)** occurs $\mathcal{O}(100\,\text{ns})$ to $\mathcal{O}(10\,\mu\text{s})$ after the main pulse. Electrons in the multiplication chain can ionize residual gas molecules (typically H, He, CO_2, Cs, Ar) remaining in the imperfect vacuum. The resulting positive ion is accelerated backwards towards the photocathode. Upon impact, the ion liberates multiple secondary electrons, creating a large amplitude MPE pulse. This mechanism generates a distinct second pulse that follows the main pulse, where the delay is determined by the specific drift path, the supply voltage and the mass of the ion species. The probability of afterpulsing is generally highest in the region between the photocathode and the first dynode due to the large volume and unimpeded ion trajectories.

Part II

Developing the optical Sensor for the next Generation Neutrino Telescope at the South Pole

5 Towards the IceCube-Gen2 optical Module

As introduced in Section 3.7, the IceCube-Gen2 baseline design requires the production and deployment of 9600 new optical modules (OMs) across 120 strings to instrument a volume of 8 km^3. The scale of this expansion imposes stringent requirements on the OM design. With an increased inter-string spacing of 240 m, each module is expected to provide at least four times the photon detection capabilities of current DOMs. At the same time, deployment constraints play an important role: drilling is a major cost driver and logistical as well as infrastructural limitations at the South Pole restrict the maximum outer pressure vessel diameter to 12.5 inches, which is even narrower than the 13-inch diameter of the original IceCube DOM. The central challenge is therefore to maximize the photosensitive area and achieve high photon detection efficiency within a compact glass pressure vessel while ensuring scalability to mass production.

Previous optical module designs informed the development process. For instance, modules with multiple smaller PMTs, such as the KM3NeT module [122] and the multi-PMT digital optical module (mDOM) [147], offer uniform angular coverage. However, with an outer diameter of approximately 14 inches, the mDOM exceeds the IceCube-Gen2 borehole diameter constraints. In contrast, the D-Egg [118] introduces an elongated pressure vessel that increases internal volume while maintaining a narrow profile but does not inherently provide the required photon detection efficiency for IceCube-Gen2. Nonetheless, the uniform coverage, intrinsic directional information and improved saturation characteristics of multi-PMT modules remain desirable. The Gen2-OM concept integrates key features from both the mDOM and D-Egg architectures. It merges the uniform coverage and high photocathode area of a multi-PMT layout with the increased internal volume enabled by an elongated pressure vessel. Furthermore, the novel concept of *gel pads* eliminates the need for mDOM-style reflectors (which necessitated matching neighboring PMT potentials) and avoids the use of complex 3D-printed holding structures that were associated with difficult-to-resolve bubble formation and delamination in the optical gel during environmental testing.

Several names have been used for this module design during its development, including LOM, Gen2-DC (design candidate) and Gen2-DOM. The designations LOM16 and LOM18 specifically refer to the two prototype versions deployed as part of IceCube Upgrade. Throughout this thesis, the term **LOM** will be used to refer to these specific IceCube Upgrade module geometries. The term **Gen2-OM** will be reserved for the future design iteration intended for large-scale production for IceCube-Gen2, which will incorporate lessons learned from LOM production.

Concept Introduction of LOMs

This overarching module architecture represents a novel approach to OM design, developed through close collaboration between institutions in Chiba (Japan), Madison (USA)

and Münster (Germany). A core tenet of the design philosophy is ensuring production scalability by securing multiple vendors for key components and developing fallback solutions. This strategy applies to the PMTs (Hamamatsu and NNVT), the pressure vessel (Nautilus Marine Service and Okamoto Glass) and even the overall module layout (LOM16 and LOM18) [148].

Driven by the narrower borehole limit, the module concept resulted in two primary design candidates currently deployed as prototypes in IceCube Upgrade: the long optical module variants with 16 PMTs (LOM16) and 18 PMTs (LOM18). These designs serve as in situ testbeds, with five and four modules deployed respectively. The LOM18 utilizes a slightly longer pressure vessel to maximize the photocathode area, while the LOM16 features a more compact and mechanically simpler design. Optimization studies combining CAD-based component placement with Geant4 performance simulations (see Chapter 8) guided the finalization of the layout. These studies concluded that newly developed compact 4-inch, short-neck PMTs (see Chapter 6) enable a dense, back-to-back layout within the pressure vessel, maximizing the internal photosensitive area.

Figure 5.1 shows the schematic views of the LOM16 and LOM18. The LOM16 features a module diameter of 312 mm and a height of 444 mm, while the more elongated LOM18 measures 318 mm in diameter and 540 mm in height. In both designs, each hemisphere contains a central support structure made of simple sheet metal. This structure holds the central electronics boards, flasher LED units and the individual PMT subassemblies (each comprising a bare PMT, a pre-cast optical gel pad and a custom readout base). These subassemblies are classified by their position within the module as either *polar*, pointing towards the vertical ends of the pressure vessel, or *equatorial*, near the sealing



Figure 5.1: Exploded view CAD renderings of the LOM16 (left) and LOM18 (right).

surfaces. A key difference is that the LOM18 support structure is anchored by central shafts on the top and bottom PMT assemblies, whereas the LOM16 omits this central shaft, simplifying the internal mechanics. The harness clasps the optical modules and the entire assembly is anchored to the main deployment cable via four wire ropes. The assembled LOM16 and LOM18 modules have masses of 21.3 kg and 30.8 kg respectively, including the wire ropes, with bare pressure vessel weights of 10.8 kg and 16.9 kg. Their net buoyancies in water are 5.8 kg and 5.5 kg [127].

A key innovation of this design is the optical coupling of the PMTs to the pressure housing. Instead of a complex mechanical holder, the modules use custom-shaped optical gel pads pre-cast directly onto the PMT face prior to integration. These gel pads are made of a UV-transparent silicone elastomer that exhibits negligible attenuation compared with the pressure vessel. Simulation studies demonstrate that utilizing these discrete gel pads yields optical performance surpassing that of a module completely filled with silicone gel [10 (p. 114)], representing a cost-effective and lightweight approach, enabling a simplified holding structure compared with the mDOM. To prevent detachment and ensure stable optical coupling to the pressure vessel housing, the *partially pre-cast gel pads* currently utilized in the deployed LOMs feature a wall-shaped protrusion (rim) along their front edge. During assembly, this rim is pressed against the inner surface of the pressure vessel, forming a sealed internal cavity. This cavity must subsequently be *backfilled* with liquid gel via injection tubes, necessitating complex and time-consuming assembly as well as degassing procedures.

Beyond these mechanical coupling requirements, the specific geometry of the gel pads is crucial for the overall optical performance. Regardless of the specific casting procedure, the conical shape of these gel pads is universally designed to enhance the module's photodetection efficiency through total internal reflection at the gel-glass interface, effectively collecting photons that would otherwise miss the photocathode, while still detecting photons incident from the side. The shape is optimized to maximize the active surface area on the pressure vessel's inner wall, providing more uniform photon collection, which is advantageous in a scattering medium like deep glacial ice.

The electronics system is designed to be compact, low-power and capable of handling signals from single photoelectrons up to the very bright events expected from high-energy neutrinos. Each PMT is equipped with a custom base, the wuBase [149], which integrates a Cockcroft-Walton high-voltage generator with data acquisition features. It reads out both the anode and the eighth dynode signals using a dual-channel 60 MSPS ADC. Within each hemisphere, a Fanout board acts as a hub, handling command distribution and data collection from the PMT subassemblies. Communication with the surface is managed by Mini-Mainboards, which connect to the main data cable via the penetrator cable assembly (PCA). The integration process is simplified by the fact that only three ribbon cables cross between the two hemispheres [149].

Based on physics simulations of cascades at different distances and orientations relative to the PMTs, Figure 5.2 shows the dynamic range capabilities of the readout electronics. Even for PMTs located within the Cherenkov cone and facing the cascade, the wuBase waveforms typically remain in the linear region up to 100 TeV at a distance of 25 m, or 1 PeV at 40 m. For muons, which produce cascades stochastically along their path, simulations indicate that a 100 TeV stochastic energy deposition can commonly occur within a 50 m segment for a 5 PeV muon, and the signal would similarly remain in the linear region as close as 25 m [127].

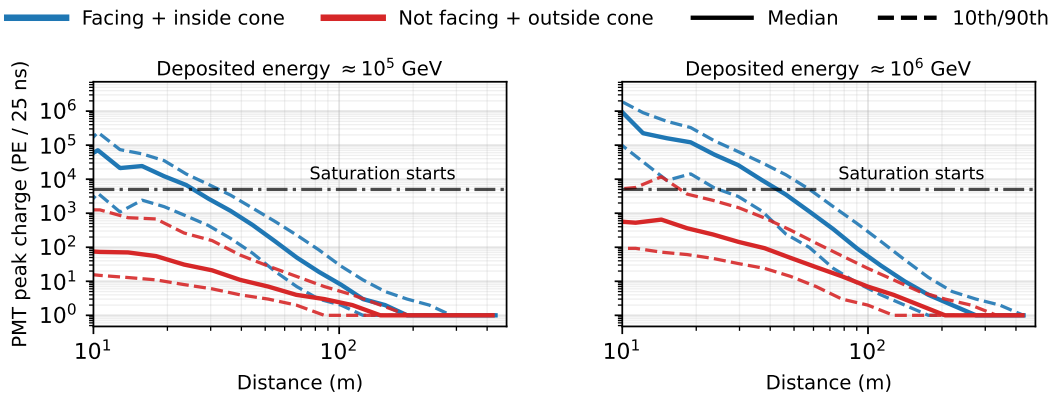


Figure 5.2: Simulation of the expected PMT peak charge per 25 ns time bin as a function of distance from cascades following neutrino interactions. The left panel corresponds to a deposited (visible) energy of 100 TeV and the right panel to 1 PeV. The blue bands represent PMTs facing the emission and located within the Cherenkov cone, while the red bands represent PMTs outside the Cherenkov cone or facing away. Saturation of the electronics for anode readout typically begins around 10^2 photoelectrons, while dynode eight readout remains linear up to $\sim$ 5500 photoelectrons per 25 ns per PMT as indicated by the black line. Taken from [127].

For in situ calibration, eight flasher LED modules are installed in the gaps of the equatorial PMT layer. These flashers emit light at approximately 405 nm with a tunable pulse width typically ranging from 1 to 10 ns and are daisy-chained and connected to the Fanout board in the upper hemisphere. However, as concluded during IceCube Upgrade's calibration campaign (see Section 3.6), internal flashers alone are insufficient. To reduce uncertainties in the modeling of the surrounding ice, dedicated calibration devices such as the LOM-Logger [150], which were successfully used during IceCube Upgrade construction, can be used to study the localized optical properties of the deep ice environment in even greater detail. This measurement is achieved by co-deploying a dedicated laser light source 2 m below a receiving LOM on the same string, allowing the combined system to continuously record the scattered light and map the ice stratigraphy as it is lowered into the borehole.

The LOM16 and LOM18 successfully serve as the Gen2-OM precursors, validating both the multi-PMT architecture and the novel gel pad coupling. Before evaluating the mechanical assembly methods required to eventually scale this design to the 9600 optical modules needed for IceCube-Gen2, the performance of the core sensors is presented in Chapter 6 with a detailed characterization of the 4-inch photomultiplier tubes. Following this sensor validation, Chapter 7 details the mechanical prototyping efforts that guided the LOM construction and explores integration methods to optimize this module architecture for future mass production.

6 Photomultiplier Tube Investigation

Photomultiplier tubes are the central element of neutrino telescopes, making a detailed understanding of their operational characteristics essential. Following the commercially available form factors, the optical module design aims to maximize the photosensitive area within the geometric constraints of the pressure vessel (see Chapter 5). This requirement drove the development of novel 4-inch PMTs with a compact envelope, limited to a maximum length of 106 mm, a width of 104 mm and a neck diameter of 55 mm. Achieving this level of compactness presents a significant design challenge, requiring a careful balance between mechanical constraints and key performance metrics, specifically detection efficiency and time resolution at the photocathode periphery. The development therefore aimed to retain performance comparable with standard PMTs while minimizing the physical envelope. To address the specifications of the IceCube-Gen2 collaboration, a partnership was established with Hamamatsu Photonics K.K. (Hamamatsu) and North Night Vision Technology (NNVT). Following an iterative prototyping phase initiated in 2020, both vendors delivered final production models: the Hamamatsu R16293 and the NNVT N2041 (see Figure 6.1). While the development cycle involved the comprehensive characterization of four Hamamatsu and five NNVT variants, this chapter focuses on the performance of the final production models. The iterative feedback from earlier measurements directly informed the design of the units evaluated here. Specifically, the following results are based on three final production units per manufacturer: serial numbers BB9780, BB9786 and BB9789 (Hamamatsu, plotted in red) and PO4049, PO4052 and PO4068 (NNVT, plotted in blue). Unless stated otherwise, performance metrics are averaged over a temperature range of −40 °C to −10 °C at nominal voltage (summarized in Section 6.6). Parts of these results have been published in [10, 151] as well as in [152], written by the author, and [127], co-written by the author.

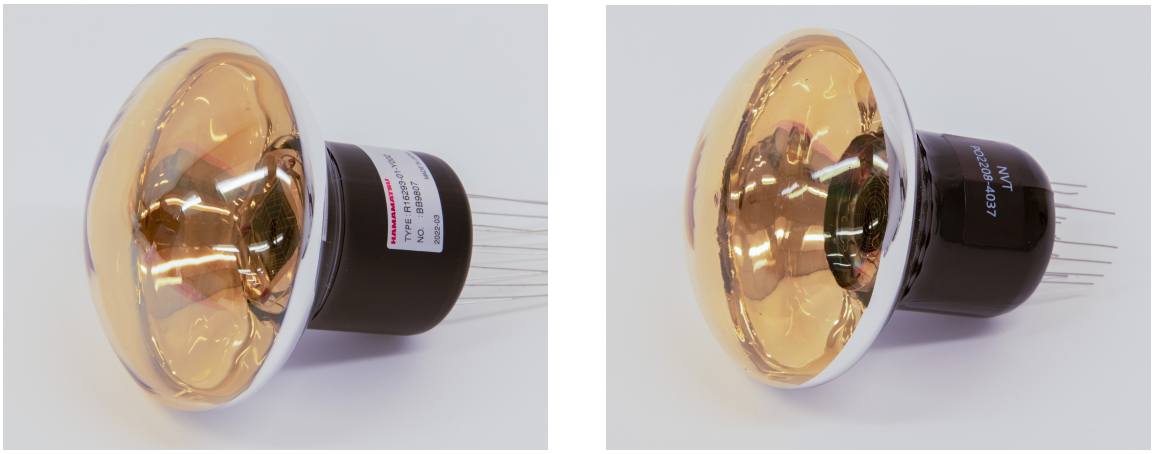


Figure 6.1: Final production versions of the 4-inch Hamamatsu R16293 (left) and NNVT N2041 (right) PMTs. Courtesy of the IceCube collaboration.

6.1 Data Acquisition

As introduced in Section 4.1, PMTs are read out in either current mode or pulse mode. Most measurements of this characterization utilize pulse mode to investigate correlations between amplitude, timing and charge for SPE levels of illumination. Data acquisition in this work distinguishes between two distinct readout modes:

- **MP-readout (multi pulse):** For every pulse contained in a waveform that exceeds a predefined threshold (e.g., 2 mV), the maximum amplitude (relative to baseline), charge (integrated over a region of interest around the maximum) and time (relative to the trigger) are extracted. This mode is typically applied to long waveforms $\mathcal{O}(1\,\text{ms})$ where multiple pulses are stored per trigger.
- **SP-readout (single pulse):** Only the pulse with the global maximum amplitude within the waveform is recorded. In addition to the aforementioned amplitude, charge and time, pulse shape parameters, specifically rise time (RT), fall time (FT) and pulse width (FWHM), are extracted. This mode is generally applied to shorter waveforms $\mathcal{O}(100\,\text{ns})$ and used in gain calibration.

Figure 6.2 illustrates these readout modes based on experimental data, including electronic baseline noise and secondary base ringing. To prevent double counting of single physical events (which can occur due to base ringing), an artificial deadtime of 10 ns is applied during acquisition. While active Cockcroft-Walton bases are employed in LOMs and intended for the Gen2-OM, the measurements presented here utilize passive bases to characterize the intrinsic PMT response.

Illumination and trigger levels are quantified in units of photoelectrons (PE). Precise knowledge of the light source is essential for characterizing the SPE response. Photoelectron emission follows a Poisson distribution $P(n, \mu) = \mu^n e^{-\mu}/n!$. At an average light intensity of $\mu = 0.7\,\text{PE}$, approximately 50 % of light flashes produce no detectable signal (the pedestal), 35 % result in a single photoelectron (SPE) and the remaining 15 % contain multi-photoelectron (MPE) contributions.

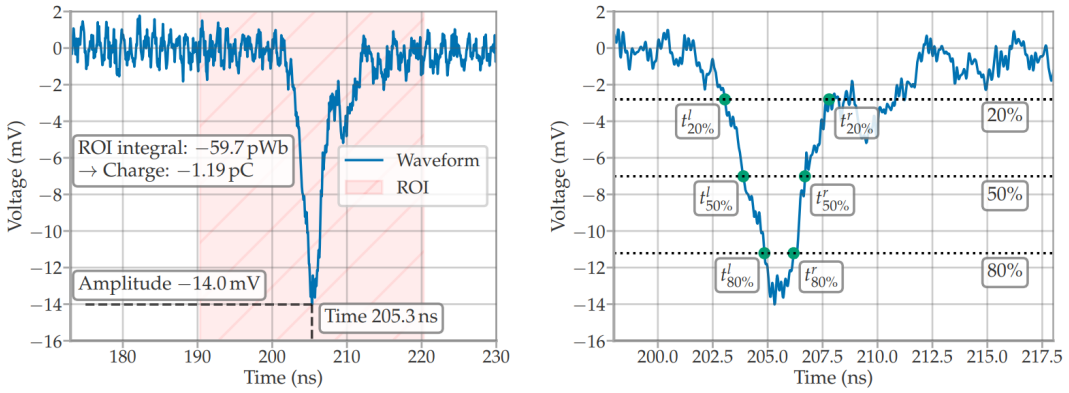


Figure 6.2: **(left)** MP-readout extracts the amplitude, charge (ROI integral) and time exceeding a predefined amplitude threshold for each pulse within a waveform. For illustration, only a single pulse from the full waveform is shown. **(right)** SP-readout additionally extracts FT, RT and pulse width (defined by the 20th, 50th and 80th percentiles). Taken from [73].

6.2 Temperature-dependent Characterization

The operational temperatures of IceCube-Gen2 range from −10 °C to −40 °C depending on the instrumentation depth, with temperatures decreasing closer to the surface. Characterizing PMT performance across this range is critical, particularly for detection efficiency as well as timing resolution, which intrinsically limit the detector's capabilities. To systematically study these dependencies, PMTs were evaluated in a climate chamber setup, as illustrated in Figure 6.3. PMTs are operated using negative high-voltage[1] and connected to an oscilloscope[2]. A pulse generator[3] provides the trigger signal while driving the LED[4] which is routed into the climate chamber[5] via an optical fiber[6] and diffused to illuminate the entire photocathode.

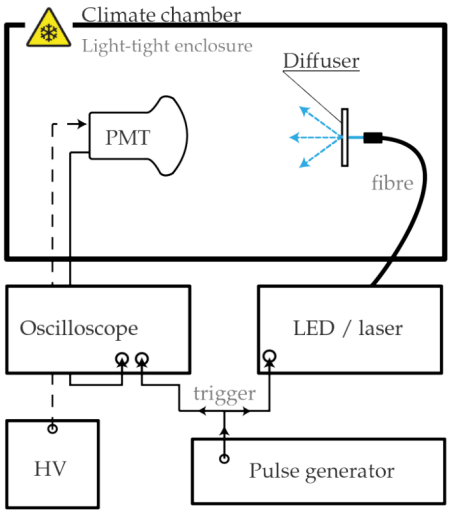


Figure 6.3: Setup for temperature-dependent PMT characterization. Taken from [73].

6.2.1 Gain Calibration and Pulse Shape Parameters

PMT characteristics depend on the dynode chain amplification, necessitating calibration to a reference gain level. For these PMTs, this corresponds to a nominal gain of $5 \cdot 10^6$ for which the nominal voltage is to be applied.

To determine the required nominal voltage, charge spectra are recorded in SP-readout mode using 250 000 waveforms (at 250 ns sampling length). A light intensity of $\mu \approx 0.7$ PE is selected to optimize the statistical contribution of the SPE peak relative to the pedestal and MPE events. This intensity ensures that, while the SPE peak remains dominant, the 2, 3 and 4 PE contributions are sufficiently present to allow for a multi-peak fit of the charge distribution. The supply voltage is varied from −875 V to −1275 V (Hamamatsu) and −650 V to −950 V (NNVT) in 25 V increments across the temperature range of −50 °C to 20 °C. The resulting charge spectra are fitted with Equation (4.6)

[1] supplied by an ISEG ECH244-type crate with EHS F4 40n module.
[2] PicoScope 6404C
[3] RIGOL DG1032Z
[4] 385 nm LED mrongen – Custom Picosecond light sources.
[5] CTS C-70/350
[6] Thorlabs FG105LCA Multimode Fiber

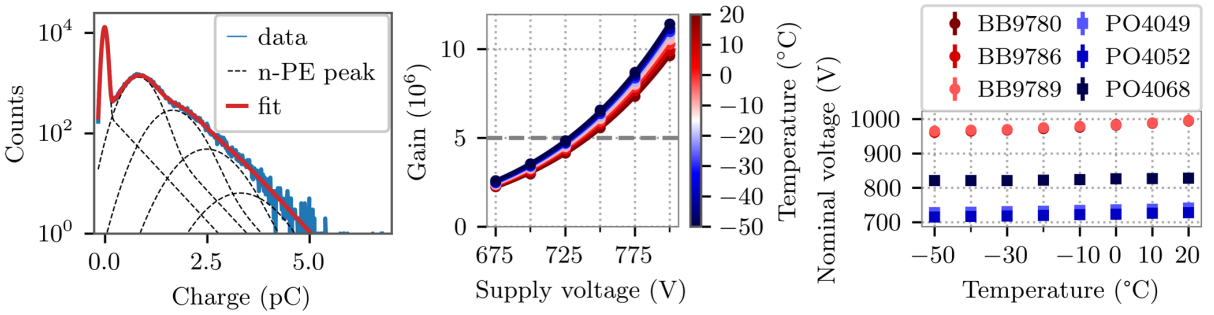


Figure 6.4: (left) Charge spectrum fit for BB9780 at −20 °C and 975 V ($\chi^2_{\text{red}} \approx 1.004$). **(center)** Gain dependence on supply voltage and temperature for PO4049. The horizontal dashed line indicates the target nominal gain. **(right)** Nominal voltage required by each PMT to achieve a gain of $5 \cdot 10^6$ as a function of temperature.

to extract the gain for each supply voltage at a given temperature, as shown in the left panel of Figure 6.4. The relationship between gain G and voltage V follows the power law $G = a \cdot V^b$ (Equation (4.5)). For improved numerical precision, linear regression is performed in double-logarithmic space ($\log G = b \log V + \log a$). The results are shown in the center panel of Figure 6.4, illustrating the fits used to identify the nominal voltage for each temperature.

As shown in the right panel of Figure 6.4, a temperature decrease of $\Delta T = 70\,°\text{C}$ results in an average gain increase of $(29.1 \pm 0.4)\,\%$ for Hamamatsu and $(13.6 \pm 0.3)\,\%$ for NNVT PMTs, attributed to the temperature-dependent secondary emission coefficient of the dynode materials [153]. Hamamatsu PMTs exhibit a smaller inter-sample spread in nominal voltage compared with NNVT (standard deviation 4.6 V vs 45.7 V), although NNVT models generally require lower absolute voltages to achieve the target gain. While the sample size within this study is limited, this observation is supported by manufacturer data for the full production batches, which report average nominal voltages of (1056.4 ± 26.8) V for Hamamatsu and (935.2 ± 103.5) V for NNVT (Appendix C). This trend is also consistent across earlier design iterations. Compliant with IceCube-Gen2 requirements, all tested units operate well below the 1500 V maximum limit. The calibrated nominal voltages are applied in all subsequent measurements in the following sections. As discussed in [73 (p. 54)], at the nominal gain of $5 \cdot 10^6$ a fitting bias of $(-5.39 \pm 0.08)\,\%$ is observed for mDOM PMTs. Since the gain is not used as a reference to calculate incident photons or infer detection efficiency, this effect is considered negligible.

Charge Resolution and Peak-to-Valley Ratio

The parameter space can further be explored for charge resolution (spread over the average charge of the 1-PE peak: σ_1/Q_1) and peak-to-valley ratio (PVR: ratio of the SPE peak to the local minimum towards the pedestal). The charge resolution quantifies the statistical fluctuation of secondary electron multiplication, where lower values imply a lower variance of the charge (at suitable gains). The PVR indicates the ability to distinguish signals from noise. Both are shown in Figure 6.5 plotted against PMT gain. As no significant temperature dependence beyond the aforementioned variation in gain could be identified, all temperatures are plotted together for compactness. At nominal gain, the mean charge resolutions are $(37.4 \pm 1.2)\,\%$ (NNVT) and $(39.5 \pm 1.1)\,\%$ (Hamamatsu), comparable with the $\sim 40\,\%$ observed in mDOM PMTs [73 (p. 58)]. All samples exceed the PVR

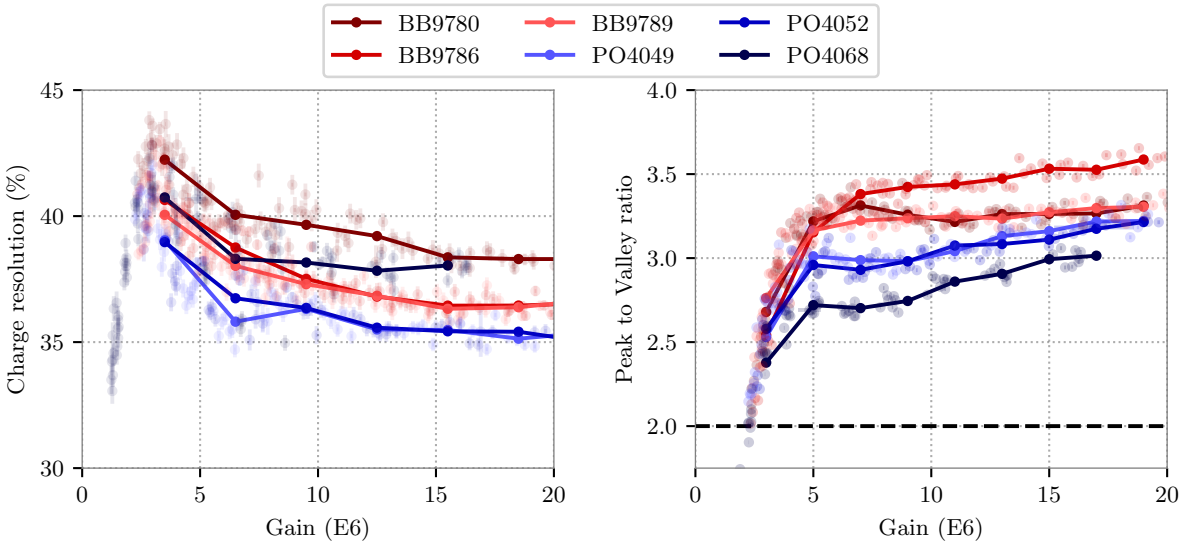


Figure 6.5: PMT performance parameters as a function of gain. Values from all supply voltages and temperatures (transparent datapoints) are binned (solid datapoints), with lines added as visual guide. **(left)** Charge resolution. **(right)** Peak-to-valley ratio (dashed line indicates the requirement at a nominal gain of 5e6).

requirement of > 2, with Hamamatsu (3.22 ± 0.07) systematically offering better noise separation than NNVT (2.85 ± 0.11), likely attributed to the higher nominal voltages.

Pulse Shape

Rise time (RT), fall time (FT) and pulse width (FWHM) are directly derived during SP-readout, where all values for a given temperature and voltage measurement are averaged. The left panel of Figure 6.6 illustrates the voltage dependence of the rise time. Higher supply voltages result in shorter rise/fall times and narrower pulses. This is physically expected, as the increased electric field accelerates photoelectrons more rapidly towards subsequent dynodes ($a \propto V$), reducing the transit time according to $t \propto 1/\sqrt{V}$. Stronger electric fields also reduce the temporal spread of secondary electrons, however, at the cost

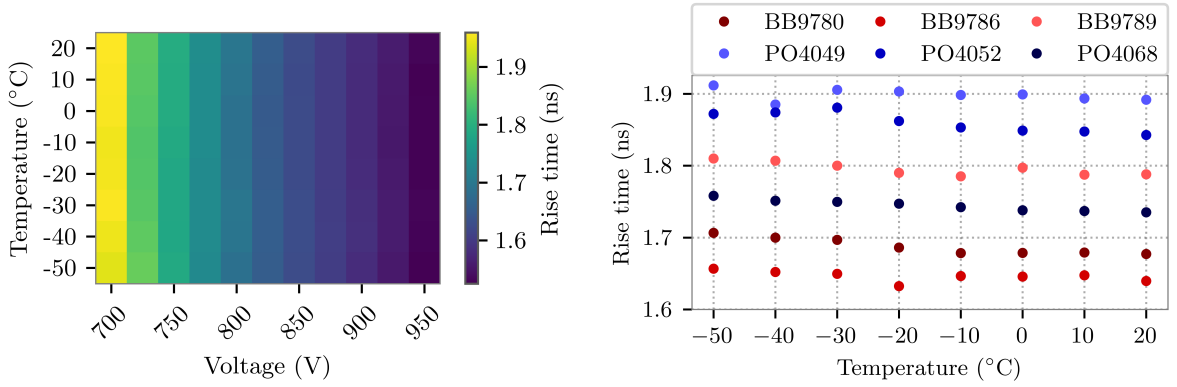


Figure 6.6: **(left)** Rise time dependence on voltage and temperature for PO4052. **(right)** Rise time interpolated to the nominal gain as a function of temperature. Error bars too small to visualize.

of slightly worsened charge resolution, as secondary electrons spread out more with higher initial energies.

Averaging over the operational range of −40 °C to −10 °C, the mean 20 % to 80 % rise times are (1.70 ± 0.06) ns (Hamamatsu) and (1.83 ± 0.07) ns (NNVT). The fall times are (1.63 ± 0.04) ns and (1.66 ± 0.04) ns, respectively, with average pulse widths (FWHM) of (3.29 ± 0.08) ns and (3.30 ± 0.05) ns. Beyond the aforementioned gain-induced shifts, no significant temperature dependence can be identified (see Figure 6.6 right).

6.2.2 Timing Resolution and Pulse Types Probabilities

The transit time spread (TTS) of a PMT directly affects the angular resolution of the entire detector, as it is the underlying timing resolution. To isolate single-photon timing characteristics and correlated noise, the setup from before is repeated at nominal voltage in MP-readout mode using 10 µs waveforms at a low light intensity ($\mu < 0.1$ PE). This occupancy (< 0.5 % MPE) ensures reliable SPE timing characterization of the TTS and accurate pulsing classification, as higher photoelectron multiplicities disproportionately increase correlated noise such as afterpulsing. A dataset of approximately 8×10^6 waveforms is acquired, yielding roughly 6×10^5 waveforms with at least one pulse to ensure sufficient statistics.

To ensure accurate parameterization of the TTS, genuine SPE pulses must be isolated from MPE contamination and electronic noise. While MPE pulses are primarily suppressed by the low light intensity, baseline noise is removed through a two-stage filtering process. As illustrated in the left panel of Figure 6.7, the correlation between pulse amplitude and charge provides a reliable separation of true SPE pulses. Data acquisition for this specific measurement initially relies on a hardware trigger of 3 mV (corresponding to ∼ 0.49 PE). Subsequently, a strict post-acquisition charge cut of 0.2 PE is applied to eliminate any remaining low-charge background fluctuations.

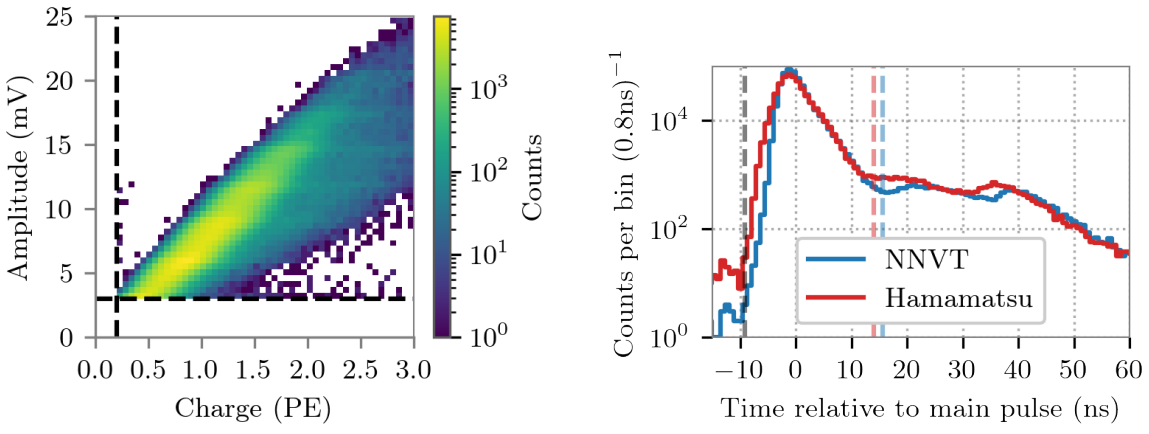


Figure 6.7: (left) Correlation between charge and amplitude for BB9780 at 0 °C. Dashed lines indicate the hardware amplitude trigger threshold and post-measurement charge cut. **(right)** Pulse arrival time distribution relative to the main peak for BB9789 and PO4068 as representative examples at −20 °C. Vertical dashed lines indicate the transition regions used for pulse classification.

With genuine SPE pulses isolated, the resulting averaged pulse arrival time distributions (Figure 6.7, right) reveal a non-Gaussian shape characteristic of PMTs with elliptical photocathodes. The pronounced tail (up to 17 ns) arises from transit time differences across the photocathode surface, heightened by the asymmetric curvature of the first dynode under full frontal diffuse illumination. Hamamatsu PMTs exhibit a generally broader main peak, whereas NNVT PMTs feature a more extended tail. Although typical, this tail is more pronounced in these PMTs than in comparable models due to the envelope shape and will be revisited in Section 6.5. Because of these non-Gaussian tails, defining a single metric for the timing resolution requires consideration, where three calculation methods of the TTS are commonly used:

$\mathrm{TTS}_{\mathrm{G,peak}}$: A Gaussian is fitted around the maximum of the peak (−2.5 ns to 4.0 ns) to extract the core timing resolution. Manufacturers typically quote the FWHM of this core peak, as their standard validation procedures often restrict illumination to the central photocathode region where extended tails are absent. However, applying this fit to data obtained under full frontal illumination underestimates the width of the distribution by neglecting the tails.

$\mathrm{TTS}_{\mathrm{G,full}}$: A Gaussian is fitted to the full pulse window (−8 ns to 17 ns). This typically yields larger values than applying the fit to just the peak, as the inclusion of the tails forces the fit to increase the core width.

$\mathrm{TTS}_{\mathrm{num}}$: The numerical standard deviation of the unbinned data over the full window (−8 ns to 17 ns). Given the pronounced non-Gaussian tails observed in these 4-inch PMTs, this metric provides the most accurate representation of the effective timing resolution and is selected as the primary figure of merit in this work.

The temperature evolution of these metrics is shown in Figure 6.8 (left). No systematic temperature dependence is observed, with most samples stable within 1.5 % across the $\Delta T = 70\,^{\circ}\mathrm{C}$. Exceptions include BB9780 (+3.5 % increase at low temperatures) and PO4052 (8 % decrease). Averaging over the operational range, NNVT PMTs exhibit a $(18.6 \pm 0.1)\,\%$ better timing resolution: $\mathrm{TTS}_{\mathrm{num}} = (2.80 \pm 0.06)$ ns (NNVT) compared with (3.36 ± 0.08) ns (Hamamatsu). The inter-sample variance is notably higher for the Hamamatsu units. Both PMT models readily fulfill the requirement of $\mathrm{TTS}_{\mathrm{G,peak}} < 3.4$ ns.

In addition to timing resolution, the respective contributions of correlated background are calculated from this measurement. As introduced in Section 4.4.2, five pulse types are distinguished based on their arrival time relative to the main peak: main peak (-8 ns to +17 ns), prepulses (-60 ns to -8 ns), late pulses (+17 ns to +140 ns), afterpulsing type I (same window as late pulses), afterpulsing type II (+140 ns to 10 µs). The probabilities for these pulse types are summarized in Figure 6.8 (right). Both manufacturers comply with the requirements, with NNVT PMTs consistently showing lower values. The probabilities of triggered prepulses are negligible ($\leq 0.05\,\%$; Req. $< 1\,\%$). The combined rate of late pulses and afterpulsing type I is $(4.83 \pm 0.04)\,\%$ for Hamamatsu and $(3.65 \pm 0.05)\,\%$ for NNVT, satisfying the $< 5\,\%$ requirement. Afterpulsing type II rates are similarly low ($\sim 4.5\,\%$), well below the $< 15\,\%$ specification. With the exception of a slight temperature dependence in the Hamamatsu afterpulsing type I, the correlated background remains constant across the investigated temperature range.

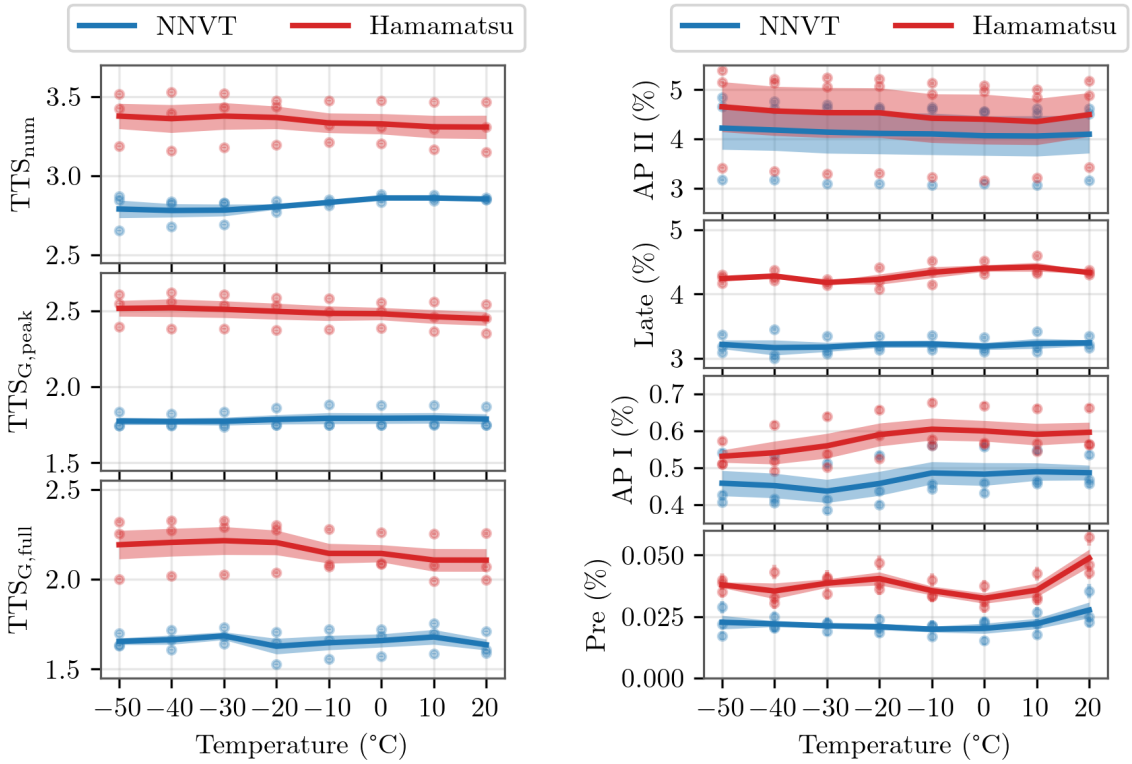


Figure 6.8: Temperature dependence of TTS and pulse type probabilities. Solid lines represent the mean; bands indicate the standard error across samples (transparent). **(left)** TTS using the standard deviation of the three definitions. **(right)** Probabilities of correlated noise types: prepulses, late pulsing and afterpulsing (type I & II).

6.2.3 Dark Rates

Dark rate measurements characterize the intrinsic noise level of the PMTs in the absence of illumination. Data is acquired using a random trigger to record 1 ms waveforms in MP-readout mode. Signal processing followed the methodology described in Section 6.2.2, utilizing a 3 mV ($\sim$ 0.45 PE) hardware trigger threshold. An additional post-acquisition charge cut of 0.3 PE was applied to suppress baseline fluctuations. As mentioned before, an artificial deadtime of 10 ns is applied, where pulses occurring within this interval are merged and treated as a single entry. The total dark rate is then calculated by averaging the number of unique pulses per waveform over the full sample set, normalizing by the waveform durations.

Noise sources are classified by analyzing the time interval Δt between consecutive pulses. As shown in Figure 6.9 (left), the distribution of $\log_{10}(\Delta t)$ allows for the separation of noise components. Uncorrelated noise, predominantly stemming from thermionic emission (see Equation (4.7)), follows a Poisson process $P(t) = \mu e^{-\mu t}$ that appears as an exponential tail in the linear time domain. In logarithmic space, this transforms to

$$f(x) = A \cdot \mu \cdot \ln(10) \cdot 10^x \cdot \exp\left(-\mu \cdot 10^x\right), \tag{6.1}$$

where $x = \log_{10}(\Delta t)$, μ is the uncorrelated rate and A is a normalization constant [73 (pp. 108–109)]. Fitting this function to the right peak ($\sim \log_{10}(\Delta t) = -1.5$) of the

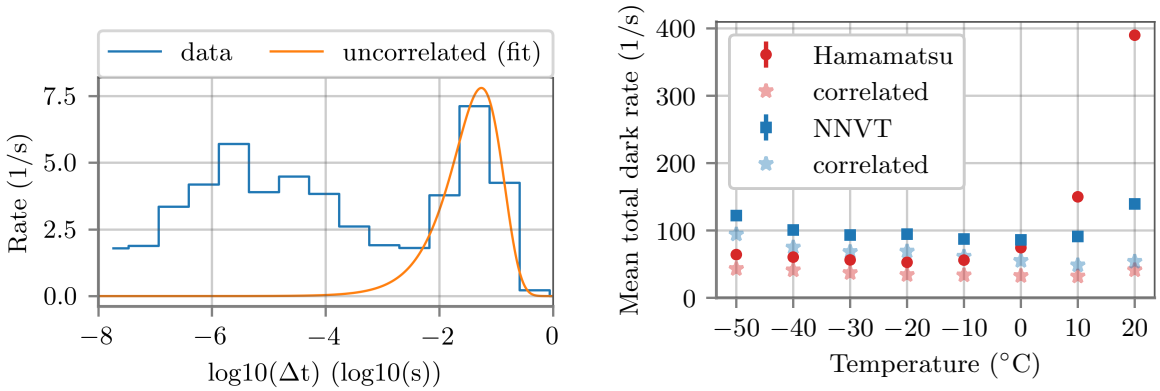


Figure 6.9: (left) Distribution of time differences between consecutive pulses ($\log_{10} \Delta t$) for BB9780 at -30°C. The fit (orange) isolates the uncorrelated Poissonian component whilst the excess at low Δt indicates correlated scintillation noise. **(right)** Mean total dark rates as a function of temperature for both vendors, including the contribution from correlated noise (standard error too small to visualize).

distribution isolates the random rate μ. The excess at short time intervals represents correlated noise, predominantly caused by scintillation induced by radioactive decays within the glass envelope (see Chapter 9) as well as PMT correlated background (Section 6.2.2) and is obtained by subtracting μ from the total rate.

The temperature dependence of the mean total dark rate is shown in Figure 6.9 (right). At room temperature, rates are driven by thermionic emission, which is higher for Hamamatsu PMTs. However, in the relevant operational range of $-40\,°\mathrm{C}$ to $-10\,°\mathrm{C}$, thermionic emission is exponentially suppressed and the background rate becomes dominated by the temperature-dependent scintillation component.

In this low-temperature regime, Hamamatsu PMTs exhibit lower dark rates of $(56.3 \pm 2.7)\,\mathrm{s}^{-1}$ compared with $(94.0 \pm 4.8)\,\mathrm{s}^{-1}$ for NNVT units. This difference is attributed to lower trace amounts of radioactive materials present in the bulb glass of the Hamamatsu samples, as evidenced by the reduced correlated noise component. While the IceCube-Gen2 requirement is soft ("as low as possible"), these values can be contextualized for an optical module comprising 16 PMTs: The total intrinsic dark rate extrapolates to approximately $900\,\mathrm{s}^{-1}$ using Hamamatsu PMTs and $1500\,\mathrm{s}^{-1}$ using NNVT PMTs. These values align well with the overall noise rates of the main IceCube Upgrade modules: The intrinsic D-Egg PMT rates exhibit $\sim 2 \cdot 400\,\mathrm{s}^{-1} = 800\,\mathrm{s}^{-1}$ [154] prior to module integration and "low noise" mDOM PMTs show dark rates of $\sim 24 \cdot 35\,\mathrm{s}^{-1} = 840\,\mathrm{s}^{-1}$ [73 (p. 107)]. Although NNVT yields a higher total rate in this sample, it remains within the acceptable range for multi-PMT optical modules. However, experience from mDOM mass production demonstrates that glass radiopurity is subject to changes in production, where dark rates have increased to $\sim 24 \cdot 320\,\mathrm{s}^{-1} = 7680\,\mathrm{s}^{-1}$ [155 (p. 128)]. A similar elevated dark rate has been observed in earlier prototype models from both manufacturers. Crucially, the PMT rates discussed here represent the bare PMT contribution, where the total in-module background, including scintillation from the pressure vessel, is discussed in Chapter 9.

6.3 Quantum Efficiency

The quantum efficiency (QE) quantifies the probability that an incident photon generates a photoelectron. The QE directly influences the energy resolution of the detector, as the statistical fluctuations in the number of detected photoelectrons are the dominant contribution to the energy uncertainty.[7] Although QE is an intrinsic property of the photocathode material, the effective QE is constrained by the glass envelope's transmissivity and governed by the photon absorption probability, which can be increased through internal reflections.

The setup, illustrated in Figure 6.10 (left), houses the PMT in a light-tight enclosure. Illumination is provided by a xenon arc lamp[8], coupled to a monochromator[9] with a spectral resolution of FWHM $\approx (1.0 \pm 0.6)$ nm. The monochromatic beam passes through an adjustable iris to change the illuminated spot diameter on the photocathode. The PMT is operated with a specialized passive base that shorts the dynode chain and anode to the first dynode, allowing the entire multiplication system to act as a photoelectron collector. A supply voltage[10] of -250 V is applied to the photocathode, ensuring stable collection efficiency (stability is typically reached above -50 V [73 (p. 68)]). As a consequence of the lack of multiplication and in contrast to the pulse-mode characterization in previous sections, the QE is determined via photocurrent measurements using a picoammeter[11]. A calibrated reference photodiode[12] with known QE, mounted on a 3-D motor scanner [156], is moved into the beam path to measure the incident illumination level.

The beam diameter is adjusted to approximately 80 mm to 90 mm to maximize photocathode coverage while preventing light leakage past the PMT envelope (outer glass diameter ≈ 102 mm). Light intensity is regulated to keep photocurrents below 0.1 nA mm^{-2}, preventing saturation effects. The wavelength is swept from 250 nm to 750 nm in 10 nm steps. For each step, the current is measured 50 times for both the PMT (I_{PMT}) and the reference diode (I_{Diode}), as well as their respective dark currents (DC). The QE is calculated as:

$$\mathrm{QE}(\lambda) = \frac{I_{\mathrm{PMT}}(\lambda) - \mathrm{DC}_{\mathrm{PMT}}}{I_{\mathrm{Diode}}(\lambda) - \mathrm{DC}_{\mathrm{Diode}}} \cdot \mathrm{QE}_{\mathrm{Diode}}(\lambda). \tag{6.2}$$

The average QE for both vendors is shown in Figure 6.10 (right). Both models meet the requirement of QE $> 25\,\%$ at 400 nm[13] at room temperature. A bachelor thesis [157] co-supervised by the author reports a maximum QE increase of $< 0.015\,\%$ per °C, consistent with literature values [134, 153]. The apparent gain in QE caused by decreasing the temperature is negligible compared with inter-sample variations.

[7] The total detection efficiency is the product of collection efficiency and quantum efficiency, both showing wavelength dependencies. The collection efficiency is relatively constant above 90 to 99 % for $\lambda > 450$ nm, decreasing slightly towards lower wavelengths (80 % to 90 %), generally attributed to higher initial velocity photoelectrons straying from the intended path [133]. The wavelength dependence of the detection efficiency is, however, governed by the QE.

[8] LOT LSE140 / 160.25.C

[9] LOT-QuantumDesign MSH-300

[10] Supplied by an Iseg NHQ 226L.

[11] Keithley Model 6482

[12] Newport 818-UV

[13] The value of 400 nm is a historical benchmark in IceCube. The spectrum reaching optical modules within IceCube's optical array depends on the propagation distance due to scattering; discussed in Chapter 8.

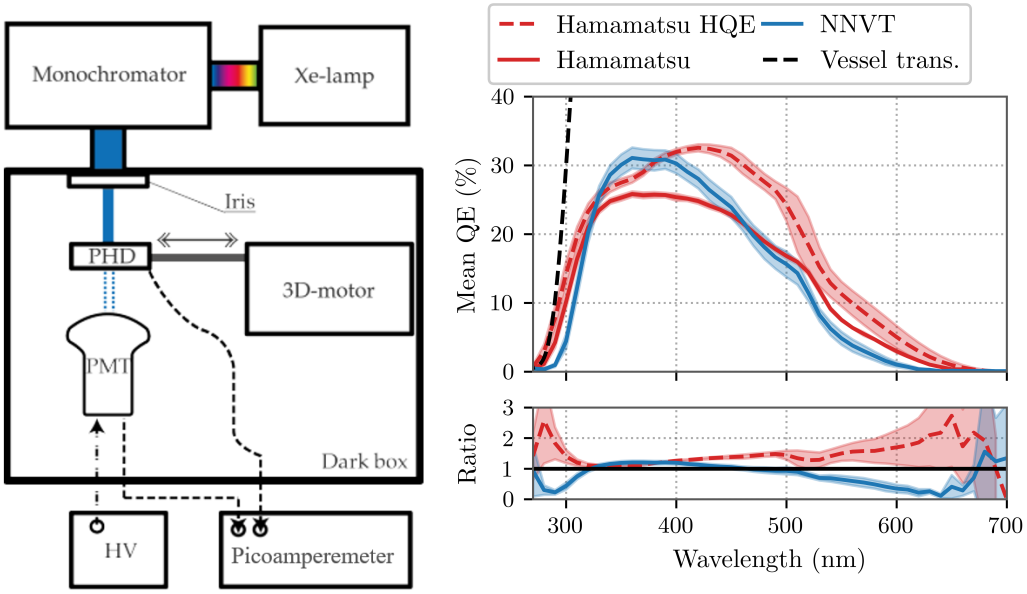


Figure 6.10: (left) Experimental setup for QE measurements. Taken from [73]. **(right)** Average QE measured under full frontal illumination at room temperature. The black line indicates the transmission cutoff of the encompassing optical module pressure vessel. The ratio below is normalized to the standard-QE Hamamatsu average. Error bands indicate standard deviation between the respective samples.

NNVT PMTs exhibit a higher peak QE of $(31.1 \pm 1.5)\,\%$ compared with $(25.8 \pm 0.4)\,\%$ for the standard Hamamatsu R16293. The spectral shape also differs: NNVT PMTs show a sharper response concentrated around the 360 nm peak, whereas Hamamatsu PMTs display a broader spectral sensitivity. This suggests slight differences in the K_2CsSb photocathode composition. Both vendors show a sharp sensitivity drop-off below 300 nm, governed by the UV-absorption of the borosilicate glass envelope. However, Hamamatsu PMTs systematically maintain higher sensitivity further into the UV range across development iterations. This cutoff aligns well with the transmission characteristics of the pressure vessels (indicated by the black line in Figure 6.10, right), ensuring that the PMT sensitivity maximizes the transparency window of the module housing.

Complementary measurements comparing central spot and full frontal illumination reveal that PMTs exhibit a higher QE under full frontal conditions than under centered spot measurements (an absolute increase of 4.5 % at the peak for Hamamatsu PMTs). This effect is attributed to higher reflectivity of the internal components, which enhances the effective efficiency for wide-beam measurements. However, given that the performance requirements are defined for full frontal illumination, the values presented here focus exclusively on the results obtained with the ~ 80 mm beam geometry.

Hamamatsu additionally offers a High-QE (HQE) variant of the R16293 measured with two samples. While this version offers enhanced photon detection efficiency (with the maximum shifted to ~ 430 nm), it entails higher costs. Furthermore, increased QE also corresponds to higher sensitivity to optical background generated within the pressure vessel glass (see Chapter 9). Thus, the trade-off between signal gain and noise background remains under evaluation.

It is critical to contextualize the performance gap observed here. Manufacturer data from the respective production batches indicates that the standard Hamamatsu PMTs

provided for this study were from the lower tail of their QE distribution, while the NNVT samples represented above-average performers. Therefore, the nominal performance gap between vendors is likely smaller than the $\sim 5\,\%$ observed in this sample. However, the general trend of NNVT PMTs exhibiting slightly higher peak QE with a higher cutoff, as well as a higher inter-sample standard deviation, has been consistent across previous prototype iterations and is supported by manufacturer data.

6.4 Performance Metrics for IceTray

Within the IceCube simulation framework, IceTray [158], the PMT response must be accurately modeled to bridge hardware characterization and software reconstruction. While fundamental properties such as quantum efficiency (Figure 6.10) and nominal voltage scaling (Figure 6.4) remain consistent with the hardware measurements discussed previously, the IceTray implementation requires specific parameterizations that differ from standard experimental conventions. This section outlines these software-specific metrics to serve as a reference for future implementations. The full detector simulation chain utilizing these models is further detailed in Section 12.1.

It should be noted that the measurements informing these specific software parameters were conducted at room temperature on pre-production PMT models. This approach is justified because the final simulated charge is inherently scaled to the target nominal gain, effectively mitigating temperature-dependent gain variations. Furthermore, the temperature-dependent intrinsic dark rate is injected independently later in the simulation pipeline as a module background rate (Chapter 9). The physical geometry and baseline performance of these pre-production units match the final integrated PMTs, ensuring the validity of this parameterization.

Charge Parameters

The data for the SPE template and linearity presented in this section were measured by Yuya Makino at WIPAC (Wisconsin IceCube Particle Astrophysics Center) during the LOM development phase. Because the wuBase alters the PMT charge response and was not available for the standalone measurements conducted in this work, these external characterizations, along with fitting procedures by the author, are documented here for completeness and reference.

To reproduce the detector response, a PDF for the SPE charge distribution is derived from low-light illumination measurements ($\mu < 0.1$). As detailed in [159], the charge response $S(q)$ is modeled as a superposition of three distinct components:

$$S(q) = \underbrace{A_1 \exp\left(-\frac{q}{w_1}\right)}_{\text{Noise/Pedestal Tail}} + \underbrace{A_2 \exp\left(-\frac{q}{w_2}\right)}_{\text{Under-amplified}} + \underbrace{A_{\mathrm{G}} \exp\left(-\frac{(q-\mu)^2}{2\sigma^2}\right)}_{\text{1 PE Peak}}, \tag{6.3}$$

where q denotes the charge in units of PE. The model accounts for the electronic noise tail (A_1, w_1), under-amplified photoelectrons (A_2, w_2) and the nominal dynode multiplication statistics (A_G, μ, σ). To extract the intrinsic SPE parameters without bias, the model must account for MPE contributions that are inherently present even in unilluminated

dark rate measurements (similar to Equation (4.6)). Fitting a pure SPE function to the charge histogram would artificially inflate the shape parameters to accommodate the high-charge MPE tail. Instead, the fitting routine dynamically constructs an MPE template by iteratively convolving the base SPE function $S(q)$ with itself. The minimizer simultaneously optimizes the relative amplitudes of these 2 PE and 3 PE components alongside the seven shape parameters of $S(q)$. A charge threshold of 0.1 PE is applied to exclude the electronic baseline noise, as shown in Figure 6.11. For IceTray integration, the extracted Gaussian and exponential widths are normalized by the fitted Gaussian mean, forcing the primary peak to strictly 1.0 PE. Together with a discrete SLC mean parameter for incomplete waveforms, these normalized values constitute the final simulation template as listed in Table 6.1.

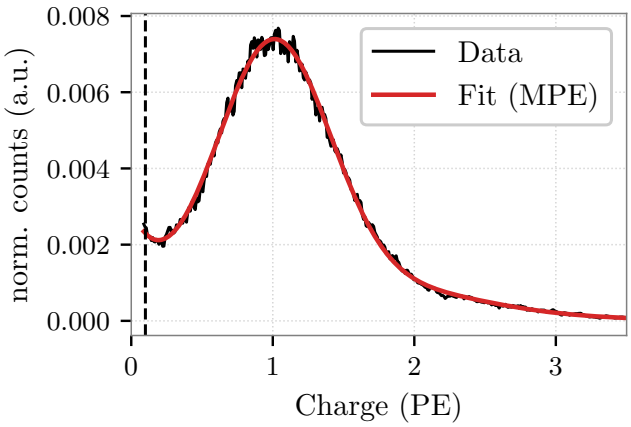


Figure 6.11: Exemplary SPE charge distribution fitting $S(q)$ to an NNVT PMT with a wuBase bias voltage of 85 V. Data provided by [160].

Table 6.1: Fit parameters for the SPE charge distribution model $S(q)$ for an exemplary NNVT PMT. Values are averaged across multiple bias voltages.

Component	Param.	Value
Noise Tail	A_1	6.9000
(Exp 1)	w_1	0.0013
Under-amp.	A_2	0.3418
(Exp 2)	w_2	0.3039
Main Peak	A_G	0.8826
(Gaussian)	μ	1.0000
	σ	0.4039

Linearity and Saturation

PMTs exhibit a linear response strictly within a limited dynamic range. For the new compact PMTs coupled with the active wuBase, the observed signal saturation is governed by two distinct mechanisms: intrinsic PMT saturation and electronics clipping. As introduced in Chapter 5, the readout employs two distinct signal paths: a standard anode channel (nominal gain 5.5×10^6 in the provided measurements) and an intermediate dynode eight channel (2.1×10^5).

At lower illuminations, the standard anode readout channel saturates at ≈ 75 PE. This occurs significantly earlier than a bare PMT on a passive base (which typically saturates at a few hundred PE), because it is governed by a hard clipping effect caused by the 12-bit analog-to-digital converter (ADC) reaching its maximum output. Conversely, true intrinsic space-charge effects within the PMT dynode cascade only dominate at much higher illuminations. This intrinsic behavior is captured by the intermediate dynode eight channel, which bypasses the early ADC bottleneck and remains linear up to roughly 5000 PE. Discriminating between these two channels to reconstruct the full dynamic range is a new feature that needs to be implemented in the processing pipeline in the future.

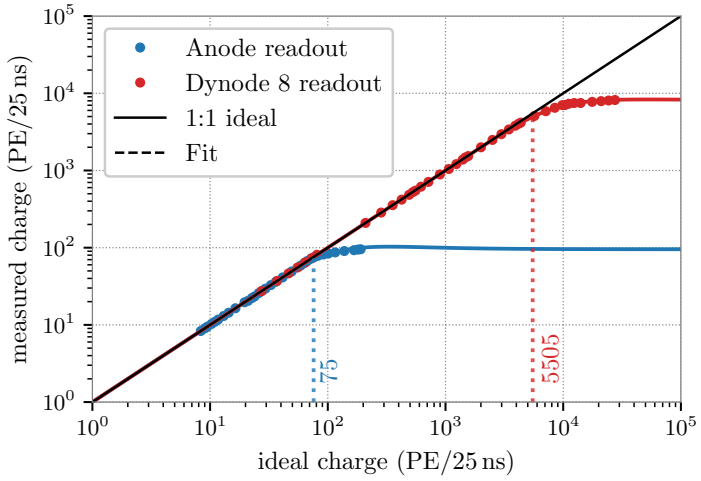


Figure 6.12: PMT linearity measurement for anode and dynode eight readouts. The measured charge N_{meas} is plotted against the expected true charge N_{ideal} in units of PE per 25 ns. Saturation occurs for higher illuminations, which is mitigated by transitioning to the dynode eight channel. Vertical lines indicate where the response drops below 90 % of the ideal charge. Data provided by [160].

The linearity of both channels was characterized by pulsing a light source from SPE level up to $\mathcal{O}(10^5)$ PE. The ideal expected charge, N_{ideal}, is extrapolated from source calibrations in the linear regime, while N_{meas} represents the actual recorded charge. Figure 6.12 illustrates the saturation behavior as a deviation from the ideal 1:1 correlation.

Within IceTray, saturation is conventionally modeled in two sequential steps: a continuous charge-attenuation function accounting for intrinsic PMT saturation, followed by a hard digital cut-off mimicking the ADC overflow. For the purpose of documenting the effective end-to-end response of the new modules in this work, the standard continuous saturation function is empirically fitted to the final measured photoelectron output of each readout channel. While this collapses the distinct physical bottlenecks into a single mathematical envelope, it provides a functional description of the module's observed dynamic range. These parameters are documented here as an empirical baseline, leaving it open for future software implementations to either utilize a unified effective parameterization or explicitly simulate the separate intrinsic and electronic limits. While the standard saturation model is typically parameterized in terms of measured current (amperes), the identical empirical function can be fitted directly to the derived photoelectron rate by extracting the parameters A, B and A' in units of PE per 25 ns. The saturation factor $S = N_{\text{meas}}/N_{\text{ideal}}$ is thus calculated using the inverse function:

$$S(N_{\text{ideal}}) = \frac{\ln(1 + A/N_{\text{ideal}})}{\ln(1 + A/N_{\text{ideal}}) + A' \cdot \exp(-B/N_{\text{ideal}})}, \tag{6.4}$$

where A dictates the saturation scale, B governs the curvature and onset threshold, and A' defines the absolute asymptotic saturation limit. To prevent mathematical degeneracy in the deep saturation regime ($N_{\text{ideal}} \to \infty$), the scaling parameter A is fixed at 1.0. Fitting Equation (6.4) directly to the PE datasets yields best-fit values of $B = 149.56$ and $A' = 1.05 \times 10^{-2}$ for the anode channel. Here, the parameterization acts as an effective empirical model for the ADC overflow. For the dynode eight readout, which represents

true intrinsic PMT saturation, the fit yields $B = 9913.87$ and $A' = 1.22 \times 10^{-4}$. This reflects its vastly extended linear range, with a 10 % deviation from linearity occurring at 5504.6 PE/25 ns for dynode eight, compared with 75.8 PE/25 ns for the anode.

Prepulses

To measure the timing shift between prepulses and main pulses, the PMT is operated in pulse mode at nominal voltage using a laser intensity of ~ 50 PE for main pulses and ~ 500 PE for prepulses. Figure 6.13 (left) compares the normalized average of over 100 000 waveforms for each domain, revealing a distinct time shift between the main photoelectron peak and the prepulse signal. Prepulses occur $\Delta t_{\text{ref}} = (17.3 \pm 0.3)$ ns earlier relative to main pulses due to the shorter electron trajectory starting at the first dynode. Since the electron transit time depends on the electric field strength, this shift is dynamically scaled in IceTray based on the operating voltage V of the module $\Delta t(V) = \Delta t_{\text{ref}} \cdot \sqrt{V_{\text{ref}}/V}$, where V_{ref} is the reference nominal voltage from Figure 6.4 and V is any PMT-specific operating voltage.

Since prepulses bypass the amplification of the first dynode ($\delta = 12 - 14$ [139]), they show a low probability of being detected using a triggered measurement (< 0.1 % from Section 6.2.2). As each pulse is sampled individually in IceTray, a threshold-free probability must be derived. For this, a reference measurement of the incident photon flux is taken using the previously calibrated photodiode, with the light intensity increased until more than 10 000 photons per $f = 100$ kHz LED flash are measured by a picoammeter ($I_D = (0.586 \pm 0.008)$ nA). The PMT is positioned behind a calibrated ND filter ($T_{\text{ND}} = (11.81 \pm 0.65)$ %). The PMT waveforms are fitted as shown in Figure 6.13 (right). The rising flank of the saturated main pulse and the baseline are fitted by an exponential function $A \exp(b(x - x_0)) + c$. The prepulse charge Q_{pre} is determined via integration within an ROI, from which the baseline fit is subtracted. The threshold-free prepulse probability P_{pre} is defined as the ratio of the number of detected prepulses N_{pre} to the expected number of primary photoelectrons N_{pe}. The number of prepulses is derived from the integrated charge divided by the charge of a single prepulse, $N_{\text{pre}} = Q_{\text{pre}}/(G \cdot e/\delta)$, where e is the elementary charge and $G = 5 \times 10^6$ is the PMT gain. The expected number of primary photoelectrons is calculated from the reference diode current, which

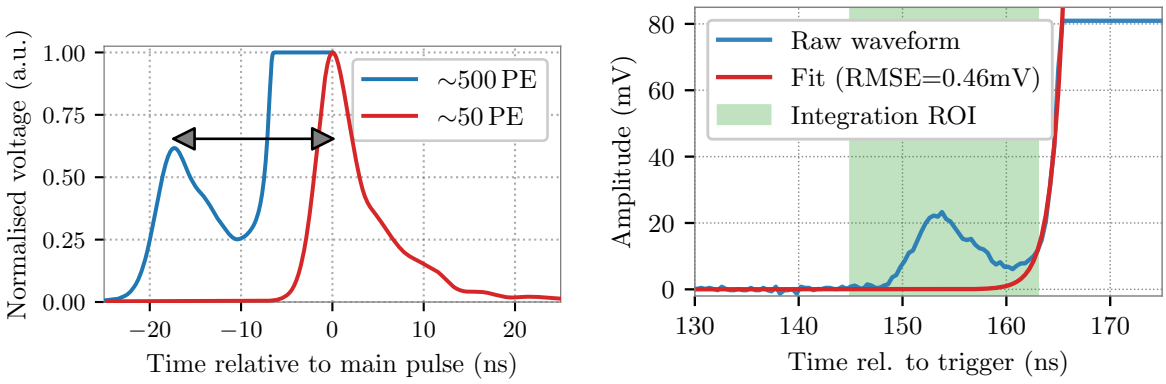


Figure 6.13: **(left)** Averaged normalized waveforms of main pulses (red) and prepulses (blue) for an exemplary NNVT PMT. **(right)** Representative determination of the prepulse charge.

is adjusted for the ND filter transmission and the respective quantum efficiencies, as per $N_{\mathrm{pe}} = [I_D/(f \cdot e \cdot \mathrm{QE_D})] \cdot T_{\mathrm{ND}} \cdot \mathrm{QE_{PMT}}$. Combining these two terms, the elementary charge e cancels out, yielding the probability

$$P_{\mathrm{pre}} = \frac{N_{\mathrm{pre}}}{N_{\mathrm{pe}}} = \frac{Q_{\mathrm{pre}} \cdot \delta \cdot f \cdot \mathrm{QE_D}}{G \cdot I_D \cdot T_{\mathrm{ND}} \cdot \mathrm{QE_{PMT}}}. \tag{6.5}$$

This results in a final, charge-based prepulse probability of $(4.41 \pm 0.34)\,\%$.

Jitter

As discussed in Section 6.2.2, while the transit time spread characterizes the timing resolution, the overall arrival time distribution around the main peak is fundamentally asymmetric. Following the standard IceTray implementation to model this delayed tail, the time shift t is parameterized using a Fisher-Tippett (Gumbel) distribution, where the probability density function is given by:

$$f(t) = \frac{1}{\beta} \exp\left(-\left(z + e^{-z}\right)\right) \quad \text{with} \quad z = \frac{t-\mu}{\beta}, \tag{6.6}$$

where μ is the location and β is the scale parameter. The scale parameter relates to the standard deviation of the distribution as $\sigma_{\mathrm{FT}} = \pi\beta/\sqrt{6} \approx 1.28 \cdot \beta$.

Figure 6.14 shows the fit of this model to the experimental data using a fit window of $t \in [-6, 15]\,\mathrm{ns}$. Averaging the results from three characterized PMTs, the model parameters are determined as:

$$\mu = (0.08 \pm 0.04)\,\mathrm{ns} \quad \text{and} \quad \beta = (1.66 \pm 0.04)\,\mathrm{ns}. \tag{6.7}$$

To evaluate the stability of this model, the fit was repeated using a narrower window restricted around the main peak region ($t \in [-5, 7]\,\mathrm{ns}$). This restriction reduced the resulting scale parameter β by only $\lesssim 1\,\%$. While a Fisher-Tippett function models the arrival time of main pulses better than a Gaussian, the tail produced by diffuse

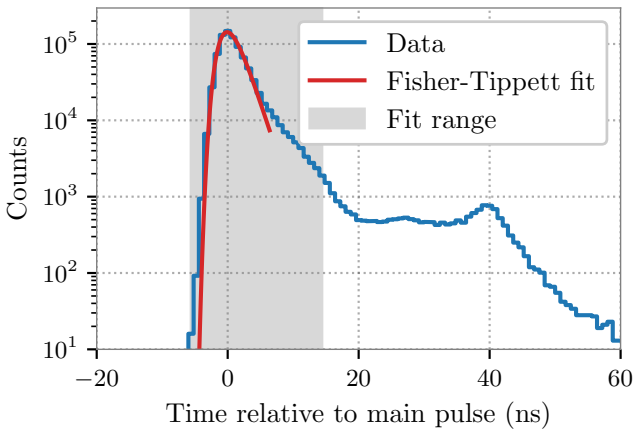


Figure 6.14: Fisher-Tippett fit to the main pulse arrival time peak for an exemplary NNVT PMT. The red line shows the model function fitted to the experimental data (blue), with the shaded region indicating the fitting range.

illumination is still not captured accurately, which is even more pronounced in these PMT models, as mentioned before. In IceTray, the specific time jitter Δt for each photon is sampled from this distribution using the inverse transform method:

$$\Delta t = \mu - \beta \cdot \ln(-\ln(u)), \tag{6.8}$$

where u is a uniformly distributed random number $u \in (0, 1)$.

Late Pulsing and Afterpulsing type I

The timing response following the main detection peak is characterized by electron backscattering at the first dynode. This is once again measured via the diffuse SPE illumination setup used in Section 6.2.2. To parameterize this behavior, a multi-component Fisher-Tippett model is fitted to the time distribution. The results are shown in Figure 6.15 with fit parameters listed in Table 6.2.

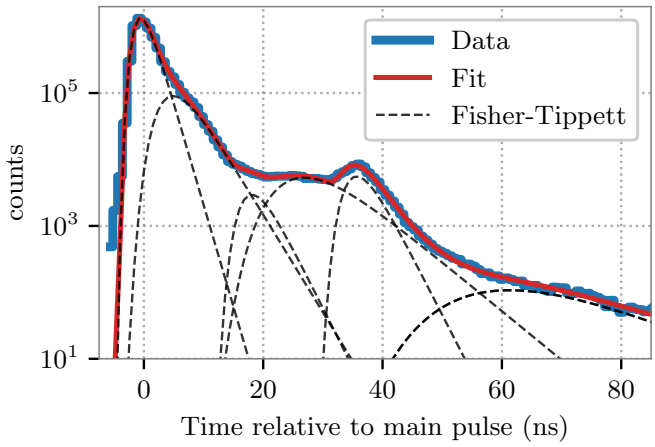


Figure 6.15: Multi-component Fisher-Tippett fit to the PMT timing distribution for an exemplary NNVT PMT. The main peak and its tail are modeled using a dual-component superposition. The subsequent components describe the inelastic shoulder and the distinct elastic backscattering peak at ~ 36 ns.

Table 6.2: Fit parameters for the multi-component Fisher-Tippett model. Component (C.) 4, with its narrow width, is assigned to the elastic backscattering peak.

C.	h	μ (ns)	β (ns)
0	5.14e6	−0.57	1.42
1	724e3	4.91	2.95
2	19.4e3	18.09	2.42
3	85.2e3	26.90	5.89
4	37.4e3	35.64	2.48
5	3.55e3	61.24	12.15

The observed structure is classified into a continuous shoulder arising from inelastic backscattering from afterpulsing type I (components 2,3,5) and distinct peaks arising from elastic backscattering from late pulses (component 4). Notably, the main peak is described by a superposition of the first two components (0 and 1): Component 0 captures the core Gaussian-like arrival time, while Component 1 models the immediate non-Gaussian tail originating from full photocathode illumination. This dual-component approach provides a more accurate representation of the detector response than the single-component jitter model discussed above.

In IceTray, this distribution is implemented as a mixture model. For each simulated photoelectron, the algorithm first selects a specific component i based on the relative probability weights derived from the fitted amplitudes h. The time delay is then sampled from

that component's specific Fisher-Tippett PDF. This method allows specific components to be tagged with physical source identifiers.

Afterpulsing type II

The simulation of ionization afterpulsing type II presents a challenge for the standard SPE probability approach. Since afterpulses often liberate multiple electrons from the photocathode, relying on a simple probability histogram based on SPE counts would significantly underestimate the total deposited charge. To address this, IceTray models the afterpulsing type II response by fitting averaged waveforms obtained under high illumination ($\sim$ 100 PE) as opposed to Section 6.2.2. This average waveform represents the product of the afterpulsing probability and the charge yield ($P(t) \times \bar{q}(t)$), ensuring that the total integrated charge in the simulation matches the physical detector output. To model the overall timing distribution, a superposition of Gaussian functions is fitted to the average waveform beyond 200 ns relative to the main peak. The resulting decomposition is shown in Figure 6.16 (left), with the corresponding fit parameters detailed in Table 6.3. Although additional Gaussian components could slightly improve the fit, the parameterization relies on a single representative PMT. Thus, the dominant source of data-to-MC deviation will ultimately stem from intrinsic PMT-to-PMT variations rather than the fit itself.

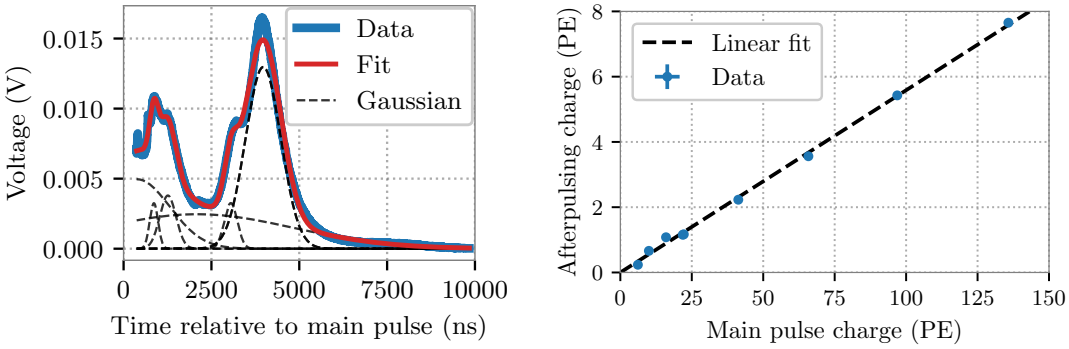


Figure 6.16: (left) Average waveforms, fitted with a sum of Gaussian components for an exemplary NNVT PMT. Note that the data is truncated below 200 ns. **(right)** Afterpulsing charge in units of PE as a function of average number of PE contained in main pulses, with a linear fit applied.

Table 6.3: Gaussian fit parameters for the afterpulsing distribution. The broad width of Component 4 suggests a continuum background rather than a distinct ion species.

Component	Weight h	Mean μ (ns)	Width σ (ns)
1	0.0050	381.9	930.42
2	0.0033	862.7	118.65
3	0.0038	1260.0	231.74
4	0.0025	2090.0	2737.53
5	0.0033	3051.8	184.53
6	0.0130	3973.7	502.90

The arrival time of afterpulsing is governed by the specific ion's mass-to-charge ratio ($t \propto \sqrt{m/q}$). The precise drift time is further influenced by the applied voltage, the internal electric field geometry and the physical dimensions of the PMT. As shown in the decomposition, the most prominent peak occurs at approximately 3980 ns. The qualitative peak structure matches initial measurements reported in [73 (p. 127)] and is attributed to Cs^+ ions [161]. Using this Cs^+ peak and its known mass ($m_{Cs} = 132.9$ u) as a calibration reference, the origin of the other main component at 1260 ns can be estimated. Using the relation $m_x = (t/t_{Cs} \cdot \sqrt{m_{Cs}})^2$, the peak corresponds to a mass of approximately 13 u. This mass is consistent with light carbon-containing ions, including CH^+, although an unambiguous species assignment is not possible based on this measurement alone. These ions may be produced by electron-impact ionization and fragmentation of residual hydrocarbon gases in the vacuum, such as methane (CH_4).

Conventionally, afterpulsing probability is defined as the ratio of afterpulse counts to primary pulses (as outlined in Figure 6.8). However, because IceTray assigns a standard SPE charge distribution to simulated pulses, a charge-based ratio is required instead. To measure this, the PMT is illuminated at varying light intensities, recording 20 000 waveforms (30 µs duration) per intensity. The signal charge Q_s is integrated over an afterpulsing window (100 ns to 12 µs after the main pulse). Baseline and dark rate contributions Q_b are subtracted using an identically sized, delayed window (15.1 µs to 27 µs), yielding the net afterpulse charge $Q_{AP} = Q_s - Q_b$. For normalization, the main pulse charge Q_{main} is integrated over a ±15 ns window around the main peak. Figure 6.16 (right) shows the mean Q_{AP} relative to Q_{main} for a representative NNVT PMT. The slope of a linear fit corresponds to the afterpulsing type II charge ratio, determined as (5.58 ± 0.06) %.

6.5 Scans across the photosensitive Area

While the measurements presented in Sections 6.2 and 6.3 characterize the overall response of the PMT under diffuse illumination, performance is not uniform across the photosensitive area. Variations in the enveloping glass, photocathode thickness, collection efficiency and electron path lengths due to geometry introduce position-dependent effects. To resolve these features, a dedicated 2D scanning setup is employed, as commissioned in [162] and further utilized in [73]. The setup, illustrated in Figure 6.17, places the PMT inside a Helmholtz coil system [163] to compensate for the Earth's magnetic field, which would otherwise deflect the electron trajectories, particularly for larger diameter

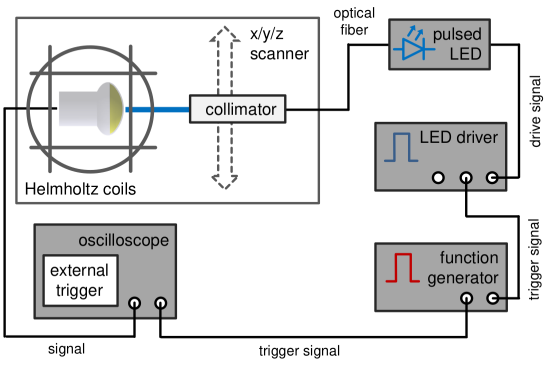


Figure 6.17: Schematic of the 2D scanning setup. Taken from [73].

PMTs. A 105 µm multimode optical fiber is mounted on a motorized XY-stage [156] at a fixed horizontal distance of 60 mm from the curved surface, corresponding to the focal point of the optical system. The light source is a pulsed LED[14] coupled to a collimator[15], producing a beam with a FWHM of ~ 0.9 mm on the photocathode. For this work, the 4-inch PMTs are scanned with a step size of 1.4 mm, resulting in approximately 4100 individual pixels across the photosensitive area.

Gain and relative Transit Time

In the first measurements, the PMTs are operated at their respective nominal voltage in SP-readout with a light intensity adjusted to ~ 0.7 PE per pulse. This allows for the extraction of gain-related parameters, transit time and pulse shape parameters for each individual pixel, analogous to Section 6.2.1. These scans were performed sporadically throughout development iterations, with only the key results presented here.

Figure 6.18 (left) displays the spatial gain distribution for a Hamamatsu 4-inch PMT. The response is homogeneous across the central region. Towards the edges ($r > 45$ mm), a characteristic drop-off is observed. Since gain is determined solely by the multiplication of photoelectrons that successfully trigger an avalanche, this spatial variation is driven by the internal electron optics rather than optical transmission or absolute collection efficiency. Photoelectrons originating from the periphery follow asymmetric trajectories through the focusing electric field. Consequently, they frequently strike the first dynode at non-optimal positions or grazing angles, yielding a locally reduced secondary emission coefficient. Furthermore, the subsequent collection efficiency between the first and second dynodes can be degraded for these peripheral trajectories. Towards the outer photocathode edge, photoelectrons miss the active multiplying structure entirely, causing the effective gain to gradually drop to zero.

The radial projection in Figure 6.18 (right) compares the NNVT and Hamamatsu 4-inch models. Both exhibit a stable plateau up to $r \approx 42$ mm with better focusing around

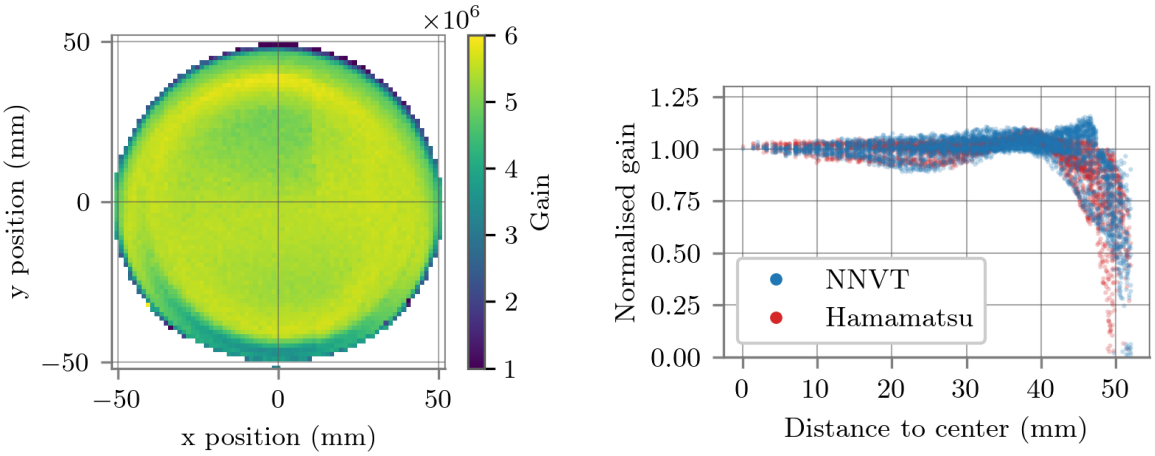


Figure 6.18: (left) 2D gain map of a Hamamatsu 4-inch PMT (BB9786). **(right)** Radial projection of the gain normalized to the central region ($r < 30$ mm), comparing the Hamamatsu (BB9786, red) and NNVT (PO4049, blue) 4-inch PMTs.

[14] PicoQuant PLS-8-2-719 driven by a PDL 800-B driver ($\lambda \approx 459$ nm).
[15] Schäfter & Kirchhoff 60FC-SMA-0-A7.5-01

the neck region $r \approx 37\,\text{mm}$. The Hamamatsu PMT maintains slightly better uniformity in the edge region ($r > 45\,\text{mm}$), whereas the NNVT response shows larger localized variations, likely due to differences in the focusing dynode.

The relative transit time in Figure 6.19, defined as the shift of the main peak arrival time relative to the center ($\overline{\text{TT}}_{r<30\,\text{mm}} = 0\,\text{ns}$), reveals a distinct asymmetry aligned with the PMT axis. Since the curvature of the first dynode is oriented towards the positive y-direction, photoelectrons emitted from the upper half of the photocathode travel a longer path through the weak-field region than those from the lower half, resulting in a measurable time difference of $\mathcal{O}(10\,\text{ns})$.

This effect was systematically studied via a dedicated COMSOL simulation of the electron optics, developed in [136, 137, 138] under the author's supervision. The simulation, exemplarily visualized in Figure 6.20 (left), traces individual electrons through the electromagnetic fields defined by a CT-scanned mDOM PMT geometry. It successfully reproduces the observed time structure, confirming that the asymmetry originates from the internal electron optics design.

These scans provided important feedback during the PMT development. Early NNVT prototypes exhibited severe transit time differences of $\Delta\text{TT}_{\text{max}} \approx 20\,\text{ns}$ between the opposing edges. While there is no requirement for this, feedback provided from these measurements resulted in a changed design to address this performance difference going forward. This also holds true for the other measurements, but it is easiest to demonstrate here with the noticeable design change. By modifying the curvature around the photocathode edge (see Figure 6.20, right) and focusing electrodes, the final model reduced this spread to $\approx 12\,\text{ns}$, resulting in performance similar to the Hamamatsu design.

To contextualize the performance, Figure 6.21 compares the spatial uniformity of the new 4-inch PMTs against the 3-inch mDOM (R15458-02), 8-inch D-Egg (R5912-100) and 10-inch DOM (R7081-02) PMTs. Overall, all PMTs show similar gain homogeneity (top left panel). The response is stable in the center as well as the neck region and then shows more variation towards the edges. Since these point clouds are hard to interpret, the

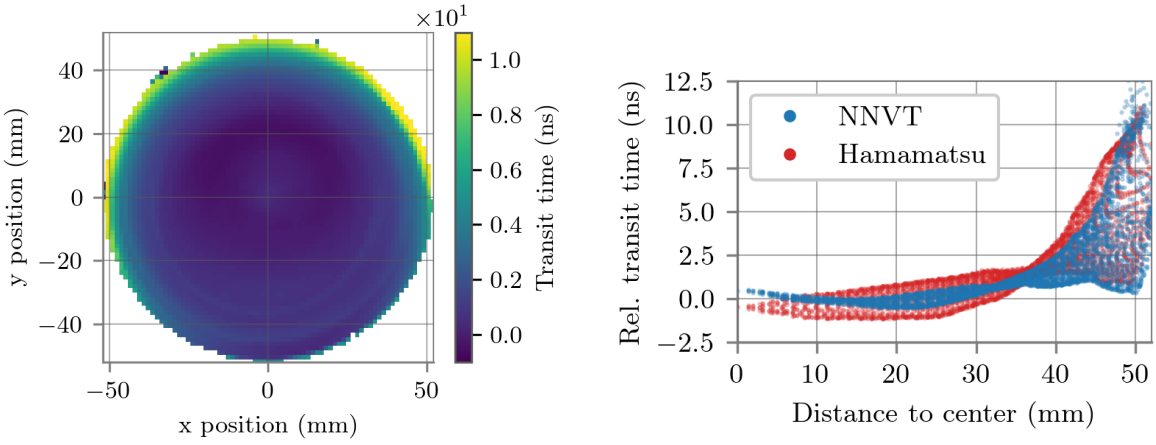


Figure 6.19: **(left)** 2D map of the relative transit time for an NNVT PMT (PO4049), showing the impact of the dynode orientation. **(right)** Radial projection of the transit times normalized to the central region ($r < 30\,\text{mm}$) for both 4-inch PMT models.

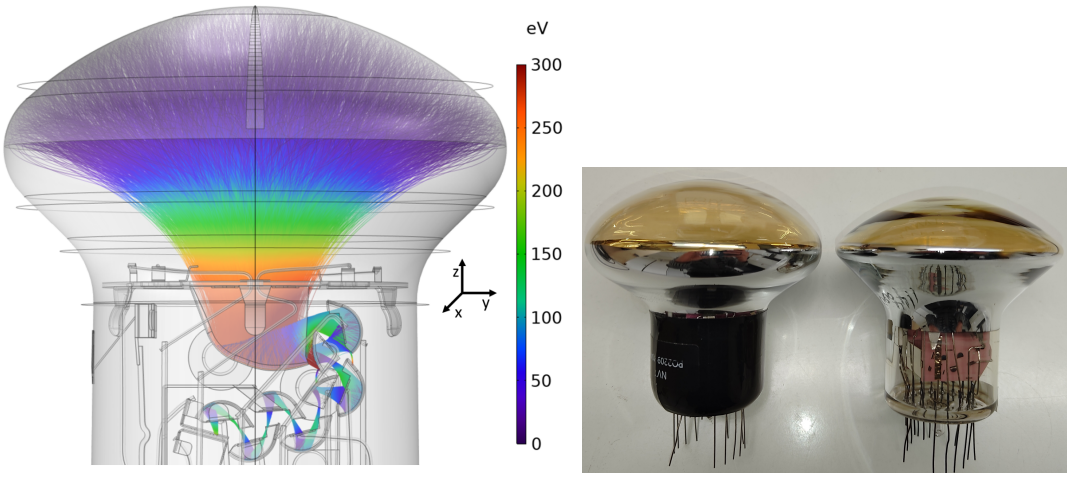


Figure 6.20: (left) Exemplary visualization of the electron trajectory simulation for 5000 photoelectrons, with color indicating kinetic energy. Taken from [137]. **(right)** Comparison of the NNVT final production model N2041 (left) and prototype (right). While the outer dimensions are unchanged, the curvature around the photocathode periphery is modified and a conductive coating is added beneath the black insulation.

standard deviation of pixels inside a relative radius window (bottom panel) quantifies the uniformity. Because D-Egg and DOM PMTs operate at twice the nominal gain, the relative gain is used for comparison (bottom left). Using this metric, mDOM PMTs prove

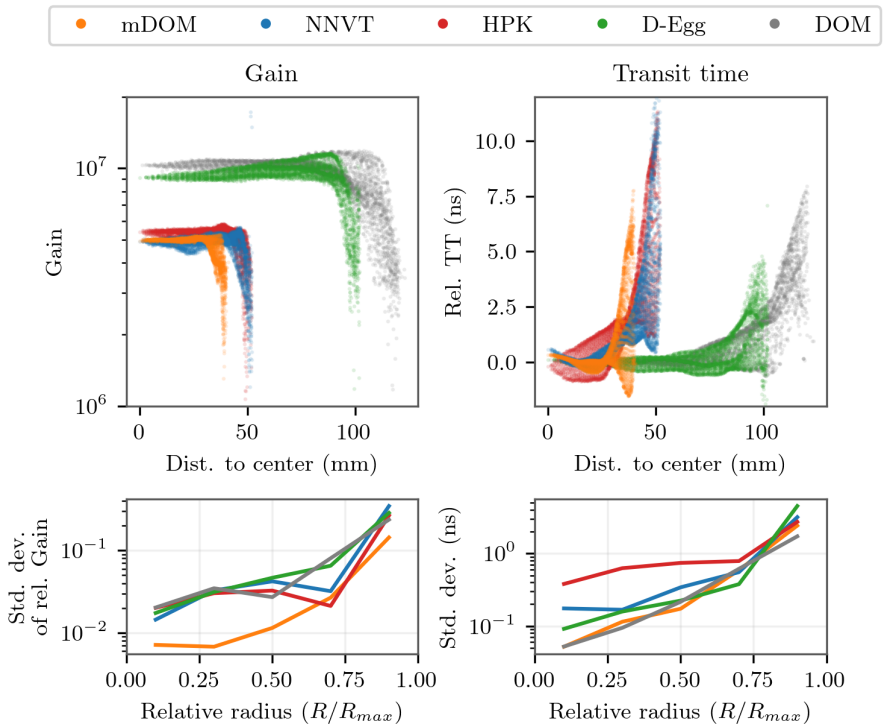


Figure 6.21: Comparison of spatial uniformity across different PMT models including mDOM, D-Egg and DOM data from [73]. **(left)** Gain. **(right)** Relative transit time. Standard deviations shown below are calculated from all pixels within the respective radial intervals. The gain values are normalized, as D-Egg and DOM PMTs are operated at twice the nominal gain.

to be the most stable, while all other PMTs are comparable, establishing both 4-inch models as viable. Regarding transit time (right panels), the 4-inch PMTs investigated in this work generally exhibit a wider maximum-to-minimum spread at the outer edges. This is an expected consequence of scaling the photocathode area while maintaining a compact dynode chain, which reduces the space available for corrective focusing fields. Consequently, a longer tail is observed in the TT distributions discussed in Section 6.2.2. Under diffuse illumination, however, this spatial variation is effectively averaged out and remains well within the TTS requirements.

As in previous sections, other parameters (TTS, RT, FT, pulse-width, SPE resolution, charge resolution) were investigated as well, yielding a similar conclusion. Both 4-inch variants are comparable in terms of overall homogeneity, with NNVT PMTs performing slightly better in the center but worse towards the edges compared with the 4-inch Hamamatsu version.

Quantum Efficiency

While Section 6.3 presents the QE over the full photocathode, spatially resolved scans reveal the local structure of the internal components and allow for further interpretation of the photocathode deposition. The same scanning setup was used, but the PMT is operated in current mode at high light intensities using the aforementioned QE-base and a picoammeter to measure the photocurrent generated by the 459 nm LED. The current of the reference photodiode is measured after every 50 scan points to correct for light source drifts. Figure 6.22 shows that in the center, until $\sim r = 30$ mm, the measured response is stable, where internal components (most notably the first dynode and front plate) are resolved. From 30 mm to 45 mm, an enhancement in the effective photon conversion probability is observed. While the intrinsic QE of the photocathode material remains constant, photons transmitted through the semitransparent photocathode are reflected by the conductive coating on the PMT *neck* back into the photocathode, increasing the interaction probability (photon recycling). This geometric effect is the reason why standard QE measurements, as performed in Section 6.3, yield higher values under full-face illumination

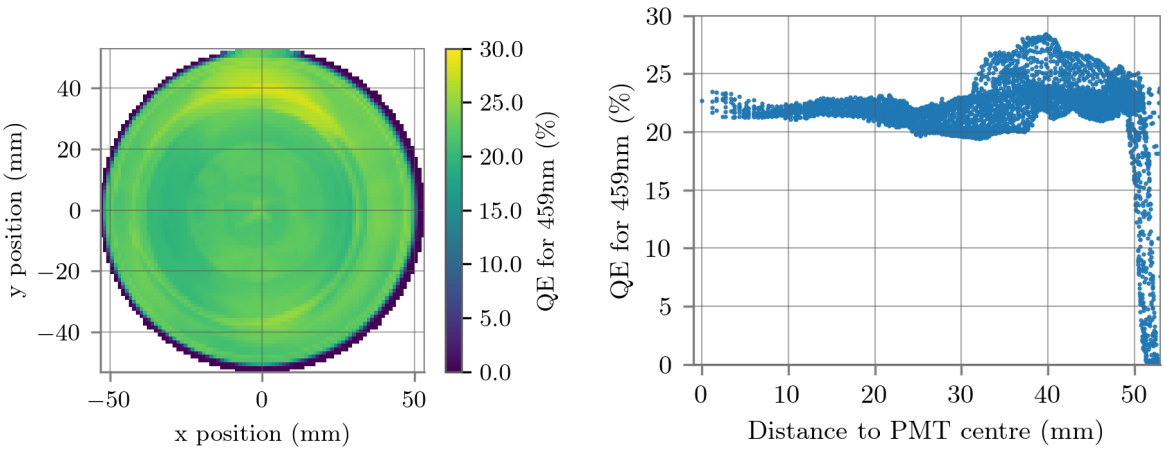


Figure 6.22: (left) Spatially resolved quantum efficiency map of PO2206 (initial production batch of NNVT's final model) at 459 nm. The distortion towards negative x values originates from a slight tilt in the horizontal PMT alignment. **(right)** Radial projection of the QE.

compared with a narrowly collimated central spot. With an average of $(22.11 \pm 1.55)\,\%$ across $r < 40\,\text{mm}$, this spatially integrated effective QE matches the results from the full cathode illumination (see Figure 6.10; $(21.90 \pm 0.08)\,\%$ at 460 nm). The sharp drop-off at the edges $r \gtrsim 45\,\text{mm}$ is primarily driven by total internal reflection. Because the collimated beam strikes the curved glass envelope at the PMT periphery at grazing angles, a significant fraction of the incident light is reflected outward rather than transmitted into the photocathode.

Localized inhomogeneities, such as the feature observed at $y \approx 40\,\text{mm}$, are attributed to uneven evaporation of the bialkali photocathode material. A dedicated study investigating the wavelength dependence of these features [164], supervised by the author, confirmed that such variations are more pronounced at longer wavelengths ($> 450\,\text{nm}$), which are more sensitive to the photocathode work function and thickness.

6.6 Summary

This chapter presents a comprehensive characterization of the 4-inch PMT models developed by Hamamatsu and NNVT, intended for IceCube-Gen2 and utilized in the LOM16 and LOM18 optical modules. Based on the detailed evaluation, both the Hamamatsu R16293 and the NNVT N2041 exceed the performance requirements set by the collaboration, as listed in Table 6.4. Thus, both models qualify as viable baselines for mass production.

Table 6.4: Performance summary of the 4-inch PMT intended for IceCube-Gen2. Mechanical specifications (first three rows) are derived from an inspection of 210 Hamamatsu and 21 NNVT PMTs measured at Chiba University. The performance parameters below represent the mean of three characterized PMTs from each manufacturer, averaged over temperatures between $-40\,^\circ\text{C}$ and $-10\,^\circ\text{C}$. The table is grouped by measurement conditions: TTS and pulse type probabilities were measured at single-photoelectron levels ($\mu < 0.1\,\text{PE}$), whereas charge properties and pulse shape statistics were measured at $\mu \approx 0.7\,\text{PE}$.

Performance metric	**Requirement**	**Hamamatsu**	**NNVT**
Mechanical length (mm)	< 106 mm	104.28 ± 0.22	104.27 ± 0.57
Mechanical diameter (mm)	< 104 mm	102.08 ± 0.23	103.82 ± 0.18
Max. neck diameter (mm)	< 55 mm	53.29 ± 0.12	53.3 ± 0.2
Top QE	-	25.8 ± 0.4	31.1 ± 1.5
QE at 400 nm	$> 25\,\%$	25.4 ± 0.5	30.3 ± 1.4
$\text{TTS}_{\text{G,peak}}$ (ns)	< 3.4	2.49 ± 0.09	1.78 ± 0.05
TTS_{num} σ (ns)	-	3.36 ± 0.08	2.80 ± 0.06
Pre-/Late-/After pulsing (%)	$< 1/5/15$	0.04/4.83/4.51	0.02/3.65/4.13
Nominal voltage (V)	< 1500	971.6 ± 4.6	757.6 ± 45.7
Peak-to-valley ratio	> 2	3.22 ± 0.07	2.85 ± 0.11
Charge resolution (%)	-	39.51 ± 1.10	37.42 ± 1.18
Rise time (ns)	< 5	1.70 ± 0.06	1.83 ± 0.07
Fall time (ns)	-	1.63 ± 0.04	1.66 ± 0.04
Pulse width (ns)	-	3.29 ± 0.08	3.30 ± 0.05
Dark rate	*As low as possible*	56.3 ± 2.7	94.0 ± 4.8

One notable characteristic of these compact 4-inch PMTs is the pronounced asymmetric tail in the timing response, quantified by the discrepancy between the Gaussian and numerical standard deviation for the TTS. This is a direct consequence of the "compactness compromise" as the limited distance from the photocathode to the first-dynode, required to fit the tight module envelope, results in larger relative path length differences for photoelectrons originating from the photocathode periphery.

Crucially, the measurements detailed in this chapter served as more than a final validation, as the feedback provided by this work played an integral part in the iterative development cycles spanning several years. The measurements presented here directly influenced the hardware design, for example driving the optimization of the NNVT photocathode edge geometry to suppress transit time asymmetries observed in early prototypes.

This successful R&D cycle yielded two mature PMT models with nearly identical geometries. It is notable that NNVT, initially less established in this form factor, has achieved competitive performance levels. In fact, a general trend of better timing resolution and slightly higher peak QE for the NNVT variants was consistently observed throughout development and verified against manufacturer quality assurance data for each production batch. However, regarding the specific QE values listed above, manufacturer data indicate that the three NNVT samples represented the high-end of their production batch, while the Hamamatsu samples represented the low-end of theirs. Consequently, the true performance gap in QE is likely smaller than observed in this study. The availability of two reliable versions is of strategic importance for IceCube-Gen2. With a production target of 10 000 optical modules over a span of approximately seven years, requiring a steady supply of $\sim 20\,000$ PMTs per year, dual-sourcing is essential to mitigate supply chain risks and ensure timely deployment.

Future experimental and simulation studies should consider several aspects that were beyond the scope of this work: The characterization presented here focused on the bare PMT performance to establish baseline parameters for selection and simulation. Future measurements may be repeated with the wuBase (see Chapter 5) attached, as the specific input impedance and signal shaping of the readout electronics will influence effective saturation limits, as well as charge and timing resolution. In particular, this affects charge resolution, peak-to-valley ratio and pulse shape parameters (Section 6.2.1) as well as TTS and pulse type probabilities (Section 6.2.2), whereas dark rates (Section 6.2.3) and uniformity scans (Section 6.5) are not expected to be affected.

PMT performance is susceptible to external magnetic fields, such as the terrestrial field of approximately 54.7 µT at the South Pole [165]. The Lorentz force deflects (photo)electron trajectories, changing gain, collection efficiency and transit time spread, where the impact of this effect scales with the internal PMT volume. While the 10-inch DOM PMTs and 8-inch D-Egg PMTs required magnetic shielding [119], measurements in [73 (p. 75)] on 3-inch mDOM PMTs showed a maximum efficiency deviation of only 5 % in DE, deeming such shielding unnecessary. The new 4-inch PMTs represent an intermediate case where the necessity of magnetic shielding has not yet been definitively established.

7 Mechanical Prototyping and Component Validation

With the performance of the 4-inch PMTs investigated in Chapter 6, their suitability as the core sensors for both the current IceCube Upgrade LOMs and the future Gen2-OM has been confirmed. As introduced in Chapter 5, the LOM16 and LOM18 serve as Gen2-OM precursor modules deployed as part of IceCube Upgrade. Having established the module concept and finalized the PMT development, the next stage was the module's mechanical realization. Achieving this required comprehensive mechanical prototyping, which is the focus of this chapter.

While the overarching Gen2-OM development is a collaborative effort, with the final IceCube Upgrade LOMs integrated by dedicated teams in Madison and Chiba, it is important to clearly delineate the work presented herein. Unless explicitly stated otherwise, the prototyping efforts, experimental setups and physical measurements presented throughout this chapter were conceived, constructed and carried out independently by the author. Within the scope of these developments, the majority of the foundational engineering was tailored for and directly enabled the production of the LOMs, while the development of the optical gel pads serves as a universal coupling solution applicable to both the LOMs and the Gen2-OM. Furthermore, a pneumatic integration mechanism is proposed to resolve the mass production bottlenecks for the future Gen2-OM. This chapter is divided into two primary areas of development: Section 7.1 details the mechanical prototyping efforts undertaken for module assembly, while Section 7.2 presents the validation of individual physical components.

7.1 Prototyping Efforts

The development of the module's mechanical realization required multiple iterative design phases. While the full scope of this engineering effort is extensive, this section focuses exclusively on the author's primary mechanical prototyping developments that directly influenced the assembly procedure of the deployed LOMs, while also proposing scalable solutions for the future Gen2-OM. Key independent contributions include the development of gel pad geometries and casting procedures (which established the thermoforming process now used for IceCube Upgrade LOM production) alongside PMT placement optimization. In addition, substantial independent R&D focused on mechanical support structures, internal module layouts and reliable pushing mechanisms to couple the gel pads to the pressure vessel.

7.1.1 Gel Pad Casting

The optical gel coupling the PMTs to the pressure vessel is a transparent Shin-Etsu silicone elastomer created by mixing two components (A and B). While a standard 1:1 ratio is typical, stiffer gel pads can be produced using a 1:1.25 ratio (as used for the LOM18; 1:1.2 for the LOM16). However, ratios exceeding 1:1.5 alter the cured gel's optical

properties [166 (p. 33)] (samples provided by the author). The liquid mixture is degassed at 5 mbar to eliminate entrapped air bubbles from mixing. Reducing the pressure below this value must be avoided, as it evaporates the hardening agent contained in component B and prevents complete curing. After pouring the gel into molds to form the gel pads (which inevitably entraps air), a secondary, geometry-specific degassing profile is applied. While mDOM mass production utilizes automated metering and mixing machinery[1], these localized prototyping procedures were executed manually.

Gel Pad Designs

Early iterations explored in situ casting, constraining the gel directly against the PMT using transparent acrylic shells. This aimed to eliminate the need for secondary optical coupling during assembly. While the material's refractive index does not alter the total internal reflection (TIR) angle, the shell walls needed to be thin to prevent photon loss at the photocathode edge. Although optical clarity at room temperature looks promising, environmental testing (500 mbar internal underpressure and temperatures down to −40 °C) revealed a mismatch in the coefficients of thermal expansion. Upon cooling, significant delamination formed between the acrylic shell and the gel, creating an additional optical air boundary that negated the gel's purpose. In combination with difficult degassing procedures during assembly, this ultimately led to the method's rejection, as no suitable transparent material could be sourced.

Alternatively, a solid acrylic lens was also considered. However, this still necessitated thin optical gel layers at both the PMT and pressure vessel interfaces to ensure genuine optical as well as mechanical coupling. Beyond the high manufacturing costs, rigid lenses lack the mechanical flexibility of cured silicone to safely compensate for the pressure vessel's inward deformation under refreezing ice (Section 7.2.1) or internal holding structure shrinkage. Both discarded designs are shown in Figure 7.1.

Following these attempts, trials utilizing a removable two-part mold inside the pressure vessel, where the gel pad was partially cast before being subsequently coupled to the glass pressure vessel using a liquid gel "puddle" during assembly, showed promising results. To further optimize for mass production, this concept evolved into the *pre-casting* method,



Figure 7.1: (left) Acrylic shell concept for in situ casting using an overflow basin, with arrows indicating gel flow. **(center)** Thermal delamination in the acrylic shell (shown for an opaque material for visibility) after 4 days at temperatures of −40 °C. **(right)** Solid acrylic lens requiring interface gel layers (coupled to a pressure vessel).

[1] MeterMix PAR 30

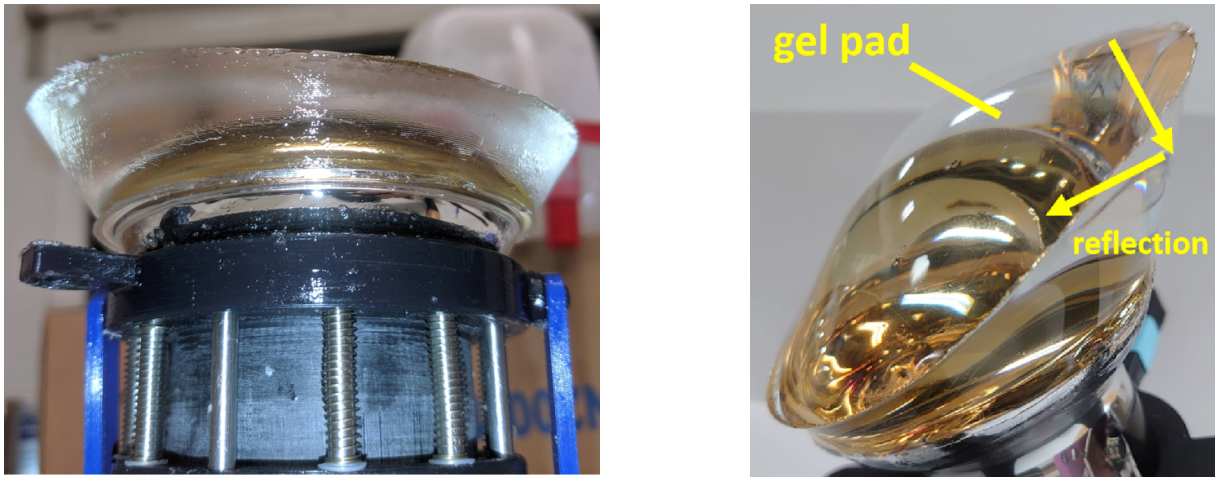


Figure 7.2: **(left)** Early pre-cast gel pad showing surface roughness from a 3D-printed mold used in a mechanical prototype utilizing springs to press the gel pads towards the inner pressure vessel surface. **(right)** Refined gel pad cast using a smooth thermoformed sheet overlay, enabling efficient TIR on and transmission through the side walls. Picture courtesy of the IceCube collaboration.

where the gel pads are fully cast and rotationally locked onto the PMTs in modular subassemblies outside the module prior to integration (see Section 7.1.2).

These two-part molds ensure consistency by constraining the PMT height and locking the gel pad rotation to the PMT's symmetry axis. Initial 3D-printed molds yielded insufficient surface finishes. A smooth side wall is critical to enable TIR, acting as a reflector for forward-going light and a clean transmission boundary for light incident from the sides.[2] To achieve this, the optimized 3D-printed molds were overlaid with thermoformed sheets, ensuring an optically smooth surface finish and mold reusability.[3] This thermoforming method was adopted for official LOM production. Figure 7.2 compares an early 3D-printed prototype with a refined thermoformed pad.

Material Selection

To identify mold materials that release cleanly while maintaining an optically smooth surface finish, dedicated adhesion tests were conducted on various polymers. Test samples were machined into smooth disks (30 mm in diameter and 10 mm thick, tapped with an M6 thread) and submerged 3 mm into freshly degassed optical gel contained in a petri dish. Following a secondary degassing step to remove any placement-induced air bubbles, the gel was allowed to fully cure. Adhesion properties were quantified using a custom testing assembly established in [73 (pp. 156–159)], which was utilized and adapted by the author to systematically evaluate the newly proposed mold materials. The petri dish containing the cured gel is secured within a rigid aluminum frame and the sample disk is attached to a force gauge mounted on a linear stepper motor. The motor extracted the sample at a constant velocity of 0.1 mm/s, continuously recording the vertical pulling force against displacement until complete detachment occurred. The resulting force-displacement curves provide insight into the demolding mechanics. Ideal mold materials exhibit a low peak

[2] While the PMT interface is always clear and the bottom of the gel pad receives an additional liquid layer during assembly, the surface finish on the side wall significantly impacts photon collection efficiency.

[3] Achieved using a *Mayku FormBox* vacuum forming machine with *Clear Sheets*.

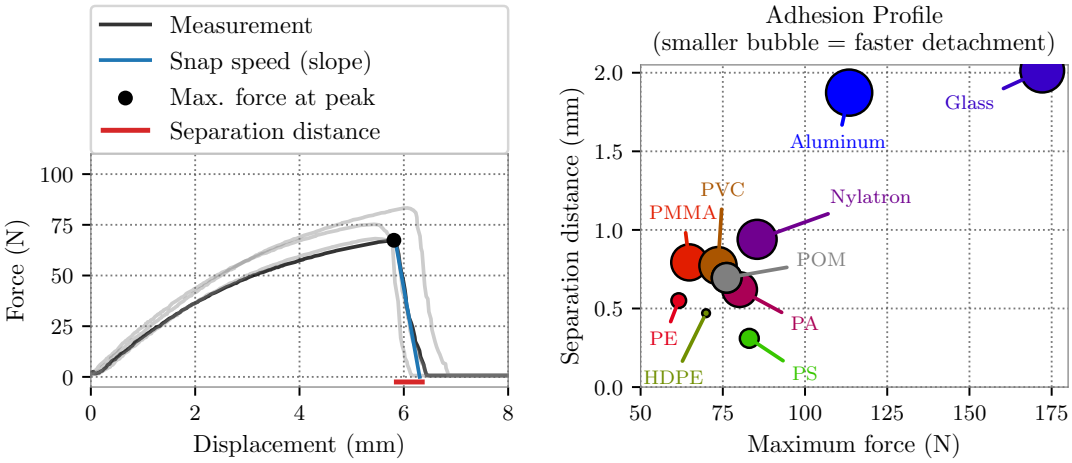


Figure 7.3: (left) Force-displacement curves for PVC, illustrating the extraction of peak adhesion force, separation distance and detachment snap speed from a representative measurement, highlighting one out of four samples. **(right)** Adhesion performance summary. Mean peak force is plotted against separation distance, with the logarithmic bubble size representing the detachment speed for each material.

adhesion force, a short separation distance (indicating little elastic stretching of the gel before detachment) and a steep force decay slope, corresponding to a clean release rather than gradual peeling. Figure 7.3 (left) illustrates the extraction of these three distinct metrics using PVC measurements as a representative example.

A performance summary of all tested candidates is presented in the right panel. According to these measurements, the most suitable mold materials are polyethylene (PE), high-density polyethylene (HDPE) and polystone (PS), which populate the bottom-left quadrant of the plot. They are plotted with noticeably smaller bubble sizes, representing a fast detachment snap speed, alongside a good balance of low separation force and minimal stretching distance. Conversely, aluminum and glass were included as baseline references because it was known beforehand that the gel strongly adheres to them. As expected, these materials yielded high resistance and appear in the top-right of the plot, demonstrating exceptionally strong mechanical coupling to the pressure vessel. Overall, the results show that the majority of the tested bare polymers perform comparably well and serve as suitable mold materials. This finding was further confirmed during later thermoforming trials, where sheet materials such as polyethylene terephthalate glycol (PETG), PE and polypropylene (PP) consistently yielded clean releases. Additionally, various demolding agents (including powders, sprays and soapy water) were evaluated. However, the results provided no empirical evidence that these agents significantly improved the release process compared with the bare thermoformed polymers like PP or PETG. To maintain a streamlined assembly procedure, demolding agents were omitted in all subsequent prototyping and the final IceCube Upgrade LOM assembly.

7.1.2 Mechanical Prototype for Gen2-OM

While the LOM16 and LOM18 successfully demonstrated the viability of the multi-PMT and gel pad concepts, their current manual assembly techniques (particularly the complex integration process involving the backfilling of partially pre-cast gel pads inside the

pressure vessel) are not readily scalable to IceCube-Gen2 production requirements. The time constraints of large-scale production therefore necessitate a streamlined assembly process. As a result, the mechanical architecture of the LOM16 was selected as the primary baseline to optimize the final Gen2-OM for high-throughput mass production. To ease internal space restrictions and relax mechanical tolerances, the Gen2-OM pressure vessel dimensions are slightly enlarged from this baseline (outer diameter from 312 mm to 318 mm, height from 444 mm to 486 mm).

This subsection presents a mechanical prototype tested within the tighter constraints of the existing LOM16 pressure vessel, designed specifically with mass production in mind. It revives the use of fully pre-cast gel pads to eliminate internal liquid backfilling, reducing the integration time. Precision component positioning and controllable contact forces are achieved using an integrated pneumatic pushing mechanism. Driven via pneumatic tubing, scissor lifts uniformly press the PMT subassemblies against the inner surface of the pressure vessel and remain operable under simultaneous vacuum and low-temperature conditions. To ensure scalability, the lifts are reusable across assemblies and the holding structure is optimized for injection molding. The prototype evaluates two optical coupling configurations: dry-fitting bare pre-cast pads directly against the glass as well as applying a thin interface layer of liquid gel.

Assembly

This setup represents the third mechanical prototype iteration designed, assembled and tested entirely by the author, replacing an earlier lever-activated spring design (see Figure 7.2, left) that lacked fine force control. In the current iteration, 3D-printed mounting brackets featuring guide grooves are glued to the PMT necks, locking the fully pre-cast gel pad's rotation to the PMT's symmetry axis during casting. Drawing on previous designs, this assembly utilizes a central vertical shaft for precise internal alignment. However, rather than anchoring the shaft (and thus all internal components) to the curved poles of the pressure vessel, it is temporarily referenced to the flat sealing surface of the pressure vessel via a metal cross during integration. The 3D-printed holding structure is loaded with four opposing PMT subassemblies (two polar, two equatorial). Each connects to a pneumatic scissor lift operated via a shared 2.5 mm inner diameter pneumatic tubing loop, which distributes 4 bar of compressed air to activate all lifts simultaneously. Crucially, this shared pneumatic line acts as an auto-balancing system. Unlike rigid mechanical springs, which would exert uneven forces if the PMT subassemblies or gel pads exhibit slight manufacturing tolerances, the interconnected pneumatic loop inherently equalizes the applied force across all PMTs, ensuring uniform optical coupling. When retracted, the lifts provide > 5 mm of horizontal clearance per PMT, accommodating mechanical tolerances even within the narrow LOM16 vessel. The entire structure, mounted to the metal cross, is lowered via guide rods into the pressure vessel hemisphere secured inside a vacuum-capable rig. Once seated, the guide rods are removed, the vacuum lid may be placed and the scissor lifts are pneumatically engaged to push the PMTs outward against the glass. The assembly rig and pushing mechanism are shown in Figure 7.4.

To evaluate coupling strategies, the two polar PMTs were prepared differently: one was dry-fitted directly to the glass to assess whether a liquid interface could be entirely omitted to accelerate assembly, while the other received a manually applied (partially cured) liquid gel layer to ensure conformal coupling. Applying a thin surface coating of liquid gel differs from and is significantly faster than the backfilling method used during

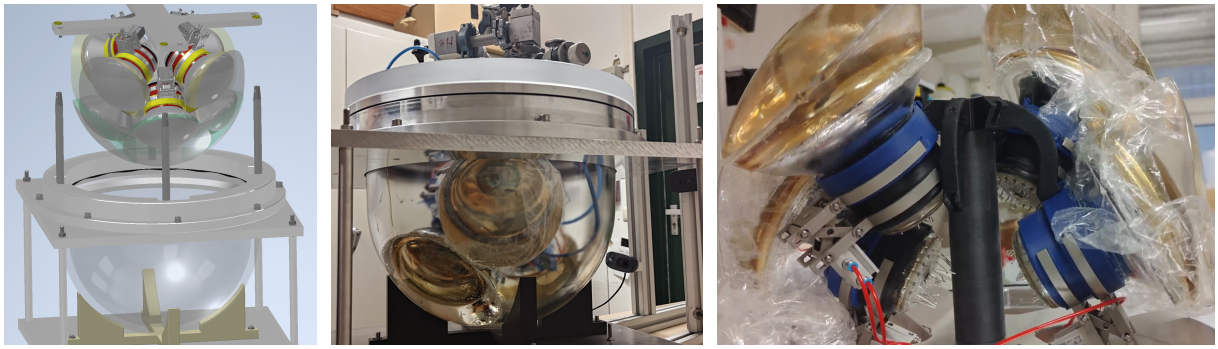


Figure 7.4: (left) CAD rendering of the assembly rig. **(center)** Photograph of the loaded internal assembly. **(right)** Photograph of the scissor lift pushing mechanism in the retracted state before assembly. Gel pads are protected by cling film or thermoformed PETG sheets.

LOM production. For future mass production, this manual application could be replaced by the motorized dip-coating method developed for the WOM [167], where withdrawal speed from a gel reservoir precisely dictates the resulting layer thickness. Once coupled, the assembly undergoes a degassing profile, which also proved to be viable during LOM production: the internal pressure is slowly ramped down to 5 mbar, held for one hour to evacuate entrapped air bubbles and then slowly returned to atmospheric pressure. The module cures for a minimum of two days, with the pneumatic lifts actively maintaining constant pressure on the gel interfaces. While the setup can utilize a heating cabinet for consistent environmental conditions, tests demonstrated that maintaining a vacuum *during* the extended curing phase severely compromises the gel's structural integrity.

Following the curing phase, the pneumatic pushing mechanism and the temporary metal cross can be fully detached and immediately cycled into the next assembly. While the internal vertical shafts of the upper and lower hemispheres can be structurally linked during final module integration, the PMT subassemblies themselves are securely anchored to the glass solely by the cured gel pads. The mechanical adhesion and structural integrity of this silicone interface are exceptionally robust, as reversible tests show that a vertically applied shear load of 10 kg on a coupled PMT induces no permanent deformation or delamination, demonstrating that the cured gel bond alone is sufficient to maintain the internal positioning without additional mechanical clamping.

In-Module environmental Testing

To verify the long-term stability of the optical coupling, the fully loaded assembly was subjected to in-module environmental conditions: an internal pressure of 500 mbar and an ambient temperature of −40 °C. For continuous visual monitoring, four strategically placed webcams[4] recorded time-lapses of the gel-glass interfaces. The initial test subjected the assembly to combined vacuum and freezing conditions for three days, as shown in Figure 7.5.

For the dry-fitted polar PMT, small air pockets emerged at the interface immediately upon the application of the vacuum. However, these bubbles were strictly confined to the boundary between the solid gel pad and the glass pressure vessel (not trapped within the gel pad) and their size and distribution remained entirely static over the three-day duration. Future studies should quantify any optical performance penalty associated with

[4]Logitech C270 (V-U0018), capturing 480 × 640 resolution images at defined intervals.

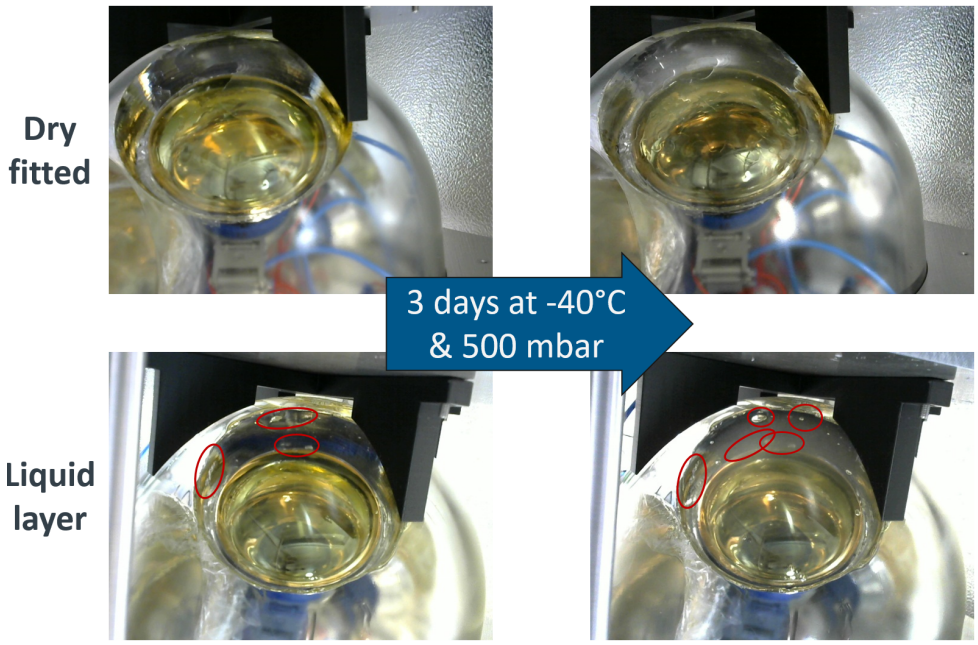


Figure 7.5: (left) Start of the test after evacuation and activation of the freezer. **(right)** After 3 days at −40 °C and 500 mbar internal pressure. The top row shows the dry-fit coupling and the bottom row shows the liquid-layer coupling, with red circles highlighting bubbles on the outer side walls rather than at the glass interface.

these boundary bubbles and weigh it against the assembly time saved by omitting the liquid interface. Furthermore, as this represents a single, first-time trial of dry-fitting fully pre-cast pads, these bubbles may simply be an artifact of the initial manual procedure. Consequently, optimizing the dry-fit method remains a promising approach worth pursuing further. For the PMT utilizing the liquid gel interface, the bubbles visible near the pad's rim in the initial frame resulted solely from manually applying excess liquid gel (a handling artifact of spreading gel onto the side walls) rather than an interface failure. Notably, no significant bubble formation or delamination occurred at the optical boundary during the test. Due to equipment limitations preventing prolonged simultaneous vacuum and cooling, a subsequent long-term test was conducted solely at −40 °C for more than seven months. As shown in Figure 7.6, the long-term observations are consistent with the short-term findings. Both the dry-fit and liquid-layer interfaces remained structurally and optically stable, with no new defects or delaminations emerging over the seven-month period.

While a single-person prototyping iteration naturally exhibits handling imperfections compared with the heavily refined IceCube Upgrade LOMs developed over years by dedicated teams, these results are promising as the initial bubble formations parallel the challenges faced during early LOM R&D. More importantly, the controllable pneumatic pushing mechanism proved effective, removable and reusable without damaging the cured assembly. With further refinement, such as implementing automated dip-coating for a consistent layer thickness of liquid gel, this fully pre-cast coupling strategy represents a potential scalable assembly approach for the Gen2-OM, eliminating the labor-intensive backfilling procedure currently required for LOM production.

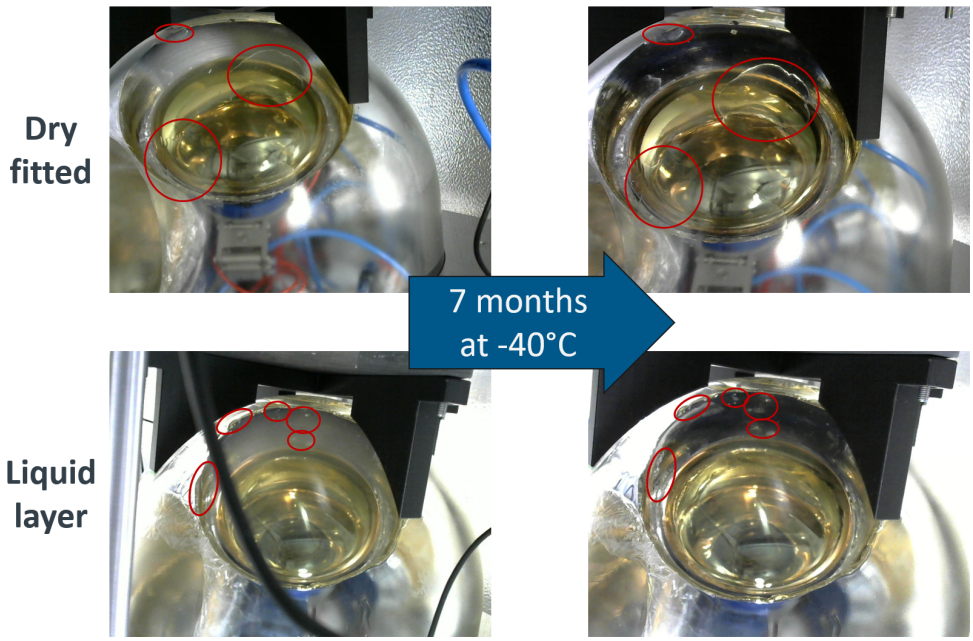


Figure 7.6: (left) Start of the test reaching −40 °C. **(right)** After seven months at −40 °C. The top row shows the dry-fit coupling, with red circles highlighting stagnant air pockets observed during monitoring. The bottom row shows the liquid-layer coupling, with red circles highlighting bubbles on the outer side walls rather than at the glass interface.

7.2 Component Testing

In addition to validating the integrated module assembly, several individual components require independent characterization. The structural integrity of the pressure vessel must be verified against the external hydrostatic pressure induced during the borehole freeze-in phase (see Section 3.4). Additionally, the optical efficiency of the gel pads must be quantified and the internal metallic holding structure used in the LOMs must be proven not to induce anomalous PMT dark rates. While the optical shadowing effects caused by the external deployment harness are addressed separately via simulation in Chapter 8, this section focuses exclusively on the mechanical and localized empirical performance of the internal components.

7.2.1 Pressure Vessel

Two viable manufacturers supply the borosilicate glass pressure vessels: Nautilus Marine Service GmbH (utilized for the LOM16) and Okamoto Glass Co., Ltd. (utilized for the LOM18). While differences in the radiopurity of these glasses inherently impact the optical background rates of the modules (detailed in Chapter 9), this section focuses exclusively on their mechanical deformation under pressure.

To withstand the deep-ice refreezing process, the pressure vessels require a fully reversible pressure rating of 700 bar. Mechanical testing of the LOM16 pressure vessel was conducted at the Nautilus pressure-testing facility. To actively monitor the pressure vessel's horizontal (radial) and vertical compression, a custom deformation logger (adapted and modified from earlier design iterations [116]) was developed and installed inside the pressure vessel, as shown in Figure 7.7. The logger utilizes linear motion potentiometers (LMPs) read out via a Phidgets InterfaceKit, providing an analog voltage (0 V to 5 V)

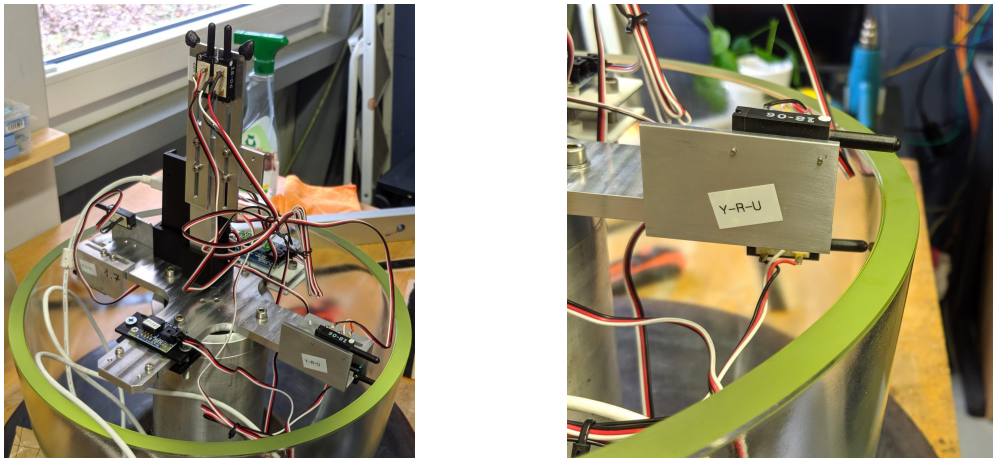


Figure 7.7: (left) Photograph of the deformation logger placed inside a LOM16 pressure vessel. **(right)** Close-up of the internal mounting structure and the LMPs.

corresponding to a 12.5 mm travel distance with micrometer precision. Multiple sensors were mounted along each axis to measure average deformations, mitigating alignment errors. The USB readout was routed entirely through the module's PCA to an external laptop.

The pressure vessel was subjected to a ramp-driven pressure profile, shown in Figure 7.8 (left). The right panel of Figure 7.8 demonstrates that the structural compression is a fully reversible linear process up to the 700 bar target. Linear fits to the data yield a horizontal deformation rate of 5.67×10^{-3} mm/bar and a vertical deformation rate of 3.53×10^{-3} mm/bar.

These measurements were compared against a finite-element (FE) simulation provided by the manufacturer, shown in Figure 7.9. For the horizontal axis, the simulation predicts a compression of 1.78 mm per side, equating to a total diameter reduction of 3.56 mm. This is consistent (within ∼10 %) with the measured 3.93 mm total horizontal deformation, with the minor discrepancy likely attributable to a slight tilt in the internal sensor assembly. Conversely, the vertical axis exhibited a more pronounced deviation. The simulation predicts a total vertical compression of 1.73 mm (0.87 mm per hemisphere), while

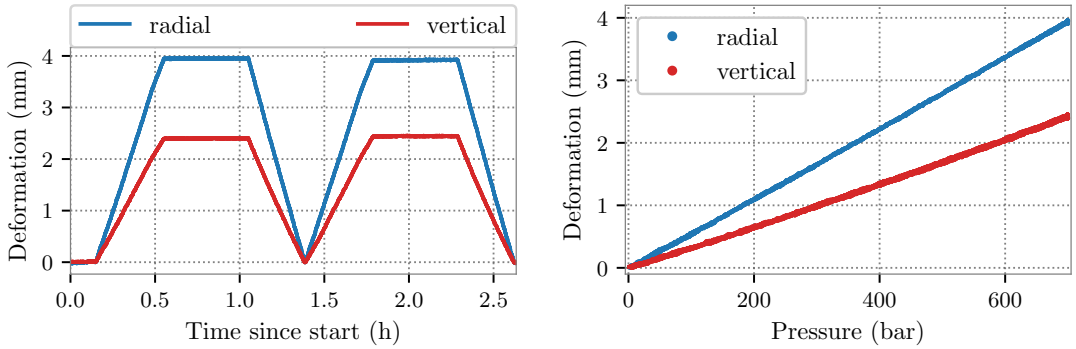


Figure 7.8: (left) Results of the pressure test showing the total structural deformation as a function of time across two subsequent pressure ramps. **(right)** Measured deformation as a function of applied pressure, with measurement points from all ramps overlaid.

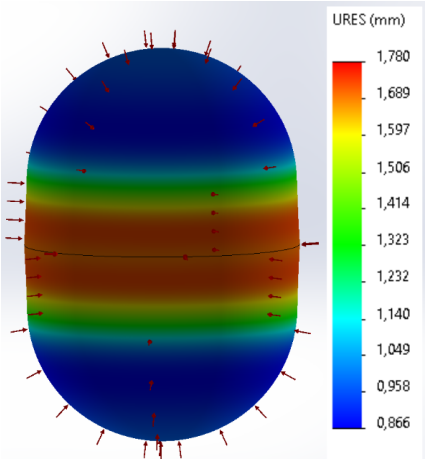


Figure 7.9: Finite-element simulation of the LOM16 pressure vessel under an external hydrostatic pressure of 700 bar, provided by the manufacturer [168].

the empirical measurement recorded 2.42 mm. According to the manufacturer, this underestimation in the idealized FE model likely arises from the glass hemispheres' ground equatorial sealing surfaces; during their first deep pressurization cycle, these microscopic surface irregularities are pressed more firmly together than the continuous material model predicts [168].

Finally, it should be noted that these 700 bar pressure tests were conducted at ambient room temperature, where the temperature in the chamber continuously decreased from 17.5 °C to 15.5 °C over the duration of the test, as measured by temperature sensors inside the pressure vessel. In situ, the optical module faces this hydrostatic load simultaneously with temperatures down to −40°C. While the borosilicate glass exhibits negligible thermal contraction (< 0.1 mm diameter reduction for $\Delta T = -60$ K), the internal components and silicone gel possess significantly higher thermal expansion coefficients. The true in situ stress environment is thus a combination of external pressure driving the pressure vessel ~ 4 mm inward, while thermal contraction simultaneously pulls the internal components away from the glass. Because simultaneous deep-freeze and 700 bar testing of integrated modules is experimentally infeasible, thorough validation relies on combining isolated long-term thermal-vacuum tests (Section 7.1.2) and compression studies, as presented here. Ultimately, the integrity of the optical bond entirely depends on the compliance of the gel pads to elastically absorb these conflicting dimensional shifts.

7.2.2 Gel Pad Efficiency

Beyond verifying the structural integrity and optical clarity of the gel pads, their quantitative impact on the photon detection efficiency must be evaluated. To measure this performance gain, PMTs equipped with pre-cast gel pads (provided by Chiba University from the LOM18 production) are mounted on a custom, high-precision, stepper-motor-controlled rotation table. The central axis of the rotation table is aligned with the center-point of the PMT photocathode. The PMTs are illuminated by a diffuse light source emitting ~ 10 PE at nominal gain and the signal is read out in current mode utilizing a picoammeter. This experimental setup is shown in Figure 7.10 (left), where the methodology is adapted from angular acceptance measurements previously performed for the mDOM reflectors [116 (pp. 166–174)].

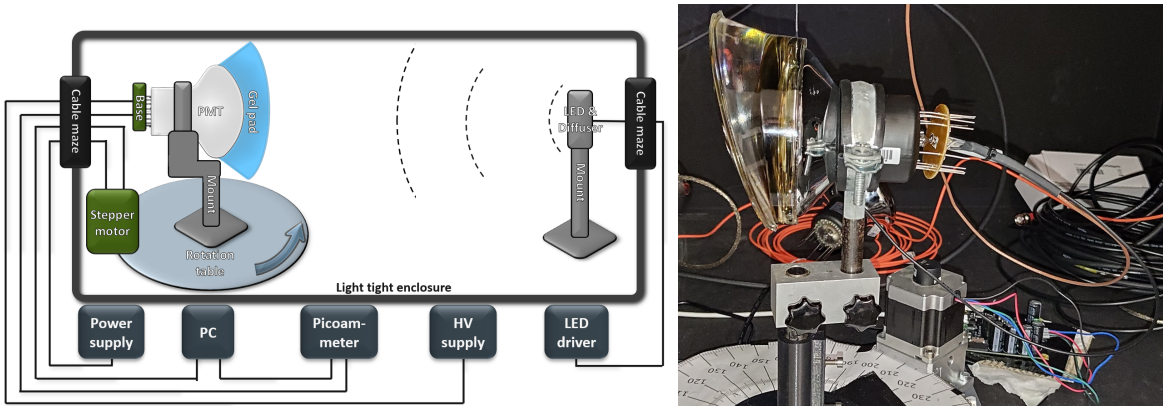


Figure 7.10: (left) Experimental setup for measuring the angular acceptance of a PMT with attached gel pads. **(right)** Photograph of a PMT with the asymmetric polar LOM18 gel pad attached in the upright position within the measurement setup.

Measurements are recorded from $-180°$ to $180°$ in $0.5°$ increments, where $0°$ is defined as the PMT symmetry axis oriented parallel to the incident light. To mitigate systematic errors, the rotation is performed twice per measurement sweep and background dark currents are recorded immediately before and after illumination for baseline subtraction. Each gel pad configuration is characterized three times with identical supply voltage: coupled to the gel pad in both upright and sideways orientations, and finally as a bare PMT after the gel pad's removal. Both partially pre-cast gel pads from the LOM18 assembly feature high-quality TIR surfaces but lack a backfilled center, leaving a slight internal cavity.

The resulting data, plotted as dark-current-corrected relative angular acceptance normalized to the $0°$ response of the respective bare PMT, is shown in Figure 7.11. A photograph showcasing the asymmetric polar gel pad in the measurement setup is shown in Figure 7.10 (right). The gel pads enhance forward acceptance, achieving nearly identical maximum relative values of 1.42 (sideways) and 1.41 (upright) at $0°$ compared with the bare PMT baseline for the polar configuration, as well as 1.34 (sideways) and 1.39 (upright) for the equatorial configuration. However, as the incident angle diverges from the central axis, the two orientations exhibit distinct acceptances towards higher angles. Because the LOM18 pressure vessel is elongated rather than perfectly spherical and the PMTs are tilted with respect to the pressure vessel's nominal surface, the gel pads are mechanically asymmetric to match the varying curvature of the glass. This asymmetry is further compounded by the relative tilt of the PMT axis for the equatorial gel pads. Sweeping across the pad's distinct major and minor axes thus yields differing boundaries for refraction and TIR.

This geometric asymmetry causes the sideways orientation to maintain a broader acceptance profile at wider off-axis angles than the upright orientation, which mathematically accounts for the discrepancy in their integrated overall photon capture rates (increasing by a factor of approximately 1.82 versus 1.67, respectively, for polar and 1.74 versus 1.70 for equatorial). This divergence is attributed to the orientation of the asymmetric pad: the wider, thicker portion of the gel pad (visible at the top in Figure 7.10, right) points towards positive rotation angles during the sideways sweep. As a result, rotation in the negative direction moves the bulk of the gel pad into the incident illumination. This asymmetry culminates in a distinct secondary peak at approximately $-110°$

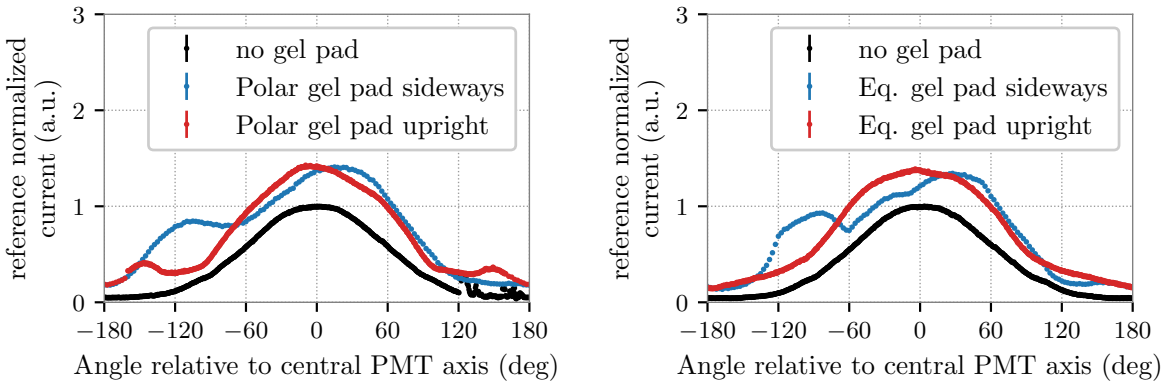


Figure 7.11: Relative angular acceptance of the tested gel pad configurations, normalized to the 0° dark-current-corrected response of the bare PMT. **(left)** Polar LOM18 PMT with gel pad. **(right)** Equatorial LOM18 PMT with gel pad. Uncertainties too small to visualize.

and induces an angular shift of the central maximum towards $\sim 25°$. Conversely, the upright orientation maintains a more symmetric profile, exhibiting less pronounced secondary peaks at approximately $\pm 150°$ for the polar gel pads. Noticeably, the response does not reach zero towards $\pm 180°$ despite the light-tight black enclosure. For the gel pad configurations, this is expected: the projected outer edge of the pad is not completely shadowed by the PMT body, allowing back-incident photons to strike the side of the gel pad and reflect internally onto the photosensitive area (converging at a relative response plateau of ~ 0.18 for all four orientations). Conversely, the bare PMT response drops to 0.045 instead of zero, representing a minor baseline offset attributable to a combination of slightly tilted manual alignment, weak stray reflections or an underestimated dark current. Finally, the underside of the PMT is partially obstructed by the mounting structure, slightly decreasing the response with gel pads for incident angles $\gtrsim |130°|$.

When interpreting these results, a geometric caveat must be considered: this experimental setup evaluates the acceptance in air, effectively representing the optics *inside* the module. In a fully assembled, deployed module, however, gel pads are coupled to the curved glass pressure vessel. This glass refracts incoming photons, fundamentally altering the incident angles at which they strike the gel pad and modifying the transmission probabilities due to a closer refractive index match between the glass and the gel. Thus, while this measurement isolates the gel pad's intrinsic performance and TIR capabilities, the overall effective area of the complete module (accounting for pressure vessel optics and in-ice scattering) is evaluated via dedicated Geant4 simulation studies in Chapter 8. As a point of reference for this scaling, prior independent full-module simulations in [10 (p. 114)] have demonstrated that the inclusion of optical gel pads increases the overall photon detection efficiency by a factor of approximately two compared with a baseline with no gel pads.

7.2.3 Influence of metallic Holding Structure on PMT Noise

The proximity of metallic components at different electrical potentials near a PMT photocathode can induce external electric fields that severely increase dark rates [135]. This effect is particularly pronounced when operating PMTs under negative high voltage (HV), necessitating mitigations such as the external conductive coating applied around the PMT

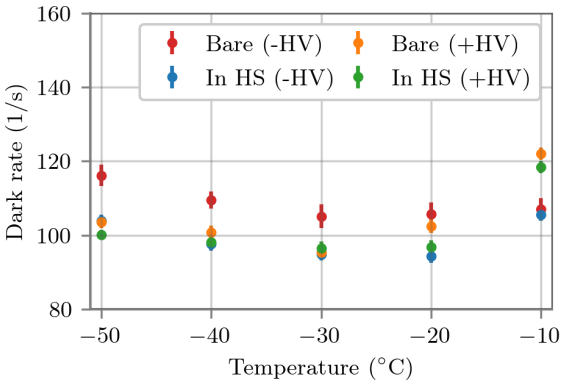


Figure 7.12: PMT dark rate as a function of temperature for a bare 4-inch NNVT PMT and the same PMT mounted within the sheet metal holding structure (HS), measured for both positive and negative HV polarities.

stems. In mDOMs, the surrounding metallic reflectors must be actively maintained at the photocathode potential to prevent additional noise inflation [73 (pp. 163–168)]. The internal support structure of the LOM16 is constructed from simple sheet metal. While it is situated behind the photocathode edge rather than protruding directly into the forward optical plane, its electrical proximity still warrants empirical verification to ensure it does not contribute to the overall module noise.

To quantify this electrical proximity effect, dedicated dark rate measurements were conducted on 4-inch NNVT PMTs (the experimental setup and procedures are detailed in Section 6.2.3). The PMT is operated in pulse mode utilizing a passive base at nominal gain and is evaluated under both positive and negative HV polarities. Measurements are conducted inside a climate chamber across a temperature range of −50°C to +20°C in 10°C increments. The PMT is tested in two configurations: completely bare and mounted within the metallic LOM16 holding structure.

The resulting temperature-dependent dark rates are presented in Figure 7.12. The data reveals no statistically significant or systematic increase in dark rate when the PMT is mounted in the metallic holding structure compared with its bare baseline. While minute deviations (on the order of $\lesssim 15\,\mathrm{s}^{-1}$) are observable, particularly for the bare PMT with negative polarity at low temperatures, these do not indicate field-emission scaling. Instead, they are attributable to minor systematic measurement tolerances: triggering the data acquisition at a fixed 2 mV threshold translates to a slightly different fractional photoelectron threshold, even while the gain is kept at a calibrated nominal value between the distinct setup configurations and temperatures. These measurements therefore demonstrate that the sheet metal holding structure used in LOM16 has no impact on PMT dark rates.

7.3 Summary and current Status

The preceding prototyping and development phase for the Gen2-OM has concluded with the LOMs. A final design review (FDR) of all components, together with design verification testing (DVT) of fully assembled modules, demonstrated the structural and optical integrity of both module designs. Ultimately, five LOM16 and four LOM18 modules were

successfully deployed into the ice as part of IceCube Upgrade. At the time of writing, the integration of these novel multi-PMT modules into the existing IceCube software framework and data processing pipelines is actively ongoing.

Looking toward the future IceCube-Gen2 array, the LOM16 architecture has been established as the baseline design. However, to ensure the feasibility of mass-producing 9600 units, the tight internal space constraints of the LOM16 are relaxed slightly to accommodate scalable assembly methods. Building upon this foundation, the finalized Gen2-OM will therefore adopt the moderately enlarged pressure vessel geometry introduced in Section 7.1.2. Given the years of close collaboration with both glass manufacturers and the consistent success of prior pressure testing campaigns (Section 7.2.1), the procurement and validation of this revised pressure vessel design are considered low-risk.

This modified pressure vessel geometry, however, necessitates a redesign of the internal layout, including the gel pad configuration. This upcoming development phase provides an ideal opportunity to further investigate the scalable assembly concepts developed and tested in this work. To transition from manual prototyping to mass production, future assembly protocols may build on the strategies outlined in Section 7.1.2:

- **Fully pre-cast Gel Pads:** Utilizing scalable thermoforming techniques to cast complete gel pads that are rotationally locked to the PMT symmetry axis. This step offers the greatest reduction in assembly complexity, as it has the potential to completely eliminate the time-consuming liquid backfilling process used during LOM production.
- **Controllable pneumatic Integration:** Employing reusable, pneumatically driven scissor-lift mechanisms to rapidly, precisely and evenly couple the PMT subassemblies to the inner glass surface under simultaneous vacuum conditions.
- **Scalable internal mechanics:** Transitioning from 3D-printed mechanical holding structures to finalized designs optimized for rapid injection molding.

Ultimately, adopting a fully pre-cast optical coupling strategy, supported by dedicated mechanical assembly tooling, provides a practical route toward reproducible, high-throughput assembly required for the unprecedented scale of IceCube-Gen2.

8 Optical Module Performance Simulations

Following the development and mechanical realization of the Gen2-OM precursor modules, the next step is to quantify their expected performance within the detector environment. For accurate event reconstruction, the wavelength- and direction-dependent sensitivity of each optical module must be precisely understood. A central performance metric describing this acceptance is the *effective area*, which provides a quantitative measure of the photon detection capability and serves as the primary benchmark for comparing the acceptance of the Gen2-OM precursor modules. The underlying Geant4 simulation framework and optical module implementations required to model this response are introduced in Section 8.1, followed by the evaluation of the effective area results in Section 8.2.

8.1 Geant4 Simulation Framework

Geant4 is a C++-based Monte Carlo simulation toolkit designed to model the passage of particles through matter [169, 170]. Its applications span various fields of particle physics, including high-energy, nuclear and accelerator physics. Throughout the years, a specialized framework *OMSim* [171] has been developed primarily for the characterization of IceCube's optical modules. This framework builds upon a simulation originally developed for two PhD theses [116, 172], which was later generalized for any optical module. Following contributions from [73], it now incorporates a detailed PMT simulation. Previously, only the mDOM was fully (and the DOM partially) implemented within this framework. In this work, the complete geometries and the necessary parameters for background estimation (see Chapter 9) have been extended to include the D-Egg, LOM16 (building on initial implementations in [173, 174]) and LOM18.

In Geant4, geometries are constructed following a definite hierarchy. First, basic shapes such as cylinders, spheres and cones are defined as *solids*. Each solid is then associated with a *logical volume*, which dictates its material properties and environmental characteristics. Finally, to position the object, the logical volume is instantiated as a *physical volume* within a larger parent volume (e.g., the surrounding ice environment). For effective area simulations, only optical processes governing photon interactions, specifically absorption, boundary processes (refraction and reflection) and Mie scattering, are simulated. These interactions are governed by the assigned material properties derived from empirical measurements.

Optical Module Geometries

For basic estimates, only the glass pressure vessel, optical gel and PMTs are required. However, the deployed modules contain additional geometries, such as harnesses, holding structures and electronics, which can physically *shadow* the photocathodes. Because these components are complex and difficult to implement natively in Geant4, they are imported from CAD files based on the method described in [175]. These CAD geometries are

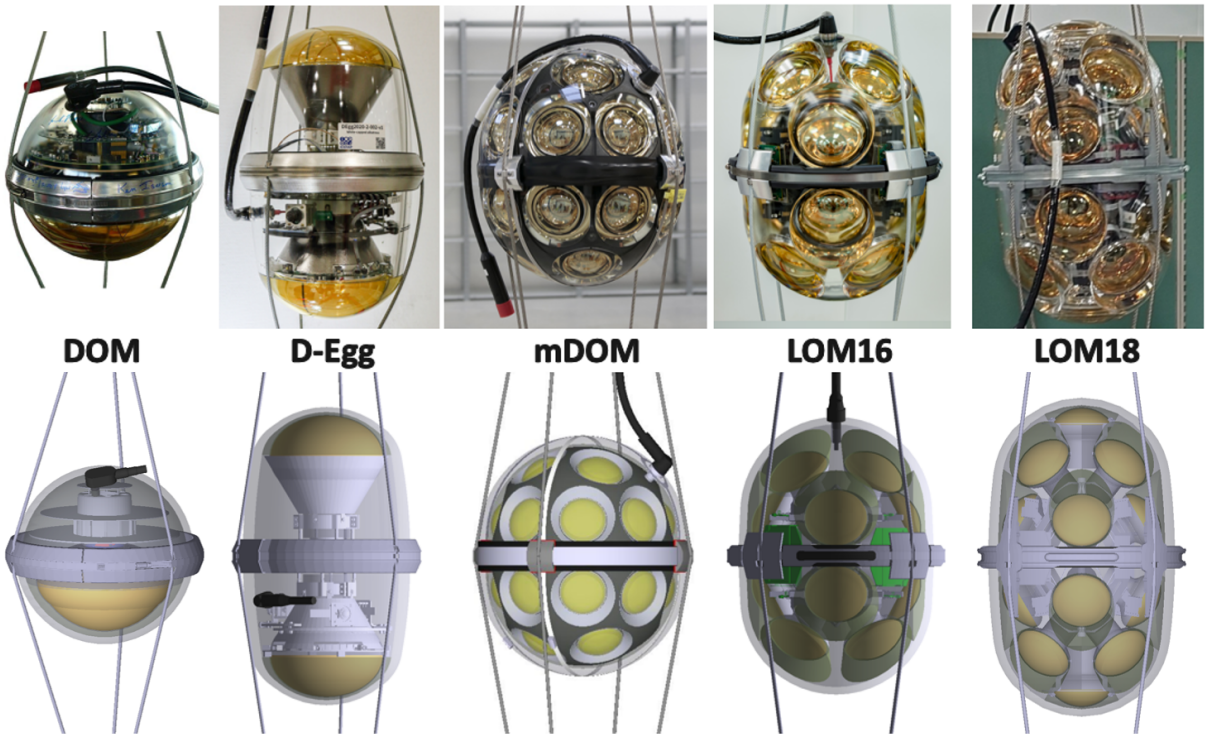


Figure 8.1: Visual comparison between real-world optical modules (top, pictures courtesy of the IceCube collaboration) and their counterparts as implemented within OMSim (bottom).

slightly optimized prior to import[1], assigned realistic material properties and placed into their respective logical volumes. A drawback of CAD-imported geometries is that they are rendered as tessellated objects, which can introduce unwanted boundary effects between volumes. Thus, this import method is restricted strictly to passive (opaque) components. An overview of the implemented module geometries, as introduced in Chapters 3 and 5, is presented in Figure 8.1.

PMT Response Simulation

Once photons reach the PMT geometry, their detection must be modeled as realistically as possible. The key implementations include:

- The PMT glass bulb geometry is generated by fitting lines, circles and ellipses to the manufacturer's technical drawings. Based on these outer dimensions, the borosilicate bulb glass is modeled with a uniform 2 mm thickness.[2]
- The photocathode is modeled as a distinct entity from the PMT glass, possessing unique optical properties. It incorporates thin-layer optics [176], as its physical thickness ($\approx$ 20 nm) is smaller than the wavelengths of the incident photons (270 nm $-$ 750 nm).

[1]Potentially overlapping geometries, such as the inner surface of the harness waistband, are thinned out, and complex details such as mounting holes or threads are removed to reduce compilation time. These modifications do not alter the optical module's response beyond statistical uncertainties. The inclusion and routing of the PCA depend on available CAD files. In total, 19 such CAD geometries are utilized across all optical module implementations.

[2]Although the CT scan of an mDOM reveals slight inhomogeneities in glass thickness, any resulting discrepancies in optical sensitivity are compensated for by the subsequent PMT matching.

- Internal components, such as the front plate[3], can reflect photons that initially pass through the photocathode onto different photocathode locations, as discussed in Section 6.3.
- Each location on the photocathode is assigned a position-dependent detection probability weight based on the empirical measurements from Chapter 6. This calibration process is referred to as *PMT matching*.

For a comprehensive explanation of this methodology, the reader is referred to [73 (pp. 128–143)], where PMT matching was pioneered for the mDOM PMT. In this work, the PMT geometries and accompanying measurement matching procedures have been finalized and extended to all IceCube PMT models. The respective PMT matching encompasses the following steps:

1. The QE measurement from Section 6.3 is recreated within Geant4 to yield the fraction of absorbed photons per simulated wavelength.[4] This establishes a baseline, simulation-native "QE" derived from initial material properties.
2. Weights are then calculated within the simulation to map the simulated absorption fractions directly to the empirical QE measurements for each wavelength. This calibrates the simulated PMT to reproduce the measured QE curve (Figure 6.10). At this stage, the overall QE is accurate, but the simulated photocathode remains uniformly sensitive across its entire surface.
3. To account for local variations in sensitivity across the photocathode, as observed in Figure 6.22, the 2D scan measurement setup is also recreated in the simulation. The relative change in DE for each spatial coordinate is assigned by mapping the simulated absorbed fraction to the measurement values (based on [73 (pp. 88–90)]).[5] Therefore, the simulated PMT accurately reflects empirical, position-dependent efficiency.
4. Finally, the 2D scan results for relative TT, TTS, gain and SPE-resolution (e.g., Figure 6.19) are mapped to their corresponding photocathode coordinates. For dedicated temporal studies utilizing these parameters, the specific rotational orientation of the simulated PMT within the optical module becomes relevant at this point. While the mDOM PMT was mapped across multiple wavelengths (from 460 nm to 640 nm), all other PMTs were scanned exclusively at 459 nm. Although parameters such as TTS and gain exhibit slight non-trivial wavelength dependence [73 (pp. 96–102)], the 459 nm scans serve as the baseline parameterization for these modules in the present simulation.

The optical properties of the internal components (reflective side walls, front plate, metallic strips and the first dynode) were recently remeasured in [177] under the author's supervision. While the updated optical properties are generally consistent with the initial results presented in [73], the slight differences observed are attributed to inter-sample variations between the tested PMTs. After rematching the mDOM PMTs using the

[3]The internal component geometries were obtained via CT scan of a 3-inch mDOM PMT. For larger PMT models, this baseline geometry is scaled accordingly. While not perfectly exact for every model, simulating these internal reflections and weighting them against measurements provides a more accurate model than artificially stopping transmission at the photocathode.

[4]Because the empirical QE values for the D-Egg and the DOM were provided by Chiba University and the University of Wisconsin–Madison respectively, their specific experimental beam geometries are recreated separately in the simulation.

[5]These measurements were conducted at high intensities (~ 50 PE) to avoid cutting out SPE pulses, as IceTray parametrizes SPE behavior downstream.

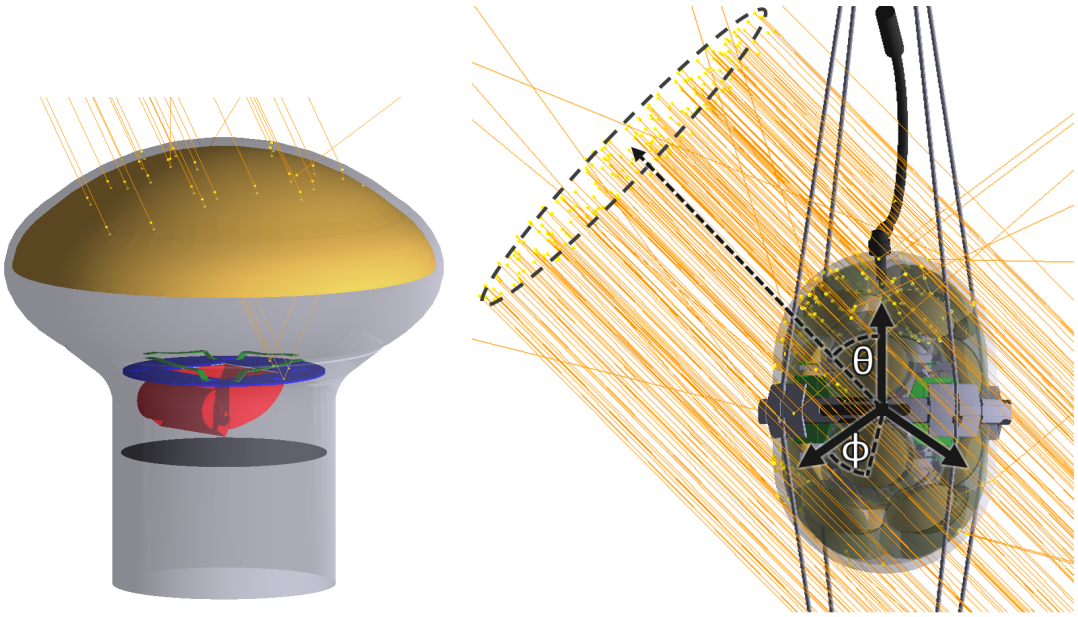


Figure 8.2: (left) Geant4 visualization of a 4-inch Hamamatsu LOM PMT, illustrating possible photon interactions beyond the initial photocathode boundary. The HA-coating and the reflective inner surface at the PMT neck are omitted for visual clarity. **(right)** Geant4 visualization of the effective area simulation for the LOM16 (with harness) for a specific incident direction defined by zenith θ and azimuth ϕ. Orange lines represent photon trajectories; yellow dots indicate interaction vertices.

updated material properties, the effective area exhibited wavelength-dependent deviations compared with the initial baseline results, bounded between $(-1.613 \pm 0.044)\,\%$ and $(2.266 \pm 0.047)\,\%$ across the simulated spectrum.

Figure 8.2 (left) illustrates the various potential photon trajectories after passing through the PMT bulb glass. These paths can result in reflection, transmission or immediate detection at the photocathode, as well as reflection or absorption by internal PMT components. Upon absorption within the photocathode, a photon is assigned a detection probability dictated by the position-dependent empirical mappings.

8.2 Effective Area

The effective area $A_{\text{eff}}(\theta, \phi)$ is conceptually defined as the equivalent cross section of an optical module operating with 100 % detection efficiency. It is calculated for a specific direction as:

$$A_{\text{eff}}(\theta, \phi) = \frac{N_{\text{det}}(\theta, \phi)}{N} \cdot A_{\text{plane}}. \tag{8.1}$$

In this simulation setup, a virtual generation disk with an area of A_{plane} (sized strictly larger than the module's maximum projected area) uniformly emits N mono-energetic photons. Photons are emitted perpendicular to the plane at an incidence angle defined by zenith θ and azimuth ϕ relative to the module center, as visualized in Figure 8.2 (right). The angular directions are distributed using *Healpy* [178], which divides a sphere into equal-area pixels. When a photon strikes the photocathode, it is evaluated against its assigned detection probability, contributing to the total detected sum N_{det} for that specific direction.

To evaluate the overall performance of a module independently of photon arrival direction, the directionally dependent effective area is averaged over the full solid angle. Because the Healpy angular distribution utilizes pixels of equal area, the mean effective area $\langle A_{\text{eff}} \rangle$ of Equation (8.1) is the arithmetic mean of all simulated directions:

$$\langle A_{\text{eff}} \rangle = \frac{1}{4\pi} \int_{4\pi} A_{\text{eff}}(\theta, \phi)\, d\Omega \approx \frac{1}{N_{\text{dir}}} \sum_{i=1}^{N_{\text{dir}}} A_{\text{eff}}(\theta_i, \phi_i), \tag{8.2}$$

where N_{dir} is the number of pixels. This mean effective area serves as the primary baseline metric for evaluating shadowing losses and overall spectral performance.

Because optical properties are wavelength-dependent (primarily driven by the PMT QE and the absorption spectrum of the glass pressure vessel), simulations are run from 270 nm to 700 nm in 10 nm increments. Each optical module is simulated within an environment matching the refractive index of ice, both with and without its mechanical harness. The parameter space covers 3,072 angular pairs, injecting 250 000 photons per angle across the specified wavelength range. For the LOMs, the measured differences in QE between NNVT and Hamamatsu PMTs (see Figure 6.10) necessitate separate performance evaluations. Unless otherwise indicated, figures in this section display optical modules including the mechanical harness and LOMs employing NNVT PMTs to align with the IceCube-Gen2 technical design report [10].

8.2.1 Angular Acceptance

To understand the overall directional response of optical modules, it is necessary to first evaluate the acceptance of individual PMTs. For the DOM, the acceptance is entirely governed by its single, downward-facing 10-inch PMT, which predominantly observes light from below. In contrast, multi-PMT optical modules exhibit complex directional acceptance profiles built from the superimposed responses of several smaller PMTs pointing in various directions. As shown by the Mollweide projection in Figure 8.3, an individual polar 4-inch PMT of the LOM16 contributes a much smaller individual effective area than the DOM's PMT (23.5 % in $\langle A_{\text{eff}} \rangle$), whereas the strategic placement of 16 such PMTs within the optical module covers distinct spatial regions. By summing the acceptance of these individual PMTs, the overall 3D directional response of the entire module is formed.

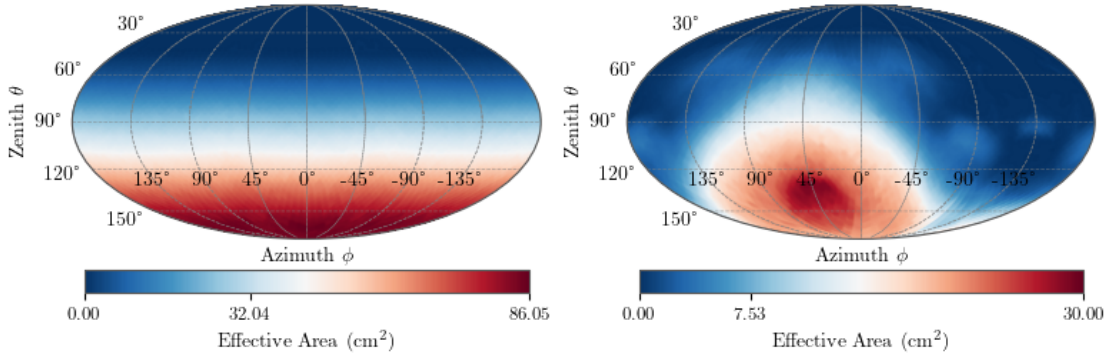


Figure 8.3: Mollweide projection of the directionally dependent individual PMT effective area at 400 nm for the single DOM PMT (left) and an exemplary polar PMT of the LOM16 (right). Red indicates higher sensitivity. Note the different colorbar scales.

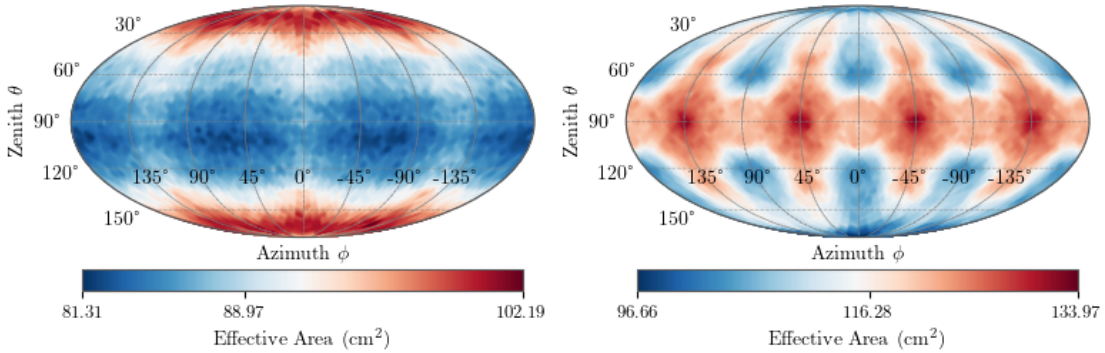


Figure 8.4: Mollweide projection of the direction-dependent effective area across all angle pairs at 400 nm, comparing the D-Egg (left) and the LOM16 equipped with NNVT PMTs (right), including mechanical harnesses. Red indicates larger effective area. Note the different colorbar scales.

Figure 8.4 contrasts the effective area (with mechanical harnesses) of the D-Egg with that of the LOM16 using NNVT PMTs at 400 nm. To quantify directional uniformity, the mean absolute relative difference (MARD), defined as the average spatial deviation from the global mean effective area, is evaluated. Multi-PMT architectures largely eliminate the directional bias of the DOM, which exhibits a MARD of 72.9 %, while the mDOM and D-Egg improve this to 5.3 % and 4.4 %, respectively. This uniform, nearly omnidirectional response fundamentally enhances IceCube's sensitivity, as Cherenkov photons originating from high-energy, through-going neutrinos scatter in the ice and arrive at the modules from all directions. Furthermore, segmenting the photosensitive area into numerous smaller PMTs mitigates local saturation and provides intrinsic directional information, enabling superior event reconstruction. The LOM architecture synthesizes the best aspects of these recent designs: it pairs the PMT segmentation of the mDOM with the narrow borehole geometry of the D-Egg. The LOM16 achieves a MARD of 5.0 %, maintaining tighter spatial uniformity than the LOM18 (10.2 %), while simultaneously exceeding the D-Egg's mean effective area ($\sim 116.3\,\mathrm{cm}^2$ vs. $\sim 89.0\,\mathrm{cm}^2$ at 400 nm).

To investigate the directional acceptance of the individual PMTs within the different optical modules, their effective area $A_{\mathrm{eff}}(\theta, \phi)$ is evaluated as a function of the local photon incidence angle α. Rather than averaging over the global azimuth, the effective area for each individually simulated direction pair (θ, ϕ) is plotted as a discrete point. The local incidence angle α for each incoming wave is derived from the dot product between the normalized incidence vector of the simulated photon plane wave $\vec{v}_{\mathrm{plane}}$ and the orientation normal vector of the respective PMT face $\vec{n}_{\mathrm{PMT}}$, such that $\cos(\alpha) = \vec{v}_{\mathrm{plane}} \cdot \vec{n}_{\mathrm{PMT}}$. Plotting the effective area of a specific PMT against this dot product reveals its intrinsic angular acceptance profile. The vertical spread of the points at a given angle exposes asymmetries introduced by the curved pressure vessel, transitions at the gel boundaries and shadowing from internal structural components (the influence of the external deployment harness is discussed separately in Section 8.2.2). This visualization reflects the non-axisymmetric geometry of the module as a whole, where different incident photon trajectories result in different optical paths.

The results are shown in Figure 8.5, highlighting the wavelength dependence for the baseline DOM, alongside exemplary polar ($\theta = 36°$) and equatorial ($\theta = 62°$) PMTs of the LOM16. For the DOM, which contains a single 10-inch PMT designed primarily to capture up-going photons, the effective area remains relatively constant up to an off-axis

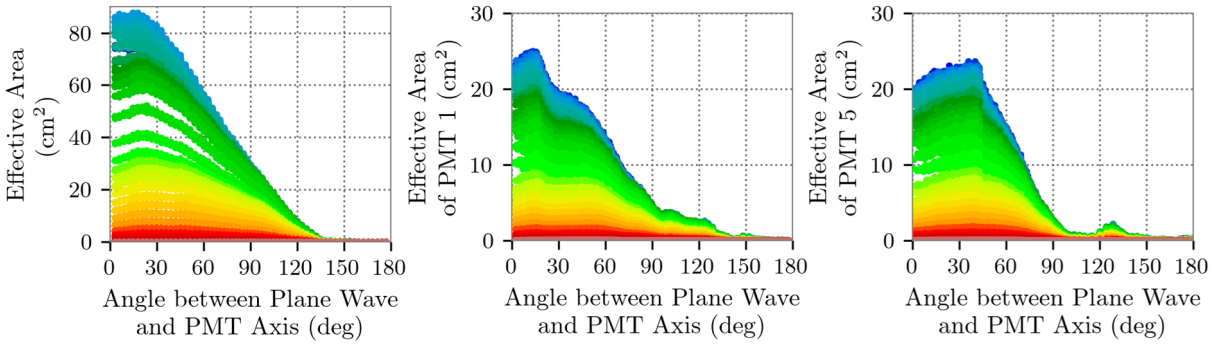


Figure 8.5: Effective area as a function of the incidence angle between the PMT axis and the incident photon plane wave. Different wavelengths are color-coded (e.g., 380 nm in dark blue, 550 nm in green, 650 nm in red) for the DOM **(left)**, a polar PMT **(center)** and an equatorial PMT of the LOM16 **(right)**. Each plot contains 3,072 simulated arrival angle pairs per wavelength. The perceived binning in x-direction results solely from the dot product calculation. Only wavelengths from 380 nm to 690 nm are shown for clarity.

incidence angle of roughly 20°. Beyond this point, it monotonically declines until reaching zero near 130°.

As noted earlier, when analyzing the multi-PMT modules at the individual PMT level, their absolute effective areas are inherently lower due to their smaller photocathode surface areas. The polar PMT of the LOM16 does not peak directly on-axis (0°), but rather exhibits a local maximum in acceptance around 15° off-axis. While the PMT itself and its gel pad (which has an opening angle of 22°) are symmetric around the PMT's central axis, they are mounted on the curved surface of the pressure vessel. Rather than decreasing monotonically with incidence angle, the acceptance curve exhibits distinct structural transitions around 30° and 90°. These features emerge from a complex combination of the pressure vessel curvature, transitions between transmission and total internal reflection (TIR) at the gel pad boundaries and the tilt of the PMT. Notably, photons with incidence angles above 130° can still result in detection. This effect occurs when photons (from the enlarged emission disk) directed toward the module's pole strike the curved pressure vessel at grazing angles; the light is refracted around the curvature of the pressure vessel and effectively guided by the gel pads onto the photocathode.

Similarly, for the equatorial PMTs, the acceptance profile exhibits an off-axis peak between incidence angles of $\sim 18°$ and 37°. This shifted acceptance is mainly driven by the 15° structural tilt of the PMT axis relative to the normal of the glass pressure vessel. Conversely, light incident parallel to the PMT axis strikes the pressure vessel at a slant, suffering higher Fresnel reflection losses. In contrast, light that strikes the pressure vessel normally (maximizing optical transmission) reaches the PMT from an off-axis angle. This tilt also systematically shifts the transition point between direct transmission and total internal reflection at the module's periphery. This optical shift is further compounded by the 22° gel pad opening angle, which preferentially alters the direct on-axis line of sight. Beyond this off-axis peak region, the effective area drops steeply until roughly 100°, after which a secondary acceptance peak is visible around 130°. This local maximum is predominantly caused by photons traversing the internal volume through the gaps between polar PMTs on the same hemisphere and striking the equatorial PMT from behind or the side, similar to the conclusion in Section 7.2.2. An equivalent internal traversal effect is present for the polar PMTs at 150°, though it is significantly less pronounced due to

increased shadowing from internal structural components.

For the remaining optical modules, the angular acceptance profiles of their individual PMTs are largely governed by the specific internal geometries of the modules. For instance, the individual PMTs within the D-Egg exhibit an angular response profile that is nearly identical in shape to the single PMT of the DOM. The mDOM demonstrates a highly homogeneous internal optical environment, resulting in minimal variation between the individual acceptance profiles of its polar (33°) and equatorial (72°) PMTs, especially when compared with the distinct positional differences seen in the LOM16. While the individual PMTs of the LOM18 share a similar overall acceptance profile with those of the LOM16, they present a sharper TIR cutoff. This more pronounced TIR edge is a direct consequence of the LOM18's elongated gel pad design, which extends further from the equator and systematically alters the photon incidence angle at the glass-gel boundary, reaching the critical angle for total internal reflection more abruptly for light arriving at grazing angles.

Furthermore, this angular acceptance exhibits a wavelength dependence visualized by the color separation in Figure 8.5, where intermediate wavelengths near 400 nm dominate the overall acceptance. Longer wavelengths are suppressed by the PMT's intrinsic quantum efficiency, while the shortest UV wavelengths are suppressed by the glass pressure vessel's absorption.

8.2.2 Shadowing Losses

While the angular acceptance is defined by the PMT layout and the glass pressure vessel, deployed modules are obscured by various mechanical structures. It is crucial to evaluate the impact of this mechanical shadowing on the module's baseline sensitivity. To quantify these losses across different viewing angles, the 3D directional data from Figure 8.4 is reduced to a zenith-dependent profile, as the detector's response is mostly symmetric in the azimuthal direction.

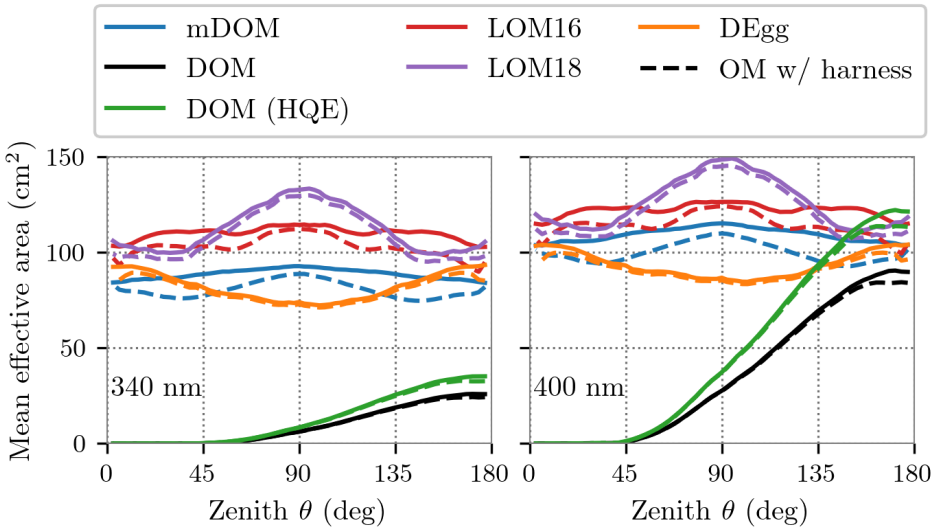


Figure 8.6: Mean effective area (averaged across azimuth angles ϕ) as a function of zenith angle θ for all optical modules, comparing configurations with (dashed lines) and without (solid lines) harnesses. Shown as an example for 340 nm (left) and 400 nm (right) with NNVT PMTs for LOMs.

Figure 8.6 presents two representative wavelength profiles showcasing the zenith acceptance for all five optical module designs. Here, the mDOM and LOM16 maintain the most uniform overall angular response among all modules, reflecting their distributed PMT architectures. In contrast, the D-Egg exhibits its peak sensitivity directly at the poles, whereas the LOM18 shows a maximum when illuminated from the side, attributed to the pronounced elongation of the pressure vessel.

However, introducing the deployment harness into the simulation (represented by the dashed lines) reveals distinct directional losses. Multi-PMT optical modules such as the mDOM and LOM16 experience equatorial shadowing, predominantly caused by the black protective tape and structural waistbands partially obscuring the horizontally facing equatorial PMTs. This creates a visible dip in effective area around the equator for angles where the harness blocks more light, which is referred to as **shadowing**. The mean shadowing loss caused by the harness is 8.5 % for the mDOM and 5.8 % for the LOM16. While the LOM18 also features equatorial PMTs that suffer from this obstruction, its elongated design positions the remaining PMTs further away from the equatorial plane, resulting in a lower overall harness impact of 2.8 %. The purely polar-focused modules, the D-Egg and the DOM, experience the lowest losses at roughly 2 %.

Moreover, deployment challenges introduce deviations from baseline CAD geometries. For example, during the preparation for IceCube Upgrade, the harness of the LOM18 required a late-stage manual rotational adjustment because the structural bridges partially blocked the internal calibration flashers. The exact rotation angle of this adjustment was not precisely constrained, which increased the shadowing of the equatorial PMTs more than originally designed, thereby slightly reducing the effective area values presented for the deployed LOM18.

In addition to the module's immediate harness, the infrastructure of the detector string itself introduces further losses. During this study, the optical shadowing impact of the inter-module harness load cables was tested and found to be negligible (< 0.15 % loss in mean effective area). However, a more important aspect is the main data cable routing (though not applicable to modules in the *deep region* below the main data cable). A preliminary simulation suggests that the main cable, running vertically alongside the modules, can induce an additional 6.2 % to 21.5 % directional shadowing depending on the optical module type (least impactful for the DOM, most impactful for the LOM18), consistent with earlier mDOM studies from [179]. Because the exact geometry in the ice (such as the cable's distance from the pressure vessel, its azimuthal position relative to the module and the module's tilt) is difficult to control, a future systematic parameter space study alongside in situ measurements will be required to account for this within detector systematics.

8.2.3 In-ice spectral Performance

To evaluate the overall sensitivity of the modules across the relevant Cherenkov spectrum, the mean effective area $\langle A_{\text{eff}} \rangle$ integrated over all angles is plotted against wavelength in Figure 8.7. When assuming the QE of the NNVT PMTs (which exhibit a higher but narrower central peak), a clear hierarchy emerges where both LOM variants noticeably outperform IceCube Upgrade's main modules across all wavelengths below 460 nm. If the Hamamatsu QE is applied, the LOMs exhibit performance roughly on par with the mDOMs. As previously discussed in Section 6.3, the true averaged response of deployed

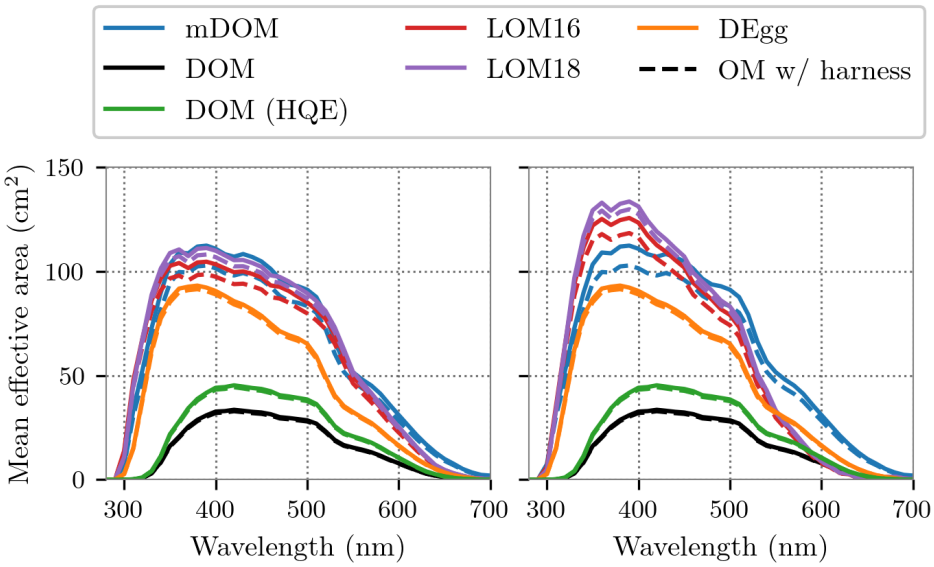


Figure 8.7: Mean effective area (averaged over all angle pairs) plotted against wavelength for all optical modules, with (dashed lines) and without (solid lines) harnesses. The left panel assumes Hamamatsu PMTs for the LOMs, while the right panel assumes NNVT PMTs.

LOMs will likely fall between these two bounds, as the measured NNVT PMTs were high-performers, while the Hamamatsu PMTs fall into the lower end of their respective manufacturing batch.

In a physical detector environment, incident photon wavelengths are not uniformly distributed. While the initial Cherenkov emission spectrum is proportional to $\sim 1/\lambda^2$ (see Equation (3.6)), this spectrum is valid only for interaction vertices in close proximity to the optical module. As photons traverse the ice, wavelength-dependent absorption and scattering heavily distort the spectrum arriving at the detector. Figure 8.8 (left) displays the expected PDFs for surviving wavelengths at various propagation distances, assuming typical deep-glacial ice properties.[6] In the proposed IceCube-Gen2 optical array, strings will be spaced roughly 240 m apart, meaning the straight-line geometric distance between an emission vertex and the nearest optical module will routinely exceed 100 m. Furthermore, because photons undergo multiple scattering events in the bulk ice (with a typical scattering length of $\sim$ 50 m at 400 nm in the physics region [96]), the actual traversed optical path length is longer than the geometric line of sight. Thus, true photon propagation distances of several hundred meters are entirely expected, filtering the surviving Cherenkov spectrum.

Weighting the raw effective areas by these distance-dependent spectra and integrating over all wavelengths yields the true physical performance ratios relative to the standard DOM, as shown in Figure 8.8 (right). All multi-PMT modules demonstrate at least a three-fold improvement over the baseline DOM. Notably, all LOM configurations outperform both IceCube Upgrade's main optical modules across all propagation distances. As discussed in Chapter 6, a high quantum efficiency (HQE) variant of the Hamamatsu

[6]Ice properties vary with deployment depth, primarily dictated by the varying concentrations of included dust. The spectral distributions presented here are calculated using the specific optical scattering and absorption lengths evaluated at a depth of 1698 m, as defined by the South Pole ice (SPICE) model [180]. As introduced in Section 3.6, the deepest *physics region* features highly transparent ice.

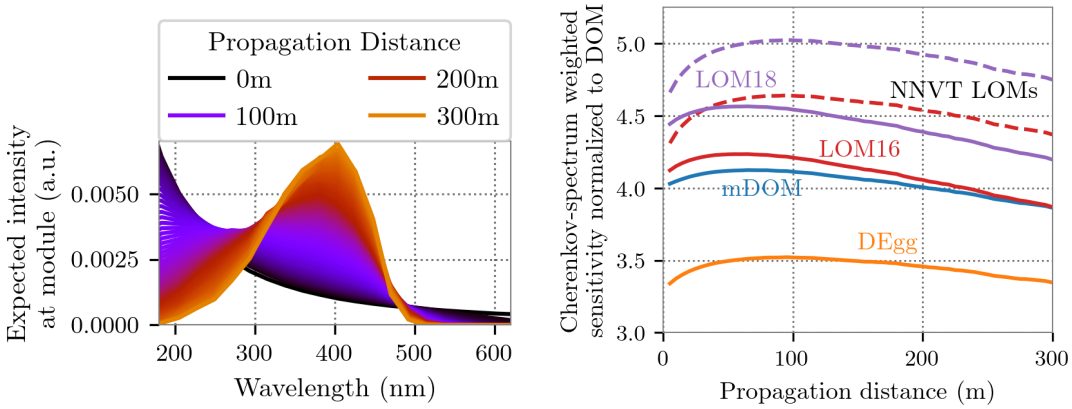


Figure 8.8: (left) Expected spectral PDFs of photons reaching optical modules at varying propagation distances, calculated using the average optical ice properties at a depth of 1698 m in IceCube's *bulk ice*. **(right)** Relative improvement in spectrum-weighted effective area (including harnesses) compared with the DOM as a function of photon propagation distance. Dashed lines indicate NNVT QE for LOMs.

PMTs was also characterized. If maximizing the effective area becomes the primary design driver, integrating these HQE PMTs would yield an expected improvement factor of 5.1 for the LOM16 and 5.5 for the LOM18 at a photon propagation distance of 100 m. However, this performance gain must be weighed against system-level trade-offs, including increased unit costs (which could necessitate deploying fewer modules overall) and elevated background noise rates (see Chapter 9).

Looking forward, the future Gen2-OM, utilizing a slightly larger pressure vessel than the baseline LOM16, provides additional internal volume for geometrical optimization. Consequently, the final Gen2-OM is expected to exceed the geometric baseline performance established here by the LOM16. Ultimately, achieving the optimal detector configuration will require balancing this optical sensitivity with the ongoing engineering challenges detailed in Chapters 5 and 7, such as spatial constraints, minimizing power consumption and optimizing mass production costs.

It should be noted that the effective area values derived in this chapter exhibit slight discrepancies from those presented in the IceCube-Gen2 technical design report [10 (pp. 119–120)]. Those specific values were calculated using DOUMEKI [181], a similar Geant4-based framework developed by the Chiba working group. The primary difference between the simulations is that OMSim includes the fully integrated mechanical geometry (e.g., holding structures and harness ropes), PMT measurement matching native to the simulation and empirically derived optical properties for all individual materials. Despite these methodological differences, the two frameworks demonstrate reasonable agreement, with maximum deviations below 12 % for a given optical module. Collaborative efforts to fully align these simulation frameworks are currently ongoing. Furthermore, discussions are underway regarding whether the direction-and-wavelength coupled results generated by OMSim, with the help of machine learning [182], should replace the current simplified, decoupled parameter tables utilized in IceTray.

8.3 Summary and Outlook

In this work, the complete mechanical geometries (integrating both native Geant4 volumes and CAD-imported internals) and empirical PMT matching procedures have been implemented for all IceCube optical modules, enabling a standardized performance comparison. The simulations incorporate detailed evaluations of mechanical shadowing, individual PMT angular acceptance and spectrum-weighted signal propagation in ice. The results validate the LOM16 as a strong baseline for the Gen2-OM. Depending on photon propagation distance and utilized PMT manufacturer, the LOM16 achieves an improvement in effective area by a factor of 3.8 to 4.3 relative to the DOM, thereby outperforming both main optical modules of the current IceCube Upgrade.

Beyond final performance benchmarking, these simulations played a crucial role during the R&D phase of the LOMs. They directly informed design decisions, including PMT orientation and positioning, as well as the geometric shaping of the gel pads to increase optical yield. Simulating the interplay of these components allowed for maximizing the effective area while accommodating realistic mechanical assembly tolerances. Looking forward, once the physical Gen2-OM mechanical prototype is completed, its geometry must be implemented into OMSim. As large-scale Gen2-OM PMT characterization proceeds, the resulting averaged QE curves and complementary 2D temporal scans must be fed back into the PMT matching pipeline to refine the effective areas. This final R&D cycle will provide an opportunity to leverage this simulation framework for fine-tuning PMT viewing angles, gel pad geometries and internal component layouts prior to mass production.

9 Optical Module Background Estimation

Having established the expected performance of the module designs, a complete characterization also requires an accurate understanding of their intrinsic background contributions. The borosilicate pressure vessels contain trace amounts of radioactive isotopes, primarily ^{40}K and elements within the uranium and thorium decay chains. Energy deposited from these radioactive decays produces Cherenkov and scintillation light, constituting the main intrinsic background for IceCube's optical modules, as detailed in Section 9.1.

The scintillation properties underlying this background are highly material-specific. While previous studies [73, 91, 183], including the author's master's thesis, have extensively characterized the mDOM pressure vessel glass, this thesis refines the scintillation emission spectrum and investigates isotope concentrations in greater detail to improve the accuracy of the background model. Moreover, this study expands the analysis beyond the mDOM to encompass the D-Egg, LOM16 and LOM18.

Section 9.2.1 measures the isotope concentrations across multiple mDOM glass vessels to assess sample-to-sample variance and ensure manufacturing consistency. Subsequently, Section 9.2.2 explores the composition of borosilicate glass by quantifying its raw constituent ingredients (such as quartz sand) to understand the origins of this background contribution. Section 9.3 presents a measurement of the scintillation spectrum as a function of temperature, extending previous room-temperature limited measurements from [183]. In Section 9.4, the scintillation yield is calibrated based on recent production mDOM measurements. Finally, Section 9.5 integrates these updated parameters into a full simulation to determine the expected in-ice background rates for the different optical module designs.

9.1 Underlying Principle

The mechanisms of scintillation are briefly outlined here. A more comprehensive description can be found in [183 (Section 4.3)] as well as in standard literature [184, 185].

Unlike crystalline scintillators, borosilicate glass is an amorphous solid. It lacks a periodic crystal lattice but exhibits short-range atomic order, allowing the application of modified energy band models [186]. It functions as an insulator, featuring a valence band (VB) and a conduction band (CB) separated by a band gap. When radioactive isotopes decay within the glass, the emitted particles (specifically α-particles and secondary electrons) deposit energy into the surrounding material. This deposited energy promotes electrons from the VB to the CB [187]. Upon de-excitation, electrons transition back to the VB. Because direct band-to-band recombination in wide-bandgap amorphous materials is inefficient, luminescence is primarily mediated by localized defect states (luminescence centers) within the band gap [188, 189]. These intermediate energy levels determine the transition energy and, consequently, the wavelength of the emitted photons.

Alternatively, non-radiative de-excitation occurs through interactions with phonons (network vibrations). This competing process leads to thermal quenching [190]. The probability of non-radiative decay increases with temperature T, following an Arrhenius-like temperature dependence proportional to $\exp(-E_a/k_BT)$, where E_a is the activation energy and k_B is the Boltzmann constant. Consequently, the overall scintillation efficiency exhibits an inverse relationship with temperature. To empirically model the background induced by these processes in the pressure vessel, the scintillation is characterized by four primary properties:

Isotope Concentrations: The specific radionuclide concentration in the material, typically measured in $\mathrm{Bq\,kg^{-1}}$ and reported separately for ^{40}K and the ^{235}U, ^{238}U and ^{232}Th decay chains (see Section 9.2).

Scintillation Spectrum: The distribution of emitted photon energies which dictates the wavelength profile of the expected optical background (see Section 9.3).

Scintillation Yield: The number of photons emitted per MeV of deposited energy. This yield depends on the type of interacting particle, as heavy, densely ionizing particles (α particles) and lightly ionizing particles (such as electrons from β decays or γ-interactions) transfer energy to the amorphous network with distinct efficiencies [184, 191] (see Section 9.4).

Scintillation Lifetime: The temporal profile of photon emission following excitation, usually described by a multi-exponential decay function. This parameter defines the timing characteristics of the intrinsic optical noise (see [73 (pp. 174–182)]).

9.2 Isotope Concentrations in the Vessel Glass

In this section, the specific activities for multiple mDOM pressure vessels and raw borosilicate glass components are measured via γ-spectroscopy. This technique relies on the interaction of incident γ-rays with a semiconductor crystal via the photoelectric effect, Compton scattering and pair production. These interactions generate electron-hole pairs in direct proportion to the deposited energy, allowing for the reconstruction of emission spectra to identify and quantify specific radionuclides [192, 193].

The common measurement setup is illustrated in Figure 9.1. Samples are placed inside

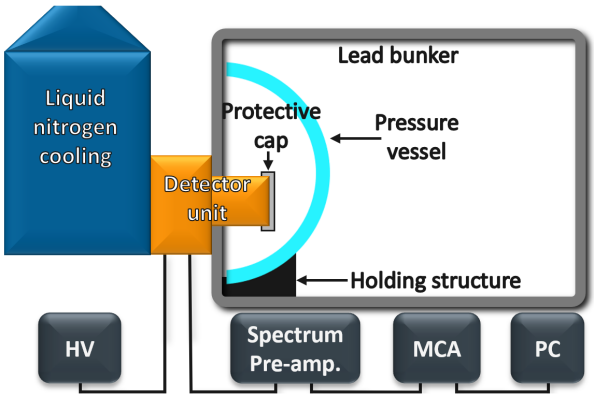


Figure 9.1: Schematic of the γ-spectroscopy measurement setup, shown for a pressure vessel sample.

a lead bunker with a cutout for a high-purity germanium (HPGe) detector[1]. A bias voltage of 3000 V is applied to the detector and its signal is guided into a spectroscopy preamplifier equipped with a pulse shaper[2] before being routed into a multichannel analyzer[3] (MCA). The MCA output is continuously recorded as a histogram with 8192 channels. To obtain sufficient statistics, measurements for each sample are conducted for a minimum of one week.

9.2.1 Concentration in mDOM Vessels

The long-term objective is to monitor the isotope concentrations throughout the Gen2-OM production, as radioactive decays ultimately scale the PMT background rate via glass scintillation. For this thesis, nine mDOM pressure vessels from subsequent production deliveries (comprising top and bottom halves, with top vessels featuring a cylindrical cutout for the PCA) are evaluated. To maximize detection efficiency, the vessels are positioned with the bulk of their mass close to the active detector component. A 3D-printed ramp matching the vessel curvature and featuring a stabilizing lip secures the samples upright. Pressure vessels are referenced against the lead bunker wall and floor to guarantee reproducible positioning across all measurements. For this evaluation, a distinction is made between two backgrounds:

- **Environmental Background:** Constant ambient radiation originating from naturally occurring radioactive materials in the surrounding environment and from cosmic rays.

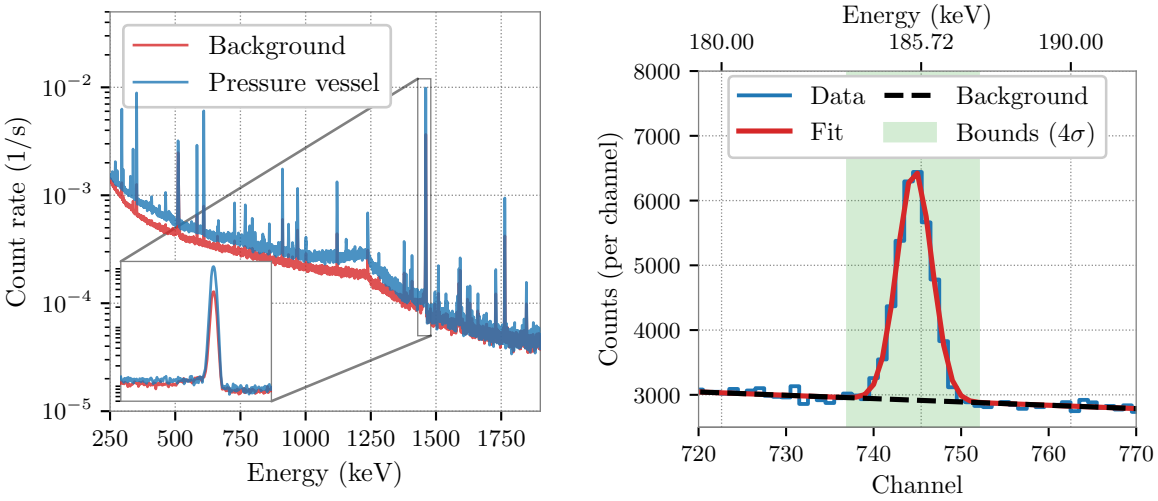


Figure 9.2: **(left)** Measurement spectra of the overall background with only the empty holder (red) and an exemplary mDOM pressure vessel (blue). The prominent ^{40}K FEP is highlighted in the inset. **(right)** Exemplary fitting procedure for the 185.72 keV FEP, showing the dynamically modeled linear background (black dashed), the total fit including a Gaussian (red) and the 4σ integration bounds (green shaded region) used for net count summation.

[1] Canberra GX3019 with a Cryostat Model 7905SL-30 and Model 2002CSL preamplifier, continuously cooled via liquid nitrogen.
[2] CAEN Mod N968
[3] CAEN Mod N957

- **Signal-induced Background:** Artifacts arising from the physical interactions of the sample's emitted radiation with the detector and surrounding shielding, generating partial-energy features such as the Compton continuum, annihilation peaks and backscattering.

An overlaid comparison of a reference background measurement (with only the sample holder present) and an exemplary mDOM pressure vessel is shown in Figure 9.2 (left). This overall background spectrum inherently contains both environmental radiation (manifesting as trace full-energy peaks (FEPs) from ^{40}K and the natural decay chains) and the associated signal-induced background continuum generated by these ambient γ-rays interacting with the detector setup [194]. These constant background components must be subtracted from the sample measurements. As shown, the introduction of the pressure vessel creates an increase in count rates across most FEPs and further elevates the overall continuum due to the sample's intrinsic radioactivity. Channels are converted to energy using a composite calibration source containing ^{60}Co, ^{137}Cs and ^{241}Am. This energy calibration is verified for each measurement using prominent natural peaks, such as ^{40}K (1460.8 keV) and ^{212}Pb (238.6 keV). Hardware limitations of the spectroscopy preamplifier constrain the dynamic range to a maximum of $\sim$ 2000 keV. As a result, high-energy lines, such as the prominent ^{208}Tl peak at 2614 keV (^{232}Th chain) and the ^{214}Bi peak at 2204 keV (^{238}U chain), exceed the dynamic range. However, a sufficient number of lower-energy transitions from each chain remain available to accurately quantify the activities of the respective parent isotopes.

To extract net counts from the raw MCA spectra, a region of interest (ROI) summation is employed (see Figure 9.2, right), where the integration window is constrained around the target FEP energy (4σ of a reference Gaussian). The local signal-induced background beneath the peak is dynamically modeled and subtracted by anchoring a linear fit to the continuum immediately outside the defined ROI. It is emphasized that the FEP itself is not fitted for integration; rather, the total counts are extracted via discrete channel summation of the background-corrected signal within the ROI. Energy conversion

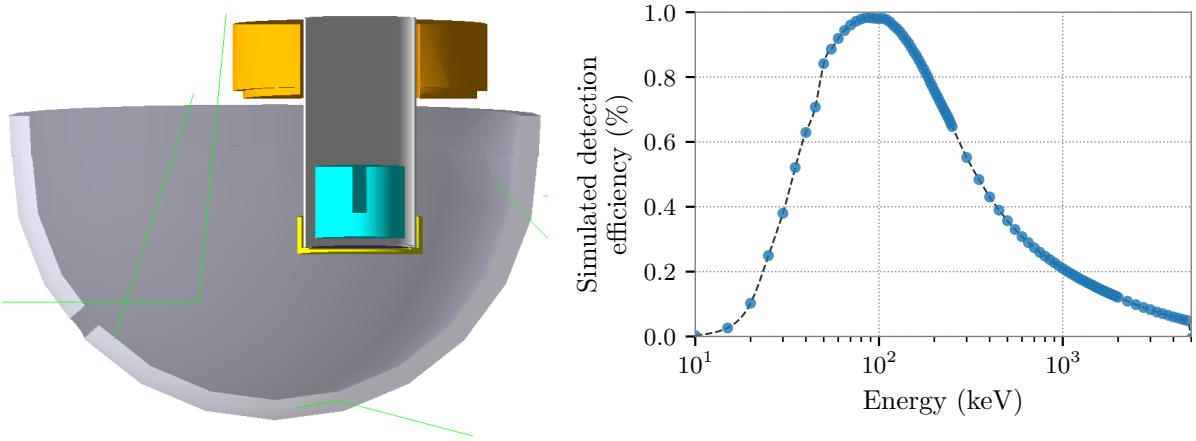


Figure 9.3: (left) Geant4 visualization of the detection efficiency simulation with a longitudinal cut plane. γ-rays (green) are emitted isotropically from the pressure vessel and the deposited energy in the active detector component (blue) is recorded. **(right)** Simulated detection efficiency $\epsilon(E)$ for the pressure vessel geometry (blue). Cubic interpolation (dashed) is used to determine efficiencies at specific energies. The inclusion of the PCA cutout produces no visible change in the simulated efficiency.

is applied solely to map peaks to specific decays and for visualization, where this method further allows for slight calibration drifts between samples. For each FEP, the specific activity A is calculated as

$$A = \frac{N_{\text{sample}}/t_{\text{sample}} - N_{\text{bkg}}/t_{\text{bkg}}}{\epsilon(E) \cdot I_\gamma \cdot m}, \tag{9.1}$$

where N_i are the net FEP counts, t_i the respective measurement livetime, $\epsilon(E)$ the detection efficiency, I_γ the absolute branching ratio and m the sample mass.

To determine $\epsilon(E)$, the detector setup and sample (parameterized with the exact elemental composition and density to model self-absorption) are reconstructed in Geant4 (see Figure 9.3, left). γ-rays with energies ranging from 10 keV to 5000 keV are simulated (10^8 events per energy), emitted isotropically from the confined volume of the sample. The detection efficiency is calculated as the ratio of detected to simulated gammas, $\epsilon(E) = N_{\text{det}}/N_{\text{sim}}$, where N_{det} represents the number of simulated gammas that deposit their full initial energy within the active detector volume. The resulting efficiency curve is shown in Figure 9.3 (right). The overall detection efficiency is low ($\lesssim 1\,\%$) primarily because the large vessel covers a minimal solid angle relative to the detector crystal and induces self-absorption of its own emitted γ-rays. The efficiency spectrum exhibits a distinct peak near 90 keV for the pressure vessel. At lower energies, γ-rays are almost entirely absorbed by the glass bulk and the detector's dead layer via the photoelectric effect. Conversely, at higher energies, the total interaction cross section decreases, causing a larger fraction of photons to traverse the HPGe crystal entirely undetected. Additionally, for the photons that do interact, the probability of Compton scattering overtakes the photoelectric effect, allowing them to escape the detector volume after depositing only a partial amount of their initial energy [192]. Cubic interpolation is applied to determine the efficiency for any specific emission energy.

To ensure the reliability of the calculated activities, Equation (9.1) is evaluated only

Table 9.1: Target isotopes, corresponding full-energy peaks and absolute branching ratios used for the activity calculation [195, 196]. Peaks marked with an asterisk (*) are excluded from the final weighted average calculations for the subsequent components (Section 9.2.2) due to low signal-to-noise ratios or background subtraction artifacts.

Chain	Daughter	Energy (keV)	Branching Ratio (I_γ)	Uncertainty
^{40}K	^{40}K	1460.82	0.1066	± 0.0017
^{232}Th	^{228}Ac	338.32	0.1127	± 0.0019
	^{228}Ac	911.20	0.2580	± 0.0040
	^{228}Ac	968.97*	0.1580	± 0.0030
	^{212}Pb	238.63	0.4360	± 0.0050
	^{208}Tl	583.19*	0.3055	± 0.0012
^{238}U	^{214}Pb	295.22	0.1847	± 0.0011
	^{214}Pb	351.93	0.3572	± 0.0024
	^{214}Bi	609.32	0.4544	± 0.0019
	^{214}Bi	1120.29	0.1490	± 0.0008
	^{214}Bi	1764.49	0.1529	± 0.0008
^{235}U	^{235}U	143.76*	0.1093	± 0.0014
	^{235}U	185.72	0.5720	± 0.0020

for statistically significant transitions. FEPs suffering from low signal-to-noise ratios or severe background-subtraction artifacts are omitted from the final analysis, leaving the decays listed in Table 9.1. Calculating the activity for ^{40}K is straightforward, as it decays to produce a single 1460.82 keV γ-ray ($I_\gamma = 10.66\,\%$). For the ^{232}Th and ^{238}U series, secular equilibrium is assumed, allowing the activity of the parent isotope to be inferred from the measured FEPs of their respective daughter isotopes. In contrast, ^{235}U is quantified directly via its primary 185.72 keV transition ($I_\gamma = 57.2\,\%$) [197]. Within the ^{232}Th chain, the absolute branching ratios I_γ provided in Table 9.1 account for branching probabilities (such as the 35.94 % branch of ^{212}Bi leading to ^{208}Tl), avoiding the need for additional branching-factor corrections. If the activity of a chain is calculated using multiple daughter isotopes or energy lines, an inverse-variance weighted mean is taken to determine the final value. Finally, the half-vessel mass of $m = (6.4 \pm 0.2)\,\mathrm{kg}$ is applied to convert the absolute activity into specific activity. The corresponding results are summarized in Table 9.2.

Comparing these recent production deliveries to previous measurements of older, cylindrical samples from [91] (which were later found to be consistent with a pre-production pressure vessel) reveals an altered overall composition. Variations in the radiopurity of borosilicate glass over long timescales are a known phenomenon. For instance, measurements of Vitrovex vessels produced in the early 2000s exhibited different activity levels compared with contemporary batches [198]. This precedent aligns with the shifts in relative isotope ratios observed in this study, indicating genuine physical changes in the glass formulation over the intervening years. Specifically, compared with the older pre-production samples from [91], the ^{40}K specific activity in the current vessels exhibits a $(23.6 \pm 1.5)\,\%$ decrease, and the ^{238}U and ^{235}U activities drop by $(34.7 \pm 3.8)\,\%$ and $(40.7 \pm 6.1)\,\%$, respectively. In contrast, the ^{232}Th activity shows a $(39.8 \pm 6.3)\,\%$ increase.

As discussed subsequently in Section 9.2.2, variations in the raw material flux or the specific melting furnace used can impact trace-isotope contamination. While the

Table 9.2: Specific activities ($\mathrm{Bq\,kg^{-1}}$) of the evaluated mDOM pressure vessels. Details include whether a vessel has a PCA cutout (top) or not (bottom), alongside the respective production delivery batch. The mean values are compared with the results from [91] of smaller cylindrical pre-production samples for reference.

Vessel sample	**^{40}K**	**^{232}Th**	**^{238}U**	**^{235}U**
1 (bottom, delivery 1)	46.89 ± 1.49	1.81 ± 0.04	3.17 ± 0.05	0.37 ± 0.02
2 (bottom, delivery 1)	46.70 ± 1.48	1.80 ± 0.04	3.18 ± 0.05	0.36 ± 0.02
3 (top, delivery 1)	47.18 ± 1.49	1.84 ± 0.03	3.25 ± 0.05	0.36 ± 0.02
4 (top, delivery 1)	46.63 ± 1.48	1.76 ± 0.03	2.91 ± 0.05	0.34 ± 0.02
Mean (delivery 1)	*46.85 ± 0.25*	*1.80 ± 0.03*	*3.13 ± 0.15*	*0.36 ± 0.01*
5 (top, delivery 2)	46.36 ± 1.48	1.74 ± 0.03	2.77 ± 0.04	0.32 ± 0.02
6 (bottom, delivery 2)	45.16 ± 1.44	1.74 ± 0.03	2.79 ± 0.04	0.32 ± 0.01
7 (bottom, delivery 2)	47.39 ± 1.52	1.84 ± 0.04	2.93 ± 0.05	0.33 ± 0.02
8 (bottom, delivery 2)	46.06 ± 1.47	1.73 ± 0.03	2.90 ± 0.05	0.35 ± 0.02
9 (top, delivery 2)	46.67 ± 1.48	1.81 ± 0.04	3.17 ± 0.05	0.35 ± 0.02
Mean (delivery 2)	*46.33 ± 0.82*	*1.77 ± 0.05*	*2.91 ± 0.16*	*0.33 ± 0.02*
Mean	**46.56 ± 0.62**	**1.79 ± 0.04**	**3.01 ± 0.17**	**0.35 ± 0.02**
Results from [91]	60.98 ± 0.86	1.28 ± 0.05	4.61 ± 0.07	0.59 ± 0.05

sequential deliveries measured for this thesis exhibit consistent activity levels with one another, long-term variations in scintillation yield across different manufacturing years have been measured previously in [183]. This variability is also reflected in the PMT dark rates observed during mass production for mDOMs [155], highlighting the importance of continuous radiopurity monitoring.

9.2.2 Components of borosilicate Glass

As discussed later in Section 9.5, the previously determined pressure vessel activities stem from Vitrovex®, a product line by Nautilus Marine Service GmbH utilized for the mDOM and LOM16. In contrast, the LOM18 and D-Egg modules employ Okamoto glass, which exhibits significantly lower trace-isotope contamination [10 (p. 116)]. To investigate the origin of the elevated radioactive contamination in the Vitrovex vessels, Nautilus provided samples of the raw borosilicate glass ingredients in powder form. The goal is to determine whether a specific constituent acts as the primary source of contamination or if the radioactive contaminants are introduced later in the production chain. Each evaluated component serves a specific structural or thermal purpose in the final glass matrix [199, 200]:

- **Quartz Sand (SiO_2):** The primary glass former and backbone of the silica network.
- **Boric Acid (H_3BO_3) and Borax Pentahydrate ($Na_2B_4O_7 \cdot 5\,H_2O$):** These compounds are the sources of boron trioxide (B_2O_3), the secondary glass former that defines the material as borosilicate. The inclusion of boron lowers the melting point of the silica network and is responsible for the characteristic low coefficient of thermal expansion.
- **Aluminum Hydroxide ($Al(OH)_3$):** Provides aluminum oxide (Al_2O_3), which acts as a stabilizer. It enhances the macroscopic hardness, durability and chemical resistance of the glass.
- **Sodium Chloride (NaCl):** Evaluated as a potential sodium-bearing additive. Sodium-containing compounds act as fluxes that lower the melting temperature and viscosity of the melt and also serve as fining agents that assist in the removal of gas bubbles.
- **Potassium Nitrate (KNO_3):** When added to the glass batch, it acts as a flux to lower the melting point and enhance the glass's chemical and thermal shock resistance.

The (background) measurement, simulation and fitting procedures described in Section 9.2.1 are repeated for these raw powder samples. Sample volumes of $\approx 810\,\mathrm{cm}^3$ are packed into a custom 3D-printed container, with masses between 0.43 kg to 0.67 kg, and positioned facing the detector endcap (see Figure 9.4, left). The mass composition, along with specific densities, is simulated as before[4], resulting in $\epsilon(E)$ as shown in Figure 9.4 (right). The evaluated specific activities as per Equation (9.1) for each component across all significant decay channels are summarized in Table 9.3.

[4]Because the samples are loose powders, utilizing standard macroscopic crystal densities in the Geant4 simulation would overestimate the self-absorption of emitted γ-rays by $\sim 20\,\%$. To accurately model the detection efficiency $\epsilon(E)$, the effective bulk density for each component was determined by dividing the respective measured sample mass by the simulated (filled) container volume.

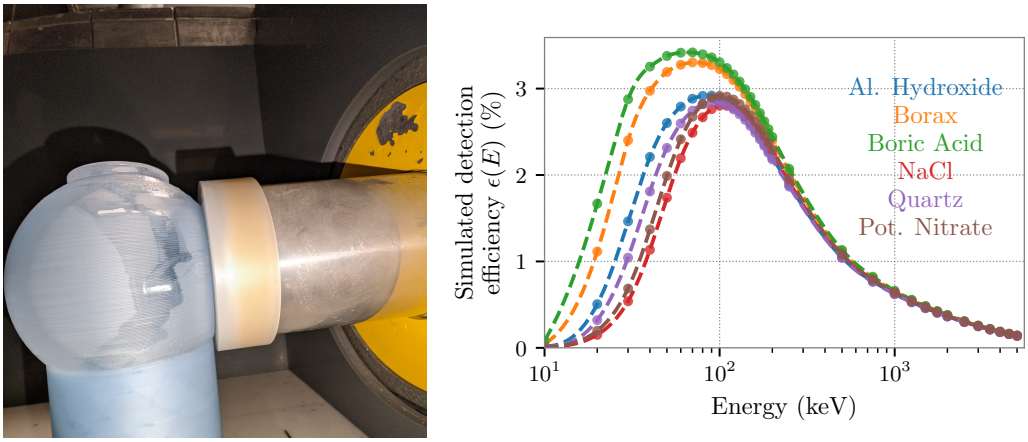


Figure 9.4: (left) Photograph of the raw powder components packed in a custom container positioned in front of the HPGe detector. **(right)** Simulated detection efficiency $\epsilon(E)$ for the respective components. Cubic interpolation (dashed) is used to determine efficiencies at specific energies.

Given the low intrinsic radioactivity of most raw materials, a decision threshold L_C is applied to the net signal to differentiate between significant sample activity and background fluctuations. Assuming that the background counting statistics dominate the net signal fluctuations, the net counting rate for an empty sample is modeled as following a Gaussian distribution centered at zero. According to the ISO 11929 standard [201], the threshold is defined as $L_C = k \cdot \sigma_0$, where σ_0 is the standard uncertainty of the net signal at zero true activity. Using $k = 1.645$ sets a 95 % confidence level, limiting the probability of a false-positive detection to 5 %. For measurements where the evaluated activity A falls below L_C, results are reported as a 95 % confidence level upper limit $L_W = \max(0, A) + k \cdot \sigma_A$, where non-physical negative net counts resulting from background subtraction are floored at zero.

Table 9.3: Specific activity ($\mathrm{Bq\,kg^{-1}}$) of the raw components and the resulting theoretical glass composite calculated via stoichiometric mass fractions. Activities for individual components are determined using an inverse-variance weighted average of characteristic gamma peaks, with background subtraction and powder-density efficiency corrections applied. For FEPs where the activity was below the 95 % confidence level decision threshold, results are reported as upper detection limits.

Component	**^{40}K**	**^{232}Th**	**^{238}U**	**^{235}U**
Quartz sand	< 1.57	2.32 ± 0.06	2.70 ± 0.06	0.35 ± 0.04
Sodium chloride	12.71 ± 0.62	< 0.05	< 0.05	< 0.04
Aluminum hydroxide	< 1.42	0.10 ± 0.04	0.11 ± 0.04	< 0.10
Borax pentahydrate	< 1.62	< 0.11	< 0.08	< 0.06
Boric acid	< 1.47	< 0.07	< 0.07	< 0.05
Potassium nitrate	11788 ± 104	< 0.15	0.22 ± 0.08	< 0.10
Theoretical composite	$\mathbf{50.64 \pm 0.45}$	$\mathbf{1.87 \pm 0.05}$	$\mathbf{2.18 \pm 0.05}$	$\mathbf{0.28 \pm 0.03}$

The dominant ^{40}K peak of potassium nitrate provides an excellent opportunity to verify the specific activity calculations via the theoretical relationship:

$$A = \frac{\ln(2)}{T_{1/2}} \cdot \left(\frac{1000}{M_{\mathrm{KNO_3}}} \right) \cdot N_A \cdot h \cdot p. \tag{9.2}$$

In this equation, A is the specific total activity of the potassium nitrate sample (Bq/kg), $T_{1/2}$ is the physical half-life of ^{40}K in seconds, $M_{\mathrm{KNO_3}}$ is the molar mass of potassium nitrate (101.1 g/mol), N_A is the Avogadro constant, h defines the natural isotopic abundance of ^{40}K (0.0117 % [202]) and p stands for the chemical purity of the powder (0.99 ± 0.01 [203]). Equation (9.2) predicts a theoretical specific activity of $(12.11 \pm 0.13)\,\mathrm{kBq\,kg^{-1}}$, where the measured value of $(11.79 \pm 0.10)\,\mathrm{kBq\,kg^{-1}}$ deviates by only (2.65 ± 1.33) % from this prediction. This close agreement between the measured and predicted activities provides an independent validation that the Geant4 efficiency calibration (including the internal γ-ray self-absorption of the sample) and the overall signal-extraction methodology within this section are modeled accurately.

To contextualize these validated results, a baseline Vitrovex formulation is considered. Although the exact batch composition is proprietary, the manufacturer specifies the final glass network composition as 80.6 % SiO_2, 12.5 % B_2O_3, 4.2 % Na_2O/K_2O, 2.2 % Al_2O_3 and 0.5 % trace elements [168]. Based on the supplied raw-material powders, a stoichiometric[5] pre-melt powder composition is reconstructed to reproduce these final oxide fractions and estimate the radioactive contributions of the individual raw materials. Because the bulk of the required Na_2O is supplied by the borax pentahydrate, sodium chloride is assumed to act only as a trace fining agent. A trace 0.2 % K_2O contribution is assumed within the manufacturer's combined Na_2O/K_2O specification, together with a 0.2 % NaCl fining addition, thereby defining the pre-melt salt allocation. This stoichiometric approximation yields an estimated pre-melt composition of 73.5 % quartz sand, 17.1 % borax pentahydrate, 5.7 % boric acid, 3.1 % aluminum hydroxide, 0.4 % potassium nitrate and 0.2 % sodium chloride.[6]

By weighting the measured specific activities of the raw powders by these calculated pre-melt mass fractions, a theoretical activity baseline for the finished glass network is approximated. Comparing this theoretical composition (Table 9.3) to the measured mean activity of the finished vessels (Table 9.2) highlights the complex dynamics of the vitrification process. The experimental data suggests that while the boron-bearing compounds and stabilizers are radiopure (often falling below the decision threshold), the intrinsic activities are distinctly split: the heavy decay chains are predominantly introduced via the quartz sand, while the ^{40}K activity is driven by the potassium nitrate flux. For potassium, the theoretical composite predicts a ^{40}K activity of $(50.64 \pm 0.45)\,\mathrm{Bq\,kg^{-1}}$, which is close to the $(46.56 \pm 0.62)\,\mathrm{Bq\,kg^{-1}}$ average measured in the pressure vessels solely due to the assumed trace amount of 0.2 % K_2O. As explicitly noted by the manufacturer,

[5] Stoichiometry is the calculation of the amounts of chemical substances involved in a reaction, enabling the determination of reactant consumption and product formation.

[6] Because raw powders lose mass via outgassing (e.g., water vapor) during thermal decomposition, this calculation predicts that producing 1 kg of finished glass requires approximately 1.1 kg of raw powder, which is consistent with literature [199]. For example, forming one mole of Al_2O_3 within the glass requires two moles of $Al(OH)_3$ raw powder. In this model, the borax pentahydrate mass is determined from the required Na_2O content, which simultaneously fixes its contribution to the B_2O_3 fraction. The remaining B_2O_3 requirement is then supplied by boric acid. The theoretical composite activities are weighted by these estimated pre-melt powder mass fractions rather than the provided final oxide percentages.

several glass constituents (including potassium) are volatile at melting temperatures and partially vaporize [168]. The manufacturer typically over-provisions KNO_3 in the initial batch mixture to ensure the final glass network meets the required thermal and structural specifications [168]. The actual raw powder batch utilized in the manufacturing process likely contained less potassium than the idealized baseline considered here.

For the heavy decay chains, comparing the theoretical composite to the pressure vessels reveals distinct behavior. The ^{232}Th activity aligns well ($(1.87 \pm 0.05)\,\mathrm{Bq\,kg^{-1}}$ predicted compared with $(1.79 \pm 0.04)\,\mathrm{Bq\,kg^{-1}}$ measured), indicating that the quartz sand serves as the source of thorium contamination. Conversely, the predicted ^{238}U amount of $(2.18 \pm 0.05)\,\mathrm{Bq\,kg^{-1}}$ underestimates the $(3.01 \pm 0.17)\,\mathrm{Bq\,kg^{-1}}$ measured in the final production vessels. This excess in the ^{238}U chain suggests that the manufacturing environment may introduce additional contamination. This hypothesis is further supported by the isotopic ratio $^{238}\mathrm{U}/^{232}\mathrm{Th}$, which increases from approximately 1.17 in the evaluated quartz sand to 1.68 in the final pressure vessels. A primary suspect for this shift is the leaching of trace isotopes from the refractory firebricks used to line the continuous melting furnaces [204, 205]. These bricks, often composed of zirconia-silica or high-alumina materials, contain elevated intrinsic levels of uranium (frequently exceeding thorium due to zircon crystal preferences) that are incorporated into the glass as the corrosive molten silicate slowly erodes the furnace lining [199]. This leaching phenomenon aligns with previously observed jumps in PMT dark rates when manufacturers transitioned between different melting furnaces for mass production [155 (p. 147)].

While ^{40}K exhibits the highest absolute activity in the glass, full simulations (see Section 9.5) show that it accounts for approximately 35 % of the vessel-induced PMT background. In contrast, the ^{238}U and ^{232}Th series generate background rates of 38 % and 21 %, respectively, despite lower absolute activity levels due to their more complex decay physics. Unlike the single-transition decay of ^{40}K, these full decay series involve multiple sequential α and β decays with higher total energy deposition and particle multiplicity, leading to a much higher probability of producing scintillation photons. Modifying the recipe to minimize ^{40}K to improve the overall radiopurity is possible [168], but the heavy chains remain the dominant background driver.

Ultimately, quartz sand and potassium nitrate are identified as the primary sources of thorium/uranium-chain and ^{40}K contamination in Vitrovex pressure vessels. However, the required volatile over-provisioning and potential furnace leaching underscore the complexity of achieving absolute radiopurity in industrial borosilicate glass.

9.3 Scintillation Spectrum

This section details the measurement of the emission spectrum of the scintillation light induced by radioactive decays. A dedicated experimental setup is designed to isolate the scintillation spectrum using narrow-bandpass filters. Subsequent corrections for optical transmission and detector efficiency are applied to extract the final wavelength distribution across various temperatures.

9.3.1 Experimental Setup

To measure the scintillation spectrum, a radioactive source irradiates the sample to induce a photon emission rate higher than the intrinsic low-intensity scintillation of the glass. An

open ^{241}Am α-source ($A = 15\,\mathrm{MBq}$) is utilized for this purpose. The Cherenkov threshold (see Equation (3.4)) for α-particles in borosilicate glass is approximately 1.3 GeV, which is orders of magnitude above the 5.49 MeV decay energy of ^{241}Am. This ensures that the observed light is predominantly produced by scintillation. Because α-particles have a range of only a few centimeters in air, the source and sample are housed in a custom vacuum chamber[7] to ensure sufficient energy deposition into the glass sample. The high activity source provides sufficient decays to overcome the low signal-to-noise ratio observed in the previous measurement [183] while maintaining temperature stability and safe vacuum handling. The α-source irradiates the sample face oriented toward the PMT. The source housing is positioned slightly off-axis to minimize shadowing of the optical path. To minimize spectral distortions arising from self-absorption during internal reflections, the sides and back of the glass sample are wrapped in absorptive black tape. The custom holders and vacuum chamber are illustrated in Figure 9.5.

Once scintillation light is produced at the sample surface, it exits the chamber through a UV-transmitting sapphire viewport[8]. The light is then detected by a UV-capable, low-dark-noise PMT[9] positioned as close as possible to the viewport to maximize geometric efficiency. To isolate specific emission wavelengths, a custom filter wheel is deployed containing five 50 mm UV-bandpass filters[10]. These filters possess a FWHM of 10 nm with central wavelengths (CWL) at 280 nm, 310 nm, 337 nm, 365 nm and 405 nm. This selection is based on the preliminary scintillation spectrum recorded in [183 (p. 32)]. A solid, blocked position is maintained to measure the background and verify baseline rates. The filter wheel utilizes a Geneva drive to precisely align each filter between the flat PMT

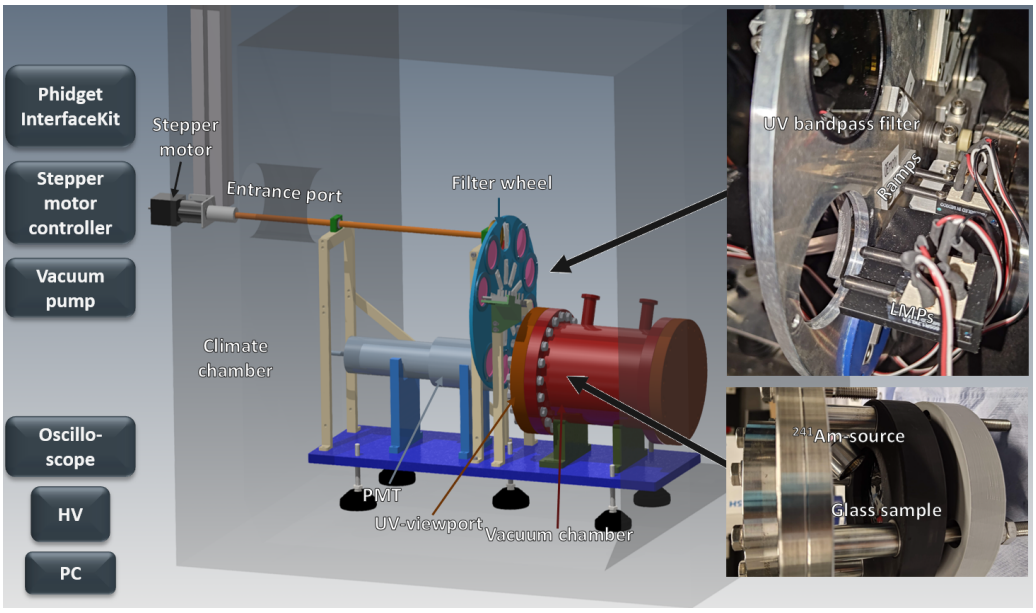


Figure 9.5: Experimental setup to measure the scintillation spectrum inside a climate chamber. Cabling for the PMT high voltage and signal, analog sensor lines and vacuum hoses routed through the chamber entrance port to the measurement equipment (listed on the left) are omitted for visual clarity.

[7] using a Leybold ECODRY 25 plus vacuum pump.
[8] VAB Sapphire CF63
[9] Hamamatsu R11410-20
[10] LC-HBP series, LASER COMPONENTS Germany GmbH

photocathode entry window and the viewport. It is actuated by a stepper motor[11], housed outside the climate chamber. To determine the absolute wheel position, reference ramps trigger linear motion potentiometers (LMPs, as utilized in Section 7.2.1), instructing the motor to stop when the corresponding LMP voltage drops. Two sets of LMPs provide redundancy for each filter position, alongside a dedicated pair exclusively for the blocked reference position.

The PMT is read out analogously to the methodology in Section 6.2.3 using pulse-mode MP-readout, calibrated to a nominal gain of 5×10^6 using positive high voltage. The nominal voltage scales between 1557 V and 1549 V across the temperature range of 20 °C to −50 °C, consistent with the trend seen in Section 6.2.1. Measurements are conducted at four distinct temperatures (20 °C, 10 °C, −10 °C, −20 °C)[12] across all six filter wheel positions. PMT pulses are recorded continuously for three days per configuration. A 2.5 mV trigger amplitude and a 0.2 PE post-cleaning charge cut are applied to isolate genuine PMT pulses from electronic noise. Wavelength-specific rates are subsequently extracted by subtracting the blocked reference position rate, which encompasses the intrinsic PMT noise as well as any scintillation light generated in the PMT glass by incident γ-rays following the ^{241}Am decays.

9.3.2 Corrections

The measured pulse rates require wavelength-dependent corrections. While the geometric efficiency remains constant across all configurations, the optical transmission of the components and the intrinsic PMT characteristics vary significantly with wavelength and, to some extent, with temperature. Tracing the scintillation light from emission to detection requires corrections for the viewport transmission, bandpass filter transmission, PMT QE and PMT afterpulsing probability. Because the vacuum inside the housing chamber is maintained at $\mathcal{O}(10^{-6}$ bar), variable atmospheric absorption of the α-particles is eliminated, ensuring a stable relative spectrum.

Transmission

The viewport transmission was previously characterized by the author in [183 (p. 31)] and the bandpass filters were measured in [177 (pp. 16–19)] under the author's supervision. Both components were evaluated using collimated light from a xenon lamp[13] emitting 200 nm to 1200 nm. The viewport was analyzed using a monochromator[14] coupled to a reference diode[15] read out via a picoammeter[16]. Measurements were recorded both with the viewport in the beam path and when removed for background normalization. The filters were evaluated using an integrating sphere[17], optically coupled via a fiber[18]

[11] Phidgets 3329_0 driven by Bipolar HC 1067_0B

[12] The bandpass filters are the only components not rated for operation below −20 °C. Tests following the main measurement with a single filter at lower temperatures resulted in mechanical cracking, as the metal housing contracts faster than the glass filter substrate.

[13] LOT LSE140 / 160.25C

[14] LOT-QuantumDesign MSH-300

[15] Newport 818-UV

[16] KEITHLEY 6482

[17] Newport 819C-IS-5.3

[18] Thorlabs M114L01

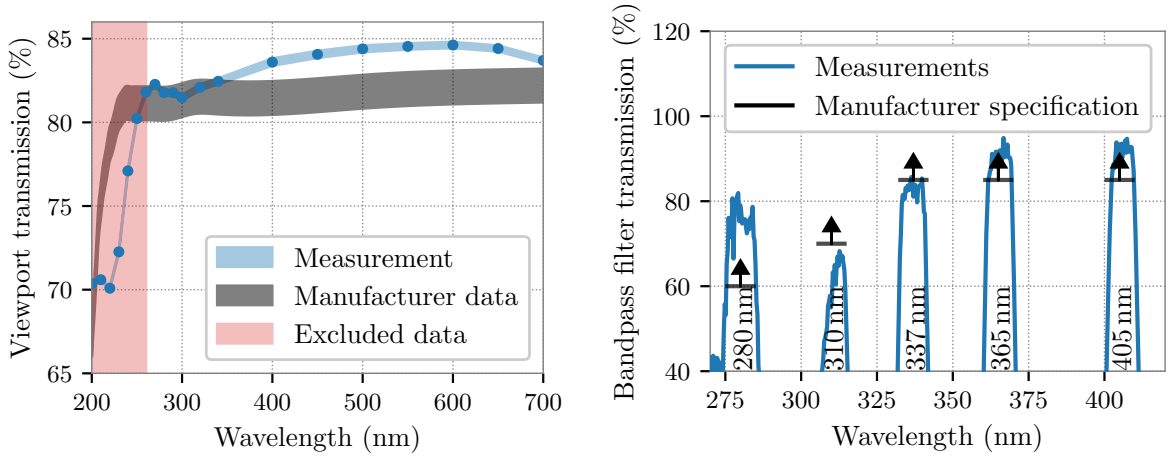


Figure 9.6: Transmission profiles of the optical components. **(left)** Viewport transmission adapted from [183 (p. 31)]. **(right)** Bandpass filter transmission adapted from [177 (pp. 16–19)].

to a spectrograph[19] equipped with a CCD[20], allowing simultaneous intensity recording across all wavelengths. The optical transmission is defined as the ratio of the sample measurement to the unobstructed reference beam. Minor deviations arising from multiple internal Fresnel reflections within the thin filters are negligible [206]. The results for both components are presented in Figure 9.6.

The measured transmission profiles generally agree with the manufacturer's specifications. The 280 nm filter exhibits slightly higher transmission than specified, whereas the 310 nm filter falls below expectations. The measured FWHM values remain consistent throughout. While temperature-dependent transmission could not be measured with the available hardware, literature suggests that spectral shifts in the CWL for bandpass filters over a $\Delta T = 40\,^{\circ}\mathrm{C}$ range are restricted to less than 0.1 nm [207], which is negligible compared with the 10 nm FWHM. Manufacturer studies of similar filters confirm no measurable degradation in transmission efficiency across this temperature span [208]. Furthermore, the optical properties of sapphire viewports remain stable over moderate temperature variations, as the refractive index of sapphire exhibits only small temperature dependence in the visible wavelength range [209]. Environmental humidity is controlled to prevent moisture accumulation during the thermal cycles. The system is rapidly cooled from room temperature to $-20\,^{\circ}\mathrm{C}$ and subsequently stepped upwards, deliberately bypassing $0\,^{\circ}\mathrm{C}$ to avert condensation. Additionally, the hard-coated optical components mitigate moisture-related degradation.

PMT Quantum Efficiency

The PMT QE is measured analogously to the procedures detailed in Section 6.3 and is shown in Figure 9.7. The experimental setup is only viable for wavelengths above approximately 280 nm due to atmospheric absorption of UV photons by oxygen [210]. Manufacturer data for this specific PMT is only available between 165 nm and 200 nm, where the provided measurement was conducted under vacuum conditions [211]. To construct

[19] Andor Solis Kymera 193i-B1
[20] Andor Solis Newton DU920P-OE

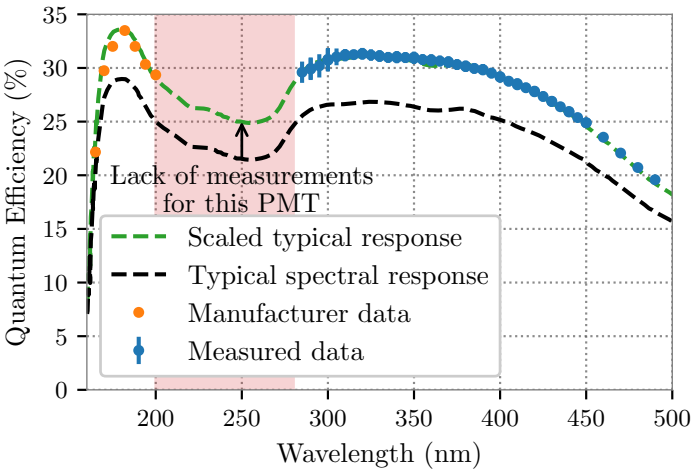


Figure 9.7: Quantum efficiency of the Hamamatsu R11410-20. A data gap (red region) exists between the laboratory measurements (blue) and the manufacturer's vacuum-UV data for this specific PMT (orange). The typical generic response curve [211] (black) is scaled to match both datasets, yielding the continuous QE profile (green).

a continuous profile across the entire emission spectrum, a scaling procedure is applied. The generic spectral response curve provided by the manufacturer for this PMT model is scaled (1.16511 ± 0.00097) to match both the experimental and the manufacturer's vacuum-UV data, generating a response curve across the full spectrum.

The quantum efficiency is also slightly temperature-dependent. Because direct low-temperature vacuum-UV measurements are unfeasible with the current setup, assumptions are derived from literature. A uniform slope of $-0.03\,\%\,°C^{-1}$ is reported for the similar R11410-10 PMT between 250 nm and 400 nm [212]. This aligns with findings for the mDOM PMTs, which exhibit a $-0.03\,\%\,°C^{-1}$ shift in detection efficiency between 350 nm and 420 nm [157] (co-supervised by the author). Thus, cooling the PMT from 20 °C to −20 °C enhances the quantum efficiency by approximately 1.2 %, a correction applied uniformly during the analysis.

PMT Triggering and Afterpulsing

When applying an amplitude-based trigger, recorded rates can be artificially skewed if the PMT gain fluctuates across temperatures. Analytical modeling of the secondary electron distribution demonstrates that signal loss from the hardware trigger threshold is strictly below 1 % at nominal gain [183 (pp. 37–39)]. Thus, complex threshold-loss corrections are deemed unnecessary. Instead, the 2.5 mV hardware threshold (see Figure 9.8, left) is mapped to a consistent 0.4 PE amplitude cut in software (alongside the 0.2 PE charge cut). While this slightly reduces the total accepted signal population, it guarantees threshold consistency across all temperature profiles (with an average SPE pulse height of ~ 7.5 mV).

The measured rates must also be corrected for afterpulsing, which artificially inflates the hit count. The raw rates are multiplied by a factor of $1 - P_{\mathrm{API+II}}(T)$, based on the probabilities shown in Figure 9.8 (right). Unlike the timing parameters discussed in Section 6.2.2, the afterpulsing probability increases below −30 °C. However, the probability remains stable within the primary measurement interval of 20 °C to −20 °C. Ultimately,

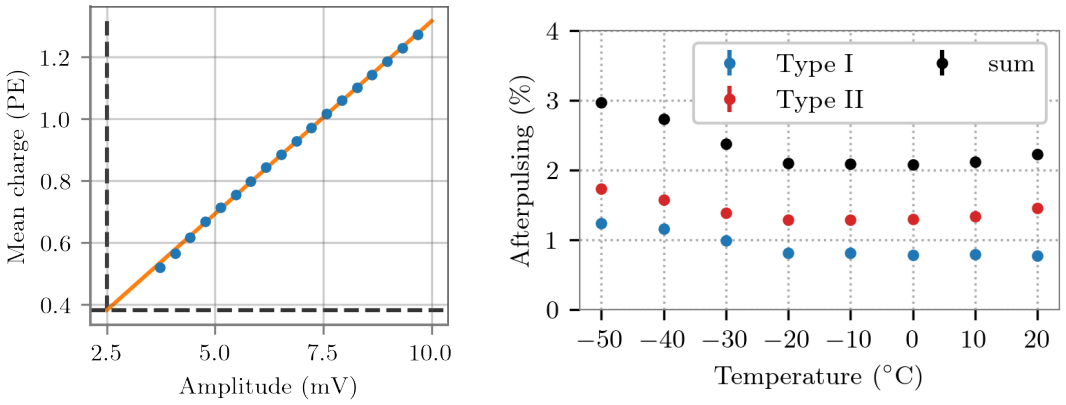


Figure 9.8: (left) Amplitude trigger threshold converted to PE equivalent. Shown for the 20 °C background measurement, where the 2.5 mV trigger corresponds to 0.38 PE. **(right)** Total afterpulsing probability (AP1 + AP2) for the Hamamatsu R11410-20 as a function of temperature.

these combined signal corrections are minor across temperatures and do not induce deformations in the relative spectral shape.

The wavelength- and temperature-dependent corrections for viewport and bandpass filter transmission, PMT QE and PMT afterpulsing are applied to the measured PMT rates to extract the true emission intensity $I_{\text{corr}}(\lambda, t)$ as per

$$I_{\text{corr}}(\lambda, t) = \frac{(R_{\text{raw}}(\lambda, t) - R_{\text{bkg}}(\lambda, t)) \cdot (1 - P_{\text{API+II}}(t))}{T_{\text{view}}(\lambda) \cdot T_{\text{filter}}(\lambda) \cdot \text{QE}(\lambda, t)}, \tag{9.3}$$

where R_{raw} and R_{bkg} denote the measured signal and reference background rates, respectively, and T denotes the optical transmission of the specific hardware component alongside the previously introduced QE and $P_{\text{API+II}}$.

9.3.3 Corrected Spectrum

Before normalizing the spectra, the impact of the applied corrections and the thermal environment on the data is shown in Figure 9.9 (left). The decline in PMT QE towards the UV regime, combined with lower filter transmission, suppresses the raw signals. Additionally, Figure 9.9 (right) highlights the temperature and wavelength dependence of the signal-to-noise ratio (SNR). Because intrinsic PMT dark noise drops while the scintillation yield simultaneously increases towards lower temperatures, the SNR improves significantly for colder measurements.

Normalizing the corrected spectra from Figure 9.9 (left) across all temperatures yields the reconstructed spectra shown in Figure 9.10 (left). The normalized spectral shape remains largely unchanged across all measured temperatures. This temperature independence supports the applicability of the previous room-temperature spectral parameterizations within the simulation framework (see Section 9.5). The emission spectrum of borosilicate glass is inherently broadened by its amorphous structure. Unlike a crystal, where all luminescent centers experience an identical environment, the glass network contains local variations in atomic arrangement and defect configurations. These variations

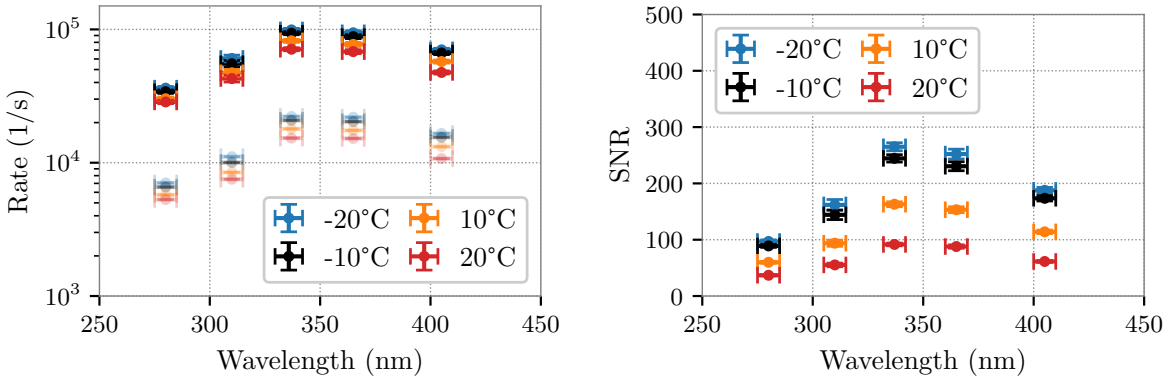


Figure 9.9: (left) Background corrected PMT rates before (transparent) and after (solid) applying wavelength-dependent corrections from Equation (9.3). **(right)** SNR of the corrected signal rates for different temperatures.

lead to a distribution of emission energies and dominate the observed spectral width [213]. Within the investigated temperature range, this static structural disorder is expected to be much larger than any temperature-dependent changes of the emission linewidth. Temperature variations mainly affect non-radiative relaxation processes and interactions with trapping centers, thereby changing the scintillation yield [184]. However, the local environment of the luminescent centers remains sufficiently stable that the relative emission spectrum is preserved. This provides a physically consistent explanation for the invariant spectral shape observed here.

Figure 9.10 (right) reveals a broader spectral profile than previously mapped, with the normalized intensity remaining notably high at 405 nm (~ 0.7 compared with ~ 0.25 from the measurement in [183]). Because this extended red tail diverges from the prior baseline, its physical validity requires verification. The previous measurement was limited

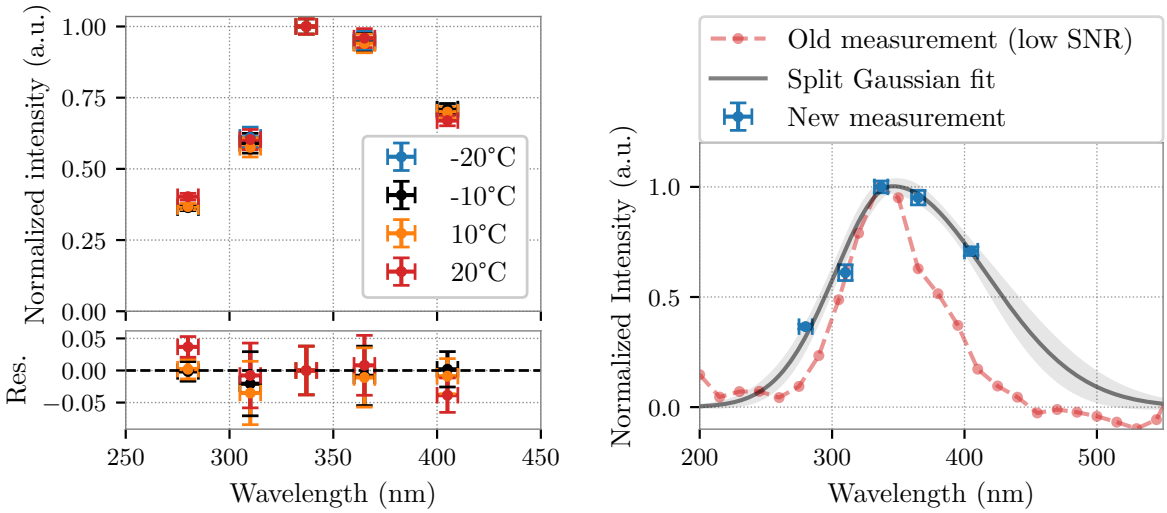


Figure 9.10: (left) Corrected and normalized scintillation spectra for all evaluated temperatures, with residuals relative to the −20 °C spectrum. **(right)** The −20 °C spectrum (blue points) fitted with a split Gaussian model (solid line). The shaded region represents the 1σ fit uncertainty. The previous measurement (red dashed) from [183] is overlaid for reference.

by a very low SNR (<1.1 compared to $\mathcal{O}(100)$ in this work). The current high-statistics evaluation uses the same Vitrovex glass sample and confirms this broadened profile across multiple independent measurement runs and varied temperatures. To exclude sample-specific anomalies or localized instrumental artifacts, a subsequent validation measurement was conducted using a different Vitrovex sample. Despite operating with a reduced four-filter configuration (necessitated by the thermal fracture of one filter during low-temperature testing), this independent run reproduced the broadened spectral shape. Thus, this higher-statistics measurement suggests that the extended visible emission is a genuine material property rather than an experimental artifact, establishing a new baseline for the background simulation.

To parameterize the asymmetric emission profile for the simulation framework, a split Gaussian[21] is fitted to the data using orthogonal distance regression, accounting for uncertainties in both wavelength and measured rate:

$$f(\lambda) = A \cdot \exp\left(-\frac{(\lambda-\mu)^2}{2\sigma^2}\right) \quad \text{where } \sigma = \begin{cases} \sigma_L & \lambda < \mu \\ \sigma_R & \lambda \geq \mu \end{cases}. \tag{9.4}$$

The fit yields a peak at $\mu = (345.8 \pm 13.9)\,\text{nm}$ with width parameters $\sigma_L = (43.5 \pm 12.0)\,\text{nm}$ and $\sigma_R = (70.5 \pm 23.5)\,\text{nm}$. This distinct width asymmetry ($> 50\,\%$) invalidates symmetric standard Gaussian models and suggests a rapid UV onset followed by a slower decay extending into the visible spectrum.

While the core emission peak is well-resolved, the exact boundary conditions of the UV and visible tails carry higher uncertainty due to the sparse sampling of the five bandpass filters. To mathematically constrain these spectral tails and quantify the uncertainty, a

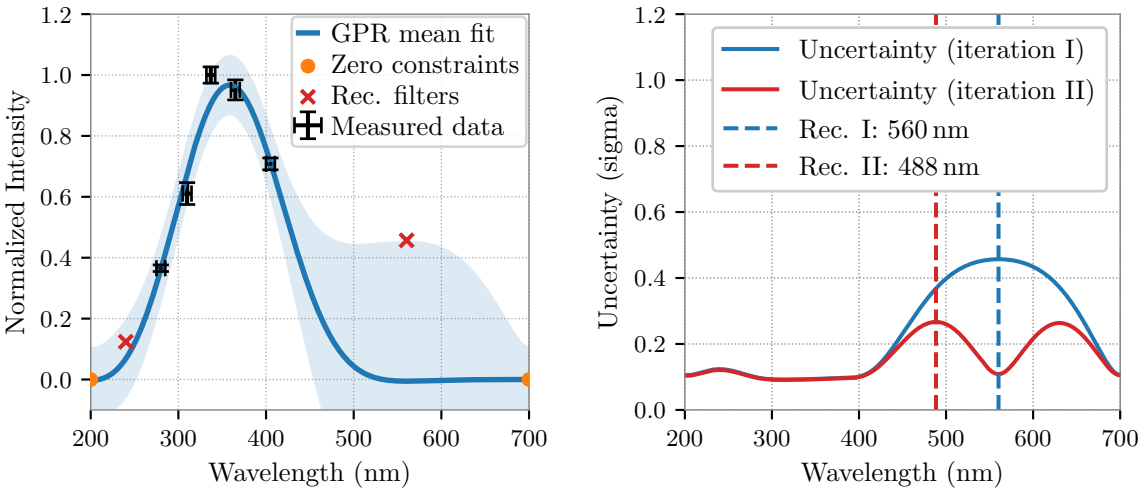


Figure 9.11: Bayesian optimization for filter recommendation. **(left)** Iteration 1 identifies the global uncertainty maxima at the spectral edges (240 nm and 560 nm). The error band represents the predictive model uncertainty σ. **(right)** Iteration 2 sequentially targets the right tail (488 nm) after theoretically constraining the red-tail endpoint.

[21] Although luminescence bands are often approximately symmetric in energy space, the nonlinear Jacobian transformation to wavelength space ($I(\lambda) = I(E)\frac{hc}{\lambda^2}$) introduces an asymmetry that typically broadens the red side of the emission band [214]. The split Gaussian model accurately parameterizes this wavelength-dependent asymmetry.

Gaussian process regression (GPR) is applied. A radial basis function kernel ($\ell \approx 40\,\text{nm}$) models the spectral smoothness, anchored by absolute physical cutoffs at $I(200\,\text{nm}) = 0$ and $I(700\,\text{nm}) = 0$. By evaluating the model's posterior standard deviation, this Bayesian framework identifies the highest remaining uncertainties at 240 nm, 488 nm and 560 nm (see Figure 9.11). These wavelengths define the regions where additional narrow-band measurements would provide the highest information gain for the reconstructed emission spectrum.

The temperature independence should also be experimentally confirmed for Okamoto glass. Since this study demonstrated the thermal stability of the Vitrovex emission spectrum, the previous measurement setup from [183] could be repurposed, as utilizing the higher-activity ^{241}Am source from this work would overcome the prior SNR limitations. Using this setup, a monochromator would replace the discrete wavelength sampling of the filter wheel assembly. Prior to this, its UV transmission range would need to be verified to ensure that the relevant spectral region, including the UV cutoff, remains accessible. Although this investigation fell outside the timeframe of this thesis, it remains a recommendation for future characterization.

9.4 Scintillation Yield

With the specific isotope concentrations (Section 9.2) and the emission spectrum (Section 9.3) of the Vitrovex glass established, the final parameter required for the optical model is the absolute scintillation yield (photons per MeV). Historically, the temperature-dependent scintillation yield of Vitrovex glass was determined using small cylindrical samples in [91] and the author's master's thesis [183], establishing a linear temperature parameterization (detailed for reproducibility in Appendix D). While these measurements capture the intrinsic material scintillation yield, translating this yield to predict background rates in a massive optical module would require precise modeling of the specific sample surface finishes and necessitate identical material properties between the measured sample and the complete pressure vessels.

To avoid these systematic uncertainties and streamline the simulation for background predictions, this work extracts the effective scintillation yield directly from integrated module-level background data. In the simulation framework (detailed subsequently in Section 9.5), the absolute scintillation yield acts as a direct linear scalar for the number of scintillation photons generated per radioactive decay. By matching the simulated hit rates to empirical mDOM target rates from [215], module-specific optical boundary conditions are inherently accounted for. This provides a systematically corrected baseline that maximizes the predictive accuracy for unmeasured or upcoming optical module architectures utilizing the same glass.

Because the measurements from [215] encompass the total integrated module background, the specific simulation target rate stemming solely from pressure vessel scintillation $R_{\text{target,vessel}}$ must be isolated. This is achieved by subtracting the intrinsic PMT dark noise R_{dark} and afterpulsing probability $P_{\text{API+II}}$:

$$R_{\text{target,vessel}}(T) = \frac{R_{\text{total}}(T) - R_{\text{dark}}(T)/f_{\text{gel}}}{1 + P_{\text{API+II}}(T)}, \tag{9.5}$$

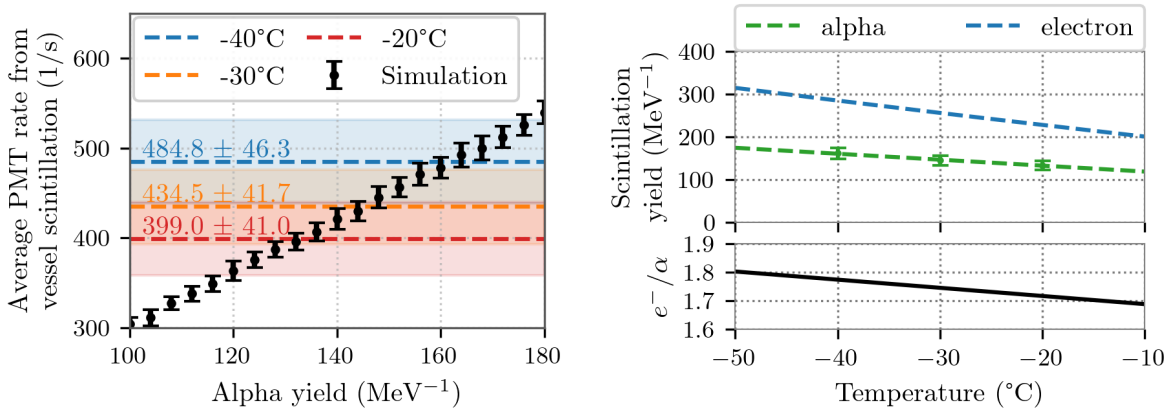


Figure 9.12: (left) Calibration of the scintillation yield. The black error bars represent the simulated PMT hit rates, averaged across the three evaluated temperatures for visual clarity, which are intersected with the empirical target bands for −20 °C, −30 °C and −40 °C derived from [215]. **(right)** Resulting parameterization of the absolute scintillation yields, incorporating the temperature-dependent e^-/α scaling established in [183].

where $P_{\text{API+II}}(T)$ is the afterpulsing probability (see Section 6.2.2), $R_{\text{dark}}(-20°\text{C}) = (340 \pm 15)\,\text{s}^{-1}$ is the intrinsic PMT dark rate measured in air at −20 °C [155 (pp. 147–148)] and $f_{\text{gel}} = 1.725 \pm 0.021$ is a correction factor to account for the decrease in internal reflections when a PMT is optically coupled via gel, as determined in [73 (p. 191)]. Moreover, an artificial dead time of 10 ns is applied between individual simulated PMT detections to replicate the data acquisition of the measurement.

When modeling scintillation photon generation, a distinction must be made based on the excitation mechanism, as densely ionizing α-particles and lightly ionizing electrons (particularly from γ-interactions) exhibit distinct scintillation efficiencies (see Section 9.1). To account for this, an electron-to-alpha yield ratio, as previously determined for Vitrovex glass in the author's master's thesis [183], incorporating a slight temperature dependence, is maintained.

As shown in Figure 9.12 (left), complete module background simulations of radioactive decays within the pressure vessel are conducted across a grid of scintillation yields. The scintillation yield for each particle type is obtained by matching the empirical target rate at each specific temperature and then extracted via a linear intersection. Because the target measurements encompass a relatively narrow thermal window (−40 °C to −20 °C), a first-order (linear) parameterization is applied to these discrete calibration points to establish a continuous baseline down to −50 °C, as presented in Figure 9.12 (right). By anchoring this parameterization to module-integrated data, the scintillation model is calibrated to forecast in situ background rates for Vitrovex vessels.

9.5 Simulation for expected in-ice Background Rates

The Geant4 framework, as introduced in Section 8.1, is utilized to simulate the expected PMT background rates in optical modules following deployment. Rather than generating optical photons directly, the simulation initializes full radioactive decay kinematics for ^{40}K as well as the ^{232}Th, ^{235}U and ^{238}U series isotropically within the pressure vessels. These simulations are executed for both air and ice environments across a temperature range from $-10\,°C$ to $-50\,°C$ for the various multi-PMT optical modules (see Figure 8.1). Ten minutes of real time are simulated per configuration, which corresponds to a set number of isotopes derived from their specific activities within the vessel, as determined in Section 9.2.1.

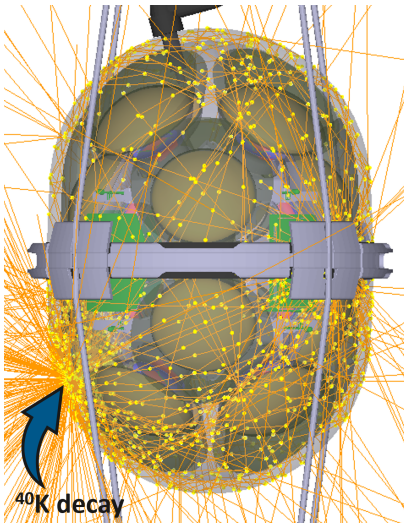


Figure 9.13: Geant4 visualization of an exemplary ^{40}K decay within the LOM16 pressure vessel. Orange lines represent photon trajectories and yellow markers indicate interaction vertices.

The ensuing α- and γ-emissions, along with subsequent secondary electrons, deposit energy into the glass. This energy transfer induces both scintillation and Cherenkov radiation, as shown in Figure 9.13, with electrons producing Cherenkov photons at kinetic energies exceeding the $\approx 160\,\text{keV}$ threshold in borosilicate glass, as per Equation (3.4). For each photon detection at the PMT photocathode, spatial and temporal information, as well as the decay origin and isotope, is recorded for analysis. The complete optical and particle physics processes implemented within this simulation are summarized in Table 9.4.

The rigorous modeling of scintillation properties (isotope concentrations from Section 9.2, lifetimes from [73 (pp. 174–182)], emission spectrum from Section 9.3 and yield from Section 9.4) isolates the background signal originating exclusively from the pressure vessel. However, for a realistic comparison against measurements of integrated optical modules, intrinsic PMT dark noise and afterpulsing are accounted for by analytically increasing the simulated pressure vessel scintillation rate $R_{\text{sim,vessel}}(T)$:

$$R_{\text{total}}(T) = R_{\text{sim,vessel}}(T) \cdot (1 + P_{\text{API+II}}(T)) + R_{\text{dark}}(T)/f_{\text{gel}}. \tag{9.6}$$

While it is technically possible to incorporate intrinsic PMT dark noise dynamically on an event-by-event basis in the simulation, the required scintillation properties of the glass of the PMT bulb (including proprietary glass bulb weight and isotope concentrations) are currently validated only for the mDOM PMT [73 (pp. 183–191)]. To establish an equivalent methodological baseline across all optical module architectures, the analytical scaling from Equation (9.6) is consistently applied.

Average PMT background rates from vessel scintillation $R_{\text{sim,vessel}}(T)$ are derived by calculating the ratio of all detected photons from a given PMT to the total simulation time. These rates are then translated to measurable hit rates by applying an artificial time window that replicates the specific dead time or triggered event logic of the respective module's electronics. The following subsections evaluate the expected total background rates for each optical module design, providing an essential tool for estimating the in-ice noise background. Because in-ice measurements are unfeasible, validation is performed

Table 9.4: Physics processes and corresponding object classes implemented in the simulation framework to model scintillation and Cherenkov emission. The custom OMSimScintillation-Process and OMSimRadioactiveDecayProcess classes are adapted from the standard Geant4 G4Scintillation and G4RadioactiveDecayBase modules.

Particle	Process	Class
Optical photon	Absorption	G4OpAbsorption
	Optical processes at medium interfaces	G4OpBoundaryProcess
	Mie scattering	G4OpMieHG
Ions	Scattering	G4hMultipleScattering
	Ionisation	G4ionIonisation
	Radioactive decay	OMSimRadioactiveDecayProcess
Alpha	Scattering	G4hMultipleScattering
	Ionisation	G4ionIonisation
Electron / Positron	Scattering	G4eMultipleScattering
	Ionisation	G4LivermoreIonisationModel
	Bremsstrahlung	G4eBremsstrahlung
	Cherenkov radiation	G4Cerenkov
	Positron annihilation	G4eplusAnnihilation
Gamma	Pair production	G4LivermoreGammaConversionModel
	Compton effect	G4LivermoreComptonModel
	Photoelectric effect	G4LivermorePhotoElectricModel
All particles above except photons	Material scintillation	OMSimScintillationProcess

by comparing the simulations to optical modules measured in air and wrapped entirely in absorptive black tape. The tape absorbs photons that would otherwise be transmitted into the surrounding ice at the nearly index-matched glass-ice interface, thereby approximating the optical boundary conditions of an in-ice deployment.

To highlight geometric dependencies, rates are grouped by PMT layer (equatorial and polar for the mDOM and LOM16, plus a vertical layer for the LOM18). D-Egg PMTs remain ungrouped due to their symmetry. DOMs are omitted entirely from this evaluation, as their background rates are already documented and well understood, having operated in situ for over a decade. Thus, the predictive framework developed here is tailored to establish operational baselines for the novel, upcoming IceCube-Gen2 optical module architectures prior to their deployment. In all subsequent plots, error bands represent statistical simulation fluctuations and all plotted values reflect the total integrated background rate as per Equation (9.6), unless explicitly stated otherwise.

mDOM

The temporal distribution of consecutive PMT hits stemming exclusively from Cherenkov and scintillation photons from the pressure vessel is analyzed via Δt histograms (see Section 6.2.3), as shown in Figure 9.14 (left). The pure vessel simulation resolves three distinct kinematic regions: a prompt Cherenkov peak around $\Delta t \approx 10^{-11}\,\mathrm{s}$, a broad delayed scintillation peak between $\Delta t \approx 10^{-8}\,\mathrm{s}$ and $\Delta t \approx 10^{-4}\,\mathrm{s}$ and an uncorrelated background peak around $\Delta t \approx 10^{-2}\,\mathrm{s}$ representing temporally overlapping decays. Adjacent to the prompt Cherenkov signal, a minor secondary peak emerges at $\Delta t \approx 2 \times 10^{-10}\,\mathrm{s}$. This feature originates from coincident Cherenkov photons undergoing internal PMT reflections (see Figure 8.2, left) before detection. As expected from the varied photon path lengths, the temporal position of this reflection peak shifts proportionally with the PMT size across the different module architectures. Note that both of these prompt kinematic features are typically obscured by hardware dead times of $\mathcal{O}(10\,\mathrm{ns})$ in measurements. Furthermore, while the simulation framework is capable of accurately smearing these signals by applying the empirically matched PMT transit time spread scans (detailed in Section 8.1), this temporal convolution was deliberately omitted in these specific histograms to explicitly resolve the underlying photon kinematics.

A deviation in the shape of the uncorrelated peak is observed when compared with the empirical data (Figure 9.14, right). Since PMT dark noise is partially uncorrelated, its presence in the empirical data naturally shifts the peak to the left and increases its amplitude. Moreover, the overall distribution is governed by the scintillation lifetime distribution. The observed discrepancy is an expected consequence of analytically adding the intrinsic PMT dark noise during post-processing, rather than simulating it dynamically as in [215 (p. 85)], as well as neglecting PMT afterpulsing around $\Delta t \approx 10^{-5.5}\,\mathrm{s}$. This deviation does not impact the integrated average rates and is only highlighted for completeness. In [215 (p. 87)], it is shown that the Δt ratios between in-ice (taped) and in-air mDOMs are consistent between simulation and measurement when dynamically simulating mDOM PMT noise.

To determine the total background rates, an afterpulsing probability of $P_{\mathrm{API+II}} = (5.65 \pm 0.03)\,\%$ from PMT acceptance testing is taken [155 (p. 143)]. For the intrinsic PMT dark rates of mDOMs, one must distinguish between low-noise (LN: $R_{\mathrm{dark}} = (100 \pm 10)\,\mathrm{s}^{-1}$)

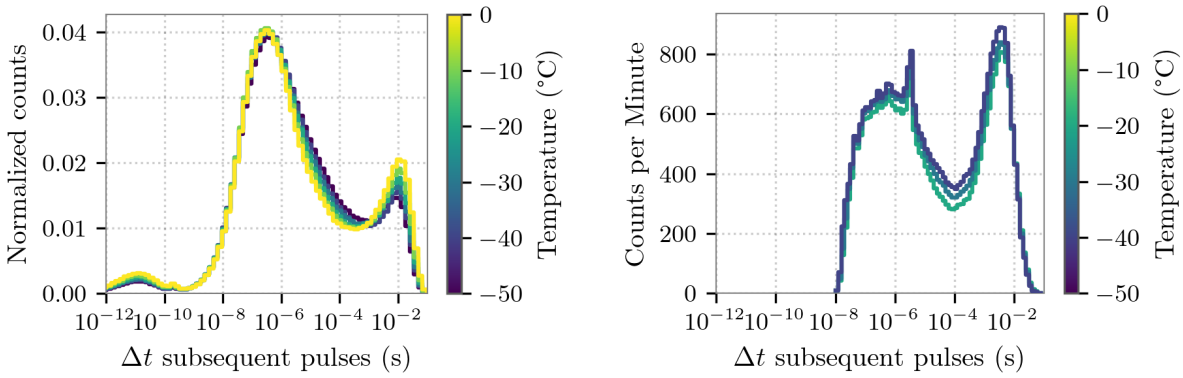


Figure 9.14: Average time difference Δt histograms for mDOM PMT background rates in an ice environment at various temperatures. **(left)** Simulation of only vessel scintillation, without artificial dead time. **(right)** Measurements from a taped mDOM mimicking the in-ice response (data adapted from [215 (p. 89)]).

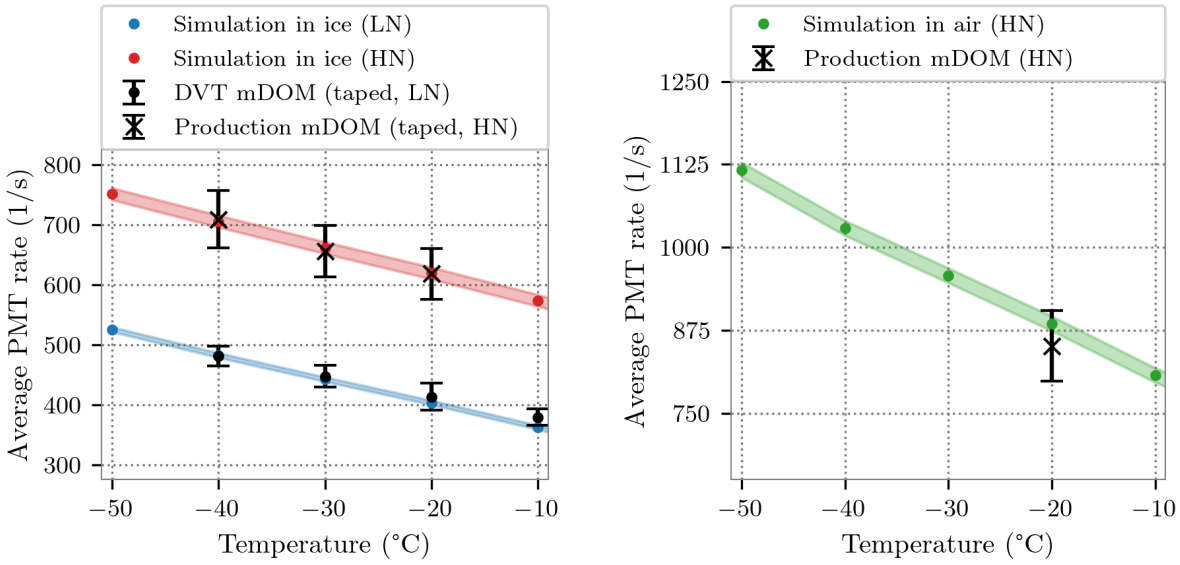


Figure 9.15: Average total mDOM PMT background rates at different temperatures. **(left)** In-ice simulation compared with measurements from a taped DVT-04 prototype (dot marker, low-noise PMTs, 80 ns dead time) [73] and a taped production mDOM (cross marker, high-noise PMTs, 10 ns dead time) [215]. **(right)** In-air simulation compared with measurements from an untaped production mDOM (10 ns dead time, high-noise PMTs) [215].

and high-noise (HN: $R_{\text{dark}} = (340 \pm 15)\,\text{s}^{-1}$, representing $\sim 90\,\%$ of all production mDOMs) PMTs [155 (pp. 147–148)] as previously introduced. Because the simulation yield was anchored to the high-noise production mDOM rates in ice, the alignment with this specific dataset is expected by design. The true validation of the simulation is demonstrated by its extrapolation to fundamentally different environments and hardware configurations. As illustrated in Figure 9.15, the simulation successfully tracks the design verification testing (DVT) mDOMs equipped with low-noise PMTs, confirming that the analytical decoupling of the PMT dark noise is accurate. Furthermore, the available in-air measurements of untaped modules exhibit strong consistency with the simulated rates. This validation demonstrates that the model accurately captures the complex geometric optics and internal reflections of the module.

While experimental results often limit reported measurements to a module-wide average, the simulation predicts slight geometric variations in the background rate between PMT layers. For example, at $-20\,°\text{C}$ in ice with a 10 ns dead time, high-noise equatorial PMTs exhibit an average rate of $(624 \pm 10)\,\text{s}^{-1}$, whereas polar PMTs register $(609 \pm 10)\,\text{s}^{-1}$. Additionally, the simulation predicts an air-to-ice rate reduction factor for mDOMs of 1.43 ± 0.03 at $-20\,°\text{C}$, which aligns with the empirically measured ratio of 1.38 ± 0.13 reported in [215].

D-Egg

As established in Section 9.2.2, the pressure vessels for the D-Egg and LOM18 are manufactured using Okamoto glass rather than Vitrovex glass. Although both are borosilicate glasses, differences in their production processes, raw material compositions and radiopurities lead to distinct scintillation properties. Consequently, the scintillation properties of Okamoto glass must be parameterized independently. While the scintillation yield and

lifetime for Okamoto glass have been constrained recently [216], a suitable glass sample for measuring the emission spectrum was not available for this thesis, necessitating the use of the Vitrovex spectrum as a functional proxy. If the intrinsic scintillation spectrum of Okamoto glass extends further into the UV regime, the yield would require remeasurement using a UV-sensitive PMT to bypass the 280 nm sensitivity cutoff of the utilized mDOM PMTs. Furthermore, the isotope concentrations were provided by an external company rather than measured in the setup from Section 9.2 due to the lack of suitable samples.

Empirical background rates at operational temperatures or detailed afterpulsing measurements for D-Egg PMTs are sparse. The intrinsic PMT dark rate at −40 °C in air is estimated to be $R_{\text{dark}} \approx (400 \pm 60)\,\text{s}^{-1}$ [217]. In [118], an air-to-ice reduction factor of 2.4 is applied to predict an overall in-ice average in-module PMT rate of $(908 \pm 87)\,\text{s}^{-1}$ with a dead time of 100 ns. Subtracting the intrinsic PMT noise isolates a target vessel-scintillation-induced background $R_{\text{sim,vessel}}$ of roughly $(508 \pm 106)\,\text{s}^{-1}$ per PMT.

Initial simulations utilizing the isotope concentrations supplied by the external company severely underestimated this target, producing only $(31.17 \pm 6.52)\,\%$ of the target vessel scintillation ($(61.49 \pm 8.86)\,\%$ in terms of total background rates) of the expected −40 °C value, as shown in Figure 9.16 (left). This discrepancy likely arises because the external company reports the activities without a detailed correction for internal geometric efficiency $\epsilon(E)$, as applied in Section 9.2. This implies that the true radionuclide concentrations within Okamoto glass may be higher than the ones stated in [10 (p. 116)].

Given these combined material uncertainties (uncorrected isotope activities, an assumed Vitrovex emission spectrum, imprecise PMT background and a lack of precise geometric coupling factors f_{gel} for the dark rate subtraction), the initial parameterization provides an incomplete picture. To establish a functional baseline for Okamoto glass in the simulation, the isotope concentrations are uniformly scaled by a factor of 3.16 to anchor the simulation to the −40 °C measurement target, as demonstrated in Figure 9.16 (right). It is important to note that this scaling serves as a necessary phenomenological proxy to match the measurement target rather than a definitive claim regarding the absolute

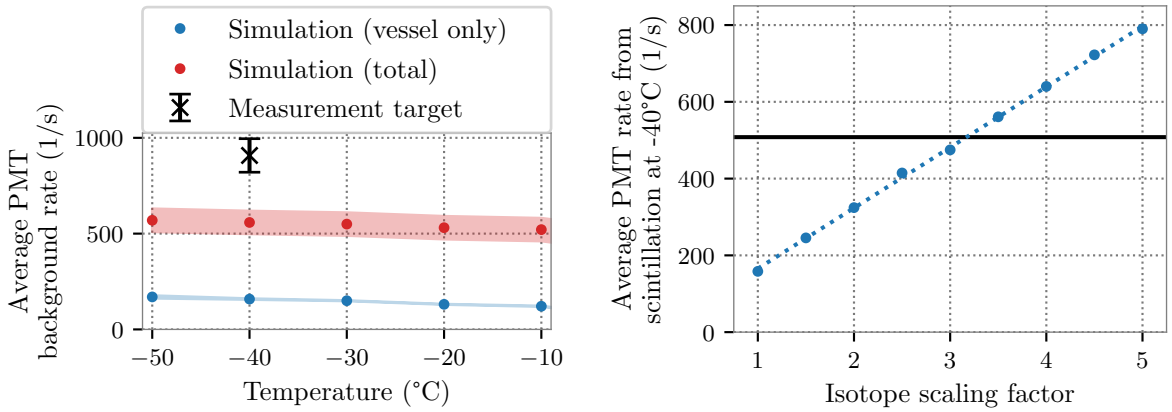


Figure 9.16: **(left)** Initial D-Egg in-ice average PMT background rates generated using uncorrected external isotope concentrations, significantly underestimating the measurement target from [118] (black cross). **(right)** Scaling of the provided isotope concentrations to match the expected average PMT rate solely due to vessel scintillation (taped, at −40 °C, 100 ns dead time).

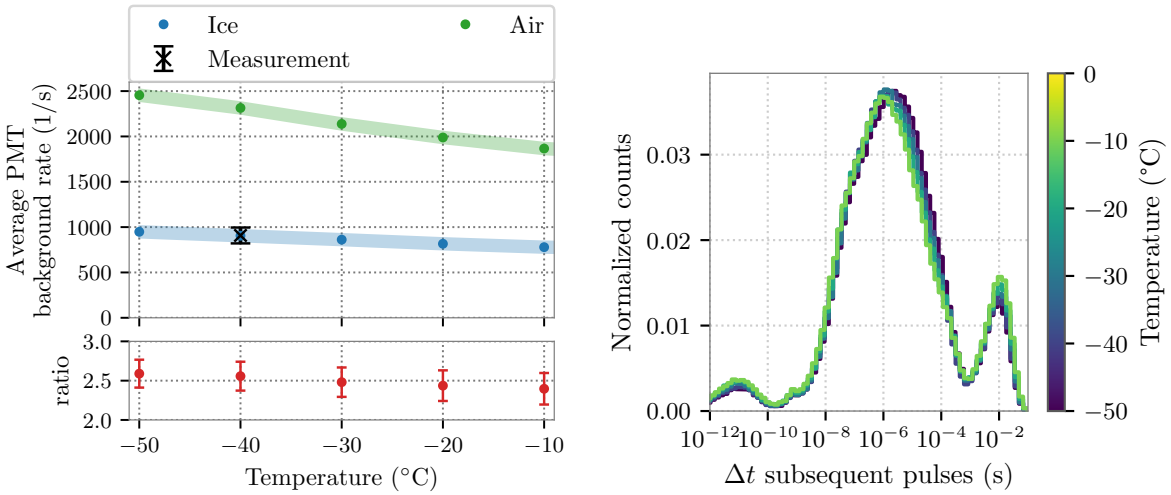


Figure 9.17: (left) Finalized D-Egg average PMT background rates in air (green) and ice (blue) compared with empirical values in [118]. The lower panel shows the resulting air-to-ice reduction ratio at different temperatures. **(right)** Average time difference Δt histograms for D-Egg PMTs from vessel scintillation in an ice environment at various temperatures.

activity of the Okamoto glass.

The simulation using this scaled isotope profile is shown in Figure 9.17 (left) and yields an air-to-ice reduction factor of 2.56 ± 0.18 at $-40\,°\mathrm{C}$. This aligns with the previously reported factor of ~ 2.4 [118]. The complementary time distributions between consecutive PMT hits stemming from vessel scintillation are shown in Figure 9.17 (right), for which no measurements of integrated PMTs at low temperatures could be found for comparison. As the intrinsic background from the large D-Egg PMTs is not included, the uncorrelated peak is likely of higher amplitude than shown. Despite the phenomenological assumptions required to bypass extensive laboratory characterization, these adjusted Okamoto scintillation properties form the basis for subsequent LOM18 modeling.

LOM16

The LOM16 uses a 10.7 kg Vitrovex pressure vessel, sharing scintillation properties with mDOM pressure vessels. This provides an excellent benchmark to evaluate the simulation's validity under a different geometry. Unlike the mDOM, the LOM employs the wuBase data-acquisition architecture. During final acceptance testing (FAT) reported in [127], module rates are stated as triggered active event windows rather than individual PMT pulse counts. When a pulse exceeds the 0.2 PE threshold, the 60 MHz digitizer (16.67 ns clock resolution) opens a pre- and post-trigger sampling window spanning 16 samples (266.72 ns) in each direction, establishing a minimum event length of 533.44 ns. Any subsequent photon detection falling within this active window dynamically extends the integration period and all enclosed pulses are cataloged as a single discrete event (see Figure 9.18, left).

Explicitly simulating this dynamically extending acquisition window on an event-by-event basis introduces computational complexity and prevents the extraction of standard pulse rates comparable with those of the mDOM and D-Egg. Instead, an analytical translation is applied to post-process the simulation output. The statistically averaged

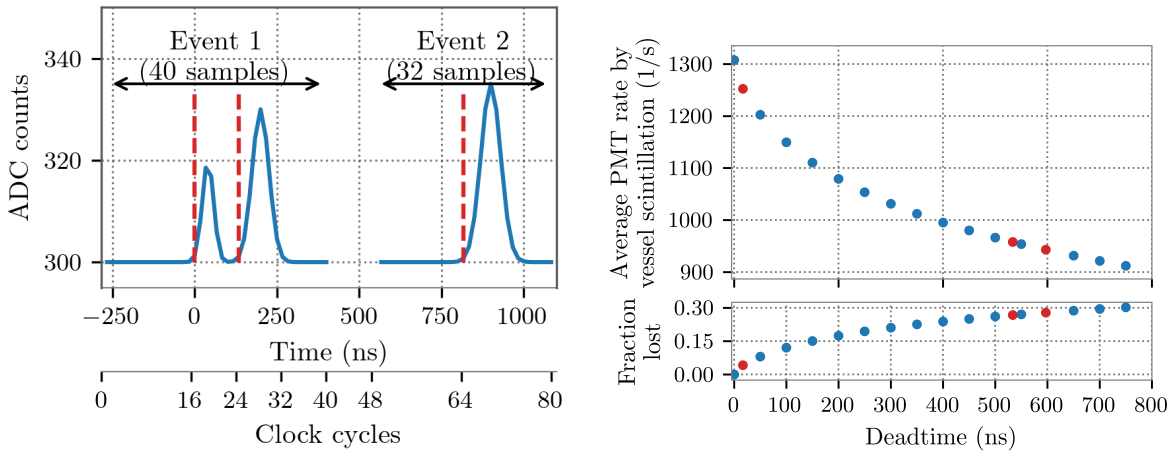


Figure 9.18: (left) Logic diagram of the wuBase DAQ triggering mechanism, demonstrating the dynamic extension of the active sampling window upon subsequent pulse detection. **(right)** Simulated average PMT background pulse rate solely due to vessel scintillation as a function of artificial dead time for the LOM16 in air at −40 °C. Red points indicate the clock resolution (16.67 ns), the minimum base window (533.44 ns) and the statistically averaged window duration (596.7 ns).

duration of these extending live-time windows is evaluated to be (596.7±1.2) ns at −40 °C in air. By applying this mean duration as an artificial, non-extending dead time to the simulated pulse timing distributions, the resulting simulation rates can be mapped to the FAT triggered event rates. Figure 9.18 (right) illustrates a representative translation curve used to correlate dead time with the resulting average PMT rate.

FAT measurements conducted in air at −40 °C provide the empirical baseline for the LOM16, as shown in Figure 9.19. The simulation accounts for PMT dark rates and afterpulsing as per Table 6.4. According to the FAT measurements, the average triggered event rate across four tested optical modules per PMT variant is $(1037 \pm 7)\,\mathrm{s}^{-1}$ for Hamamatsu-equipped modules and $(995 \pm 11)\,\mathrm{s}^{-1}$ for NNVT-equipped modules. The simulation shows

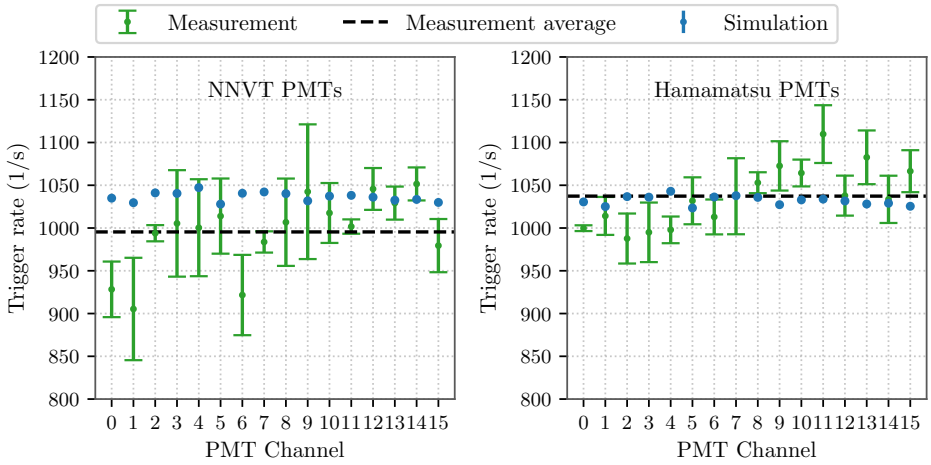


Figure 9.19: Comparison of simulated average LOM16 PMT background trigger rates with average FAT measurements in air at −40 °C. **(left)** Modules equipped with NNVT PMTs. **(right)** Modules equipped with Hamamatsu PMTs. Equatorial PMTs correspond to channels 0–7 and polar PMTs to channels 8–15.

good agreement, with total background rates of $(1032.1 \pm 1.2)\,\mathrm{s}^{-1}$ (Hamamatsu) and $(1036.5 \pm 2.1)\,\mathrm{s}^{-1}$ (NNVT). However, differences between PMT variants and layers are more pronounced according to the FAT measurements. This close agreement with empirical values, together with the corresponding agreement observed for the mDOM, provides an independent validation of the adopted scintillation parameterization for Vitrovex glass, as both optical modules employ the same pressure vessel material. The remaining deviations are largely dependent on the parameterization of the intrinsic PMT background.

The in-ice PMT rates for the LOM16 across the entire expected IceCube-Gen2 operational range (−45 °C to −10 °C) [10 (p. 89)] are presented in Figure 9.20 (left). To maintain consistency with mDOM and D-Egg parameterizations, these rates may be reported as standard PMT pulses per second (using a 10 ns dead time) rather than wuBase trigger event rates. This translates to an expected in-ice pulse background rate of $(679.4 \pm 1.6)\,\mathrm{s}^{-1}$ (Hamamatsu) and $(694.0 \pm 2.9)\,\mathrm{s}^{-1}$ (NNVT) at −40 °C. This is comparable with the mDOM's $(705.9 \pm 9.9)\,\mathrm{s}^{-1}$ average total background PMT rate, which is heavily skewed due to the higher intrinsic PMT dark rate of $(340 \pm 15)\,\mathrm{s}^{-1}$ compared with the LOMs' $(56.3 \pm 2.7)\,\mathrm{s}^{-1}$ (Hamamatsu) and $(94.0 \pm 4.8)\,\mathrm{s}^{-1}$ (NNVT). The (28.1 ± 1.1) % increase in the background rate due to vessel scintillation in the LOM16 is driven mainly by the larger 4-inch PMT photocathode area and the geometry of the optical coupling. Specifically, the larger, bulkier gel pads act as extended light guides, capturing more photons before they can scatter or be absorbed by the internal holding structures. Because physically taped LOM16s were not measured during FAT, the air-to-ice scaling relies entirely on the simulated geometry, which has been demonstrated to be accurate for Vitrovex glass using the mDOM previously. At −40 °C, the simulation models environmental suppression from air to ice by a factor of 2.14 ± 0.02. The PMT pulse timing distributions (Figure 9.20, right) stemming solely from vessel scintillation are similar to those of the mDOM, as the Vitrovex scintillation lifetime parameterization is shared. However, the uncorrelated peak

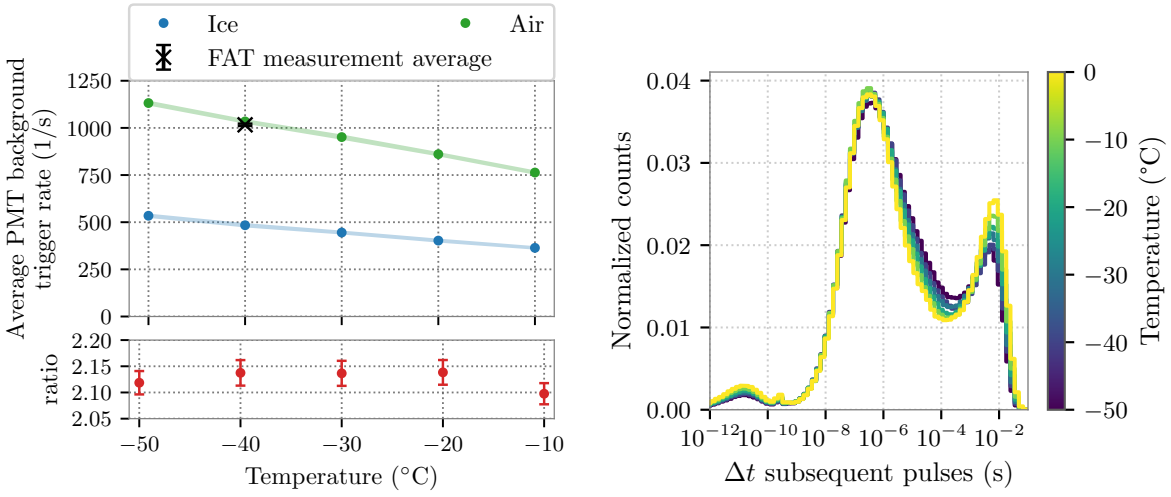


Figure 9.20: (left) Predicted PMT background trigger rate (averaged across Hamamatsu and NNVT variants) for the LOM16 over the anticipated operational temperature range in both environments. Rates are shown as wuBase event rates, in contrast to true pulse rates reported for mDOMs and D-Eggs. The lower panel shows the simulated air-to-ice reduction ratio across temperatures. **(right)** Average time difference Δt histograms for LOM16 PMTs from vessel scintillation in an ice environment at various temperatures.

around $\Delta t \approx 10^{-2}\,\mathrm{s}$ is slightly more pronounced. The wuBase event triggering logic will suppress part of the scintillation peak of the timing distribution, as an effective dead time of $\gtrsim 5 \times 10^{-7}\,\mathrm{s}$ is applied.

LOM18

The LOM18 features a 17 kg Okamoto glass vessel with the wuBase DAQ architecture introduced with the LOM16. The simulation draws upon the previously derived scintillation parameters, utilizing the 3.16 radionuclide concentration scaling factor established for the D-Egg and a $(568.6 \pm 1.1)\,\mathrm{ns}$ mean wuBase event window logic at $-40\,^{\circ}\mathrm{C}$ in air.

The simulation is in good overall agreement with the expected event rates when compared with FAT measurements averaged across five Hamamatsu- and two NNVT-equipped modules (Figure 9.21). However, it slightly overestimates the trigger rates by $(60 \pm 24)\,\mathrm{s}^{-1}$ for the NNVT PMTs and $(40 \pm 19)\,\mathrm{s}^{-1}$ for the Hamamatsu PMTs on average. These discrepancies may arise from uncertainties in the derived scintillation parameterization, differences between the PMTs characterized in this work (Section 6.2.3) and those installed in the LOM18 modules, or the limited livetime of the FAT measurements. The combined duration of all trigger windows amounts to less than one minute per PMT, resulting in relatively large statistical uncertainties.

Because the LOM18 incorporates more distinct PMT layering, Table 9.5 dissects the rates by specific PMT layers. Both the simulation and the empirical data exhibit a decreased background rate for polar PMTs relative to the equatorial and vertical PMTs. This spatial variance is attributable to the optical coupling geometry, as the polar PMTs are situated behind smaller gel pads, reducing their solid-angle acceptance of photons originating in the thicker equatorial sections of the pressure vessel.

Finally, the extrapolated in-ice rates for the LOM18 are presented in Figure 9.22 (left). A module-specific measurement is available for validation, wherein a single LOM18 was evaluated at $-40\,^{\circ}\mathrm{C}$ both bare and covered with absorptive tape. The simulation demonstrates excellent agreement with this module across both environments, further

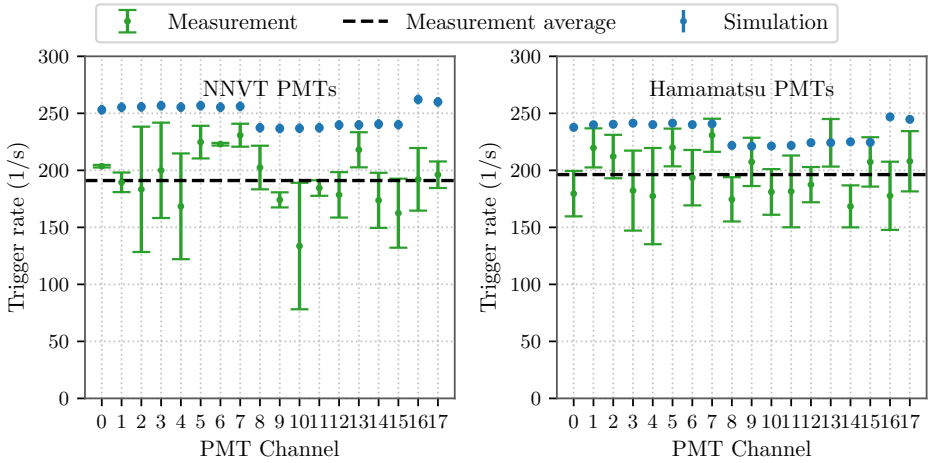


Figure 9.21: Comparison of simulated average LOM18 PMT background trigger rates with average FAT measurements in air at $-40\,^{\circ}\mathrm{C}$. **(left)** NNVT PMT variant. **(right)** Hamamatsu PMT variant. Equatorial PMTs correspond to channels 0–7, polar PMTs to channels to 8–15 and the vertical PMTs to channels 16 & 17.

Table 9.5: Comparison of simulated average background rates with FAT measurements for the LOM18 in air at −40 °C, categorized by PMT placement geometry.

PMT Type	Source	PMT Layer	Rate (s^{-1})
Hamamatsu	FAT Measurement	Equatorial	202 ± 11
		Polar	178 ± 10
		Vertical	194 ± 15
	Simulation	Equatorial	240.2 ± 1.4
		Polar	223.0 ± 1.4
		Vertical	245.8 ± 1.7
NNVT	FAT Measurement	Equatorial	202 ± 9
		Polar	192 ± 8
		Vertical	193 ± 20
	Simulation	Equatorial	255.6 ± 2.4
		Polar	238.6 ± 2.4
		Vertical	261.1 ± 2.6

validating the underlying framework, whereas the average FAT rate in air at −40 °C is slightly lower as discussed before. When translating the wuBase-triggered in-ice rates of $(114.6 \pm 5.2)\,s^{-1}$ (Hamamatsu) and $(131.4 \pm 5.5)\,s^{-1}$ (NNVT) to standard pulse rates with a 10 ns dead time, the simulation predicts rates of $(142.6 \pm 1.7)\,s^{-1}$ and $(163.2 \pm 2.8)\,s^{-1}$ at −40 °C in ice, respectively. Furthermore, the FAT measurement yields an air-to-taped reduction ratio of 2.06 ± 0.39, while the corresponding simulated reduction factor is 1.96 ± 0.08, demonstrating geometric agreement. The timing distributions (Figure 9.22,

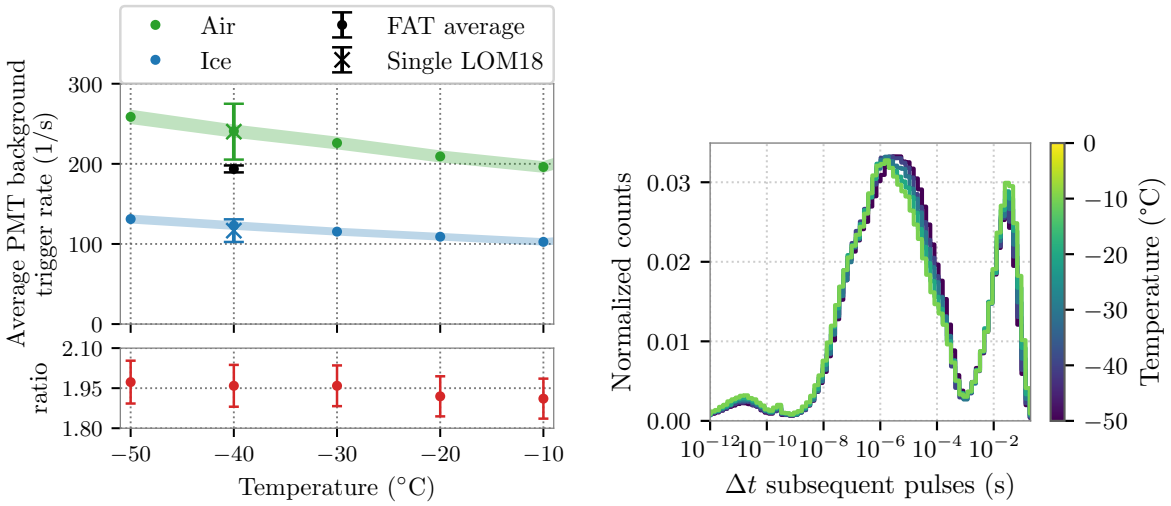


Figure 9.22: (left) Predicted average PMT background rate (averaged across Hamamatsu and NNVT variants) for the LOM18, shown as wuBase event rates across temperatures and environments. Results are validated against measurements from a single LOM18 which was also taped (cross marker) and average FAT data (dot marker) from Figure 9.21, both measured at −40 °C. The lower panel shows the simulated air-to-ice reduction ratio across temperatures. **(right)** Average time difference Δt histograms for LOM18 PMTs from vessel scintillation in an ice environment at various temperatures.

right) reveal an even larger uncorrelated component compared with LOM16s, which is expected as fewer scintillation photons are produced in Okamoto glass, increasing the relative probability of detecting coincident uncorrelated decays. Similar to D-Eggs, adding intrinsic PMT dark noise will increase the relative amplitude of the uncorrelated peak even further.

Comparing both LOM geometries directly, the expected background event rates in ice at $-40\,^{\circ}\mathrm{C}$, averaged over both PMT variants, are reduced by a factor of 3.93 ± 0.13 for the LOM18. The corresponding reduction factor in air is 4.29 ± 0.12, or 5.25 ± 0.12 when considering the FAT measurement averages. These results corroborate that Okamoto glass yields lower background rates than current Vitrovex pressure vessels. As discussed in Section 9.2.2, the elevated background in Vitrovex glass likely originates from higher trace concentrations of radionuclides in the quartz sand and a higher potassium nitrate content in the glass batch.

9.6 Summary and Outlook

This chapter establishes a comprehensive, simulation-derived optical background model for both IceCube Upgrade LOMs, providing the baseline background modeling for the Gen2-OM. The parameterization is rooted in empirical measurements, where the isotope concentrations within the Vitrovex glass pressure vessels are consistently lower than those previously reported. Notably, no significant vessel-to-vessel variations were observed, a stability metric that will be continuously monitored across multiple batches during IceCube-Gen2 mass production. Moreover, the analysis of individual Vitrovex glass components identifies quartz sand as the primary source of heavy chain contamination, challenging prior hypotheses that attributed this activity primarily to furnace leaching. In addition, the scintillation spectrum measurement achieved an excellent SNR, revealing a consistent spectral shape across different temperatures.

For the average expected total background rates, both the mDOM and LOM16 simulations align closely with measurements, providing validation for the Vitrovex scintillation properties. In contrast, modeling Okamoto glass based on the D-Egg proved more challenging due to a lack of detailed comparative measurements. Here, the isotope concentrations were scaled using empirical D-Egg target rates to establish a functional baseline. The subsequent LOM18 FAT measurement comparison confirms that this scaling approach is accurate. Both the measurements and simulations exhibit parallel trends under parameter variations. Discrepancies remain in the Δt distributions, largely stemming from the analytical treatment of dark noise and the absence of dynamic type II afterpulsing in the event generation. These simulation-derived background rates constitute the baseline noise model injected into the IceTray software architecture for IceCube-Gen2 forecasting, pending further validation from in situ IceCube Upgrade data.

Future experimental efforts should focus on completing the characterization of Okamoto glass scintillation properties, particularly its intrinsic radionuclide concentrations, as no suitable glass samples or bare vessels were available for this study. Additionally, the emission spectrum should be measured with the setup described in Section 9.3.1, ideally extending to a broader range of bandpass filters. If the Okamoto spectrum extends deeper into the UV range than that of Vitrovex, both scintillation yield and lifetimes need to be

reassessed using UV-sensitive PMTs. Furthermore, bare D-Egg PMTs should be characterized analogously to Chapter 6 to provide dark-rate and Δt histogram targets for the intrinsic PMT background simulation, while also supplying the measurements required for a future implementation of PMT afterpulsing. In addition, the scintillation properties of the 4-inch LOM and 8-inch D-Egg PMTs must be characterized to complete the overall background model. Finally, while the total module background (including atmospheric muons) for an mDOM submerged in a water tank has been successfully measured and simulated in [218] under the author's supervision, equivalent evaluations for the LOMs remain an important objective for future studies.

Part III

IceCube Upgrade Sensitivity Study on Neutrino Emission from active Galaxy NGC 1068

10 NGC 1068: A Key to understanding high-energy Astrophysics

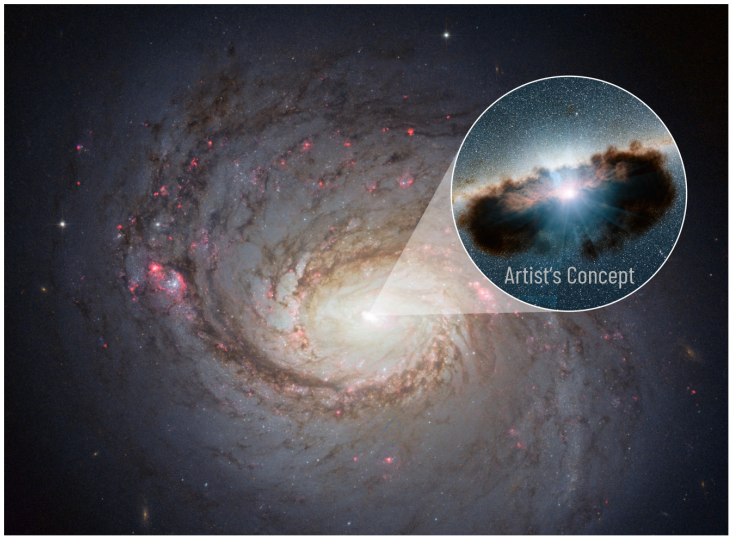


Figure 10.1: NGC 1068 as imaged by the Hubble Space Telescope. The inset shows an artistic interpretation of the AGN based on NuSTAR X-ray observations. Taken from [219].

NGC 1068 (Messier 77), as depicted in Figure 10.1, is a nearby barred spiral galaxy and the archetype of a Seyfert II active galactic nucleus. Located in the constellation Cetus at a commonly adopted distance of 14.4 Mpc (∼47 million light-years), it serves as a primary target object in AGN research. Its nucleus hosts a (1–2) $\times 10^7$ $M_\odot$ supermassive black hole deeply embedded in a dense, Compton-thick circumnuclear torus [220]. From Earth's perspective, this geometry obscures the central engine in optical and soft X-ray bands, producing the characteristic Seyfert II "hidden core" signature, as introduced in Section 2.2.2. Additionally, NGC 1068 features intense starburst activity and multiphase outflows, making it an important site for studying AGN-host galaxy feedback and cosmic ray acceleration [221]. Historically, NGC 1068 has been central to the validation of the unified AGN paradigm and remains a primary target in multi-wavelength and multi-messenger astrophysics due to its rich emission spectrum.

10.1 Astrophysical Insights and high-energy Particle Acceleration

Recent high-resolution radio [222] and X-ray [223] observations have provided evidence for jet-disk interactions, identifying both relativistic jets and the accretion disk corona as promising sites for cosmic ray acceleration via diffusive shock acceleration and magnetic reconnection [45]. In this dense, photon-rich environment, cosmic rays undergo hadronic

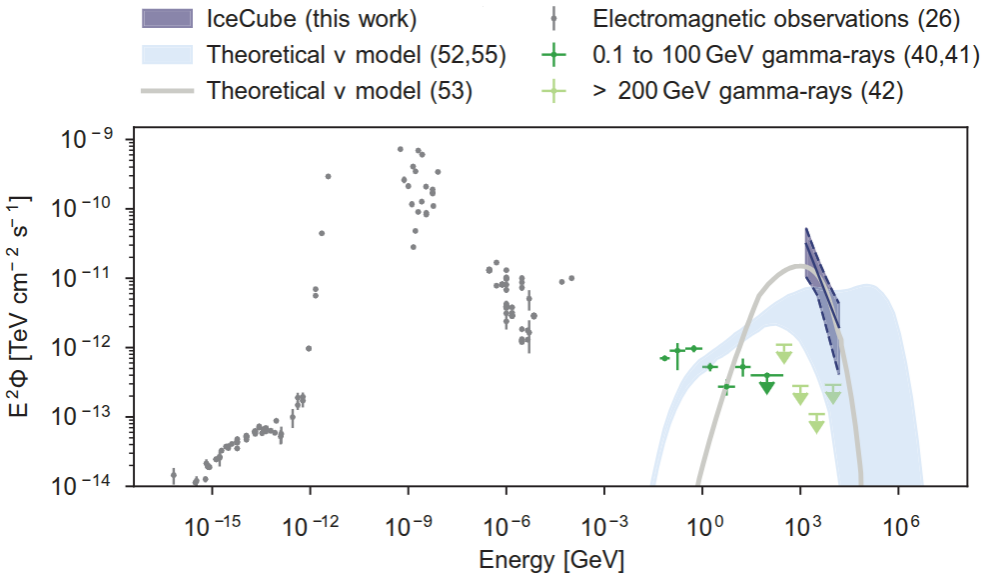


Figure 10.2: Multi-messenger spectral energy distribution of NGC 1068. The neutrino flux measurement from IceCube (dark blue band) is compared with electromagnetic observations (data points) and theoretical models (lines). Taken from [6].

(pp) and photohadronic ($p\gamma$) interactions, yielding charged and neutral pions and consequently high-energy neutrinos and γ-rays (see Section 2.2.1).

A defining feature of NGC 1068 emerging from recent observations is its "hidden" cosmic ray accelerator status. As illustrated in Figure 10.2, the neutrino flux observed by IceCube exceeds the potential $\mathcal{O}(100\,\mathrm{GeV})$ γ-ray flux limits by at least one order of magnitude [6]. While conventional hadronic models predict neutrino and γ-ray fluxes of similar magnitude, observations reveal NGC 1068 to be "neutrino-bright but γ-ray-dim" at higher energies [224]. This tension between the high neutrino flux and the γ-ray limits suggests a high-opacity environment where γ-rays are absorbed and reprocessed to lower energies [45, 224]. Indeed, γ-rays produced in the disk and corona are heavily suppressed by $e^{\pm}$ pair production on X-ray and IR fields, cascading their energy down to MeV γ-rays or hard X-rays before escaping [225]. Neutrinos, in contrast, interact predominantly via the weak force and escape the AGN core region unperturbed, providing a unique probe of the otherwise unobservable central environment.

Surveys searching for time-integrated emission from AGN populations [48] show that γ-ray-loud blazars and radio-loud AGN, despite their γ-ray brightness, contribute only a modest fraction of IceCube's observed diffuse astrophysical neutrino flux. This motivates a detailed investigation of obscured, starburst or hidden AGNs such as NGC 1068, allowing for the revision of models describing cosmic ray acceleration, which are now considered viable contributors to the high-energy astrophysical neutrino flux.

Despite intensive campaigns with Fermi-LAT [26], HAWC [28] and other γ-ray telescopes, neither photon nor neutrino emission at the PeV scale has been firmly associated with NGC 1068, suggesting it is not a *PeVatron*. The dominant acceleration and interaction processes appear to cap at the sub-PeV scale, though true ratios depend on local absorptive and geometric properties. The gap between observed neutrino and suppressed γ-ray fluxes supports models with efficient γ-ray absorption and cascade reprocessing. Computational modeling incorporating detailed AGN structure, intense photon

fields and environmental absorption [224] remains essential for interpreting observations and for probing the differences in acceleration mechanisms between obscured AGN and blazars.

10.2 The strategic Value of NGC 1068 in IceCube and multi-Messenger Astrophysics

NGC 1068 occupies a central position in high-energy neutrino astronomy and the broader multi-messenger landscape. In a landmark result in 2022 [6], the IceCube Neutrino Observatory identified a statistically significant ($> 4\sigma$) persistent flux of 79^{+22}_{-20} high-energy neutrinos from the vicinity of NGC 1068's core. Unlike earlier neutrino source candidates based on single events or transient flaring activity, this excess is supported by a population of events accumulated over many years of observation, enabling the study of NGC 1068 as a steady astrophysical emitter.

As shown in the high-resolution scan in Figure 10.3, the optical position of NGC 1068 lies within the 68 % confidence region of the neutrino excess, located just 0.11° from the best-fit hotspot. This spatial coincidence, combined with the high signal statistics, established the first statistically significant association of a steady but obscured AGN with a high-energy neutrino source. The measured best-fit spectrum from this initial discovery is a power law with a spectral index of $\gamma \approx 3.2$. While softer than canonical expectations, this index is consistent with internal absorption and cascade models. More recently, a follow-up study [226] utilizing expanded datasets suggests an even softer index of $\gamma \approx 3.4$.

NGC 1068 stands out as the most significant and best-characterized time-independent (steady-state) neutrino point source identified to date. It contrasts with earlier discoveries

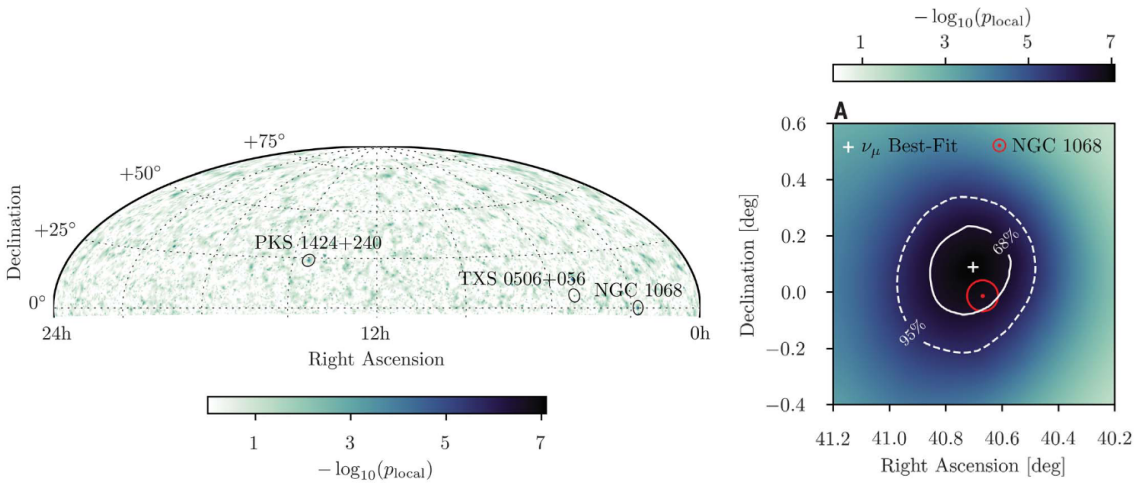


Figure 10.3: Sky map of the point-source scan in the Northern Hemisphere. The color scale indicates the logarithm of the local p-value obtained from the maximum likelihood analysis, evaluated at each location in the sky with the spectral index as a free parameter. The map is shown in equatorial coordinates using a Hammer-Aitoff projection. The black circles indicate the three most significant objects from the source list search. The circle around NGC 1068 marks the most significant location in the Northern Hemisphere, enlarged in the right inset. Taken from [6].

of transient emission, such as the flare from blazar TXS 0506+056 [5, 227]. As such, it signal provides a "standard candle" for the cross-calibration of present and next-generation neutrino detectors. The sample size permits detailed spectral analysis and model discrimination (e.g., cascade versus direct emission, or gas and dust torus absorption), while the high event statistics allow for the physical validation of detector performance metrics, such as directional reconstruction and background suppression algorithms.

While the majority of IceCube's measured diffuse astrophysical neutrino flux remains unresolved into individual point sources [228], the ∼80 neutrino events linked to NGC 1068 underscore its strategic importance. Its persistent signal provides a benchmark for extracting astrophysical events from heavy atmospheric backgrounds and testing improved analytic approaches. The combination of X-ray obscuration and strong neutrino emission motivates coordinated multi-wavelength studies involving observations from CTA [229], ULTRASAT [230] and future X-ray missions. These efforts will refine spectral models and help resolve the physical mechanisms driving the observed emissions.

The recently deployed IceCube Upgrade and the planned IceCube-Gen2 extension promise significantly improved sensitivity, energy range and angular resolution. These enhancements will deepen our view into AGN cores, constrain low-energy spectral cutoffs and improve searches for time-dependent emission. NGC 1068's flux and spectral characteristics ensure that it will remain central to IceCube's science goals and the evolving multi-messenger landscape, enabling ongoing advances in both AGN physics and the search for the origins of high-energy cosmic neutrinos.

IceCube Upgrade provides the first opportunity to study astrophysical neutrino sources with a detector configuration that includes newly developed optical module technologies. While the complete IceCube-Gen2 detector will require dedicated simulations beyond the scope of this thesis, IceCube Upgrade serves as an essential intermediate step for evaluating the expected performance of these detector concepts. NGC 1068, with its steady neutrino emission and soft spectral index, provides a well-suited benchmark for this purpose.

11 Point Source Analysis Method

The detection of astrophysical neutrino sources requires distinguishing a small population of cosmic neutrino events from the overwhelming background of atmospheric muons and neutrinos. This challenge is addressed through statistical methods that exploit the characteristic spatial and energy distributions of signal and background events. This chapter introduces the unbinned maximum likelihood method used for neutrino point source searches in IceCube, which is based on [231] unless stated otherwise. A point source analysis evaluates the significance of potential sources by testing for spatial and energy-dependent excesses of events relative to the expected atmospheric background. By leveraging individual event-level observables (specifically the reconstructed direction, angular uncertainty estimate and energy proxy), this statistical framework provides improved separation between signal and background compared with binned approaches. In this thesis, the method is implemented using the `csky` analysis package [232], a standardized software framework for IceCube point source searches.

11.1 The unbinned Likelihood Function

The core of the analysis is the likelihood function $\mathcal{L}$, which quantifies the compatibility of the observed data with a composite hypothesis consisting of a localized astrophysical signal flux and a continuous background. The physical composition of this background is energy-dependent: while it is dominated by atmospheric muons and atmospheric neutrinos at lower and intermediate energies, the isotropic diffuse astrophysical neutrino flux itself becomes the primary background component at the highest energies (typically $\gtrsim \mathcal{O}(100\,\mathrm{TeV})$ for tracks [228]). The method is *unbinned* because it utilizes individual event observables directly rather than relying on histogrammed data, thereby maximizing the spatial and kinematic information retained from the detector.

Single-Dataset Likelihood

For a dataset containing N observed events, the standard likelihood is defined as the product of the probabilities for each individual event i:

$$\mathcal{L}(n_s, \gamma) = \prod_{i=1}^{N} \left[\frac{n_s}{N} \mathcal{S}_i(\vec{x}_s, \gamma) + \left(1 - \frac{n_s}{N}\right) \mathcal{B}_i \right]. \tag{11.1}$$

Here, n_s is the number of signal events (the primary free parameter) and γ is the spectral index of the source (also a potential free parameter). $\mathcal{S}_i$ and $\mathcal{B}_i$ represent the signal and background probability density functions (PDFs) for the i-th event, respectively.

Multi-Dataset Likelihood

Since the analysis described in Chapter 12 combines data from different detector configurations (namely IceCube + DeepCore and IceCube Upgrade), the likelihood is generalized

to combine M independent datasets. The joint likelihood is the product of the individual dataset likelihoods:

$$\mathcal{L}(n_s, \gamma) = \prod_{j=1}^{M} \prod_{i=1}^{N_j} \left[\frac{n_s^j}{N_j} \mathcal{S}_i^j(\vec{x}_s, \gamma) + \left(1 - \frac{n_s^j}{N_j}\right) \mathcal{B}_i^j \right].$$

Here, N_j is the total number of events in dataset j. The expected signal count in a specific dataset n_s^j is not a free parameter but is derived from the total fitted signal flux Φ_0 (equivalently, the total fitted n_s) weighted by the relative acceptance (effective area and livetime) of the corresponding dataset.

Time-Dependent Likelihood Extension

For transient analyses where emission is confined to a short time window (see Section 13.3), the extended likelihood formulation is employed [233]. Unlike the standard likelihood, which fixes the total number of events N, this approach accounts for Poisson fluctuations in the event counts, which become statistically significant over short durations:

$$\mathcal{L}(n_s, \gamma) = \frac{(n_s + n_b)^N}{N!} e^{-(n_s+n_b)} \prod_{i=1}^{N} \left(\frac{n_s \mathcal{S}_i(\gamma) + n_b \mathcal{B}_i}{n_s + n_b} \right).$$

Mathematically, this formulation allows the total number of observed events to fluctuate according to Poisson statistics, enabling the analysis to evaluate the significance of observing N events when the expected background count n_b is small.

Probability Density Functions (PDFs)

The ability to distinguish astrophysical neutrinos from atmospheric background relies on accurate PDFs. These are factorized into spatial and energy-dependent terms to capture the distinct signatures of signal events.

Signal PDF ($\mathcal{S}_i$)

The signal PDF describes the probability of observing an event with reconstructed properties $(\vec{x}_i, E_i)$ from a neutrino originating from a source at $\vec{x}_s$ with a power law spectrum $E^{-\gamma}$:

$$\mathcal{S}_i(\vec{x}_s, \gamma) = S_{\text{space}}(\vec{x}_i | \vec{x}_s, \sigma_i) \times S_{\text{energy}}(E_i, \delta_i | \gamma).$$

The **spatial term** S_{space} quantifies directional clustering. It is modeled analytically, typically as a 2D Gaussian centered at the source coordinates $\vec{x}_s$:

$$S_{\text{space}}(\vec{x}_i | \vec{x}_s, \sigma_i) = \frac{1}{2\pi\sigma_i^2} \exp\left(-\frac{|\vec{x}_i - \vec{x}_s|^2}{2\sigma_i^2} \right). \tag{11.2}$$

Here, $\vec{x}_i$ is the reconstructed event direction and σ_i is the event-by-event angular uncertainty estimator. For high-energy tracks, this is often approximated as a 2D Gaussian. However, for analyses dominated by lower-energy events associated with larger angular uncertainties, a Kent distribution is employed [233]. The Kent distribution accounts for the curvature of the celestial sphere, providing a more accurate description of the point

spread function (PSF) tails than a planar Gaussian approximation. For extended sources, the uncertainty σ_i is combined with the source extension width σ_{src}.

Unlike the spatial term, the **energy term** S_{energy} does not have a closed analytic form. It represents the probability of reconstructing an energy E_i given a true power law source spectrum $\Phi \propto E^{-\gamma}$ and is constructed numerically using Monte Carlo simulations. In csky, this is implemented by weighting the simulation according to the trial spectral index γ and fitting a multidimensional spline to the resulting energy distribution.

Background PDF ($\mathcal{B}_i$)

The background PDF characterizes the total background observed by the detector, which is dominated by atmospheric muons and atmospheric neutrinos over most of the energy range. As these are characterized by detector effects that are difficult to model analytically, the background model is derived entirely from experimental data:

$$\mathcal{B}_i = B_{\text{space}}(\delta_i) \times B_{\text{energy}}(E_i, \delta_i).$$

For the **spatial term** B_{space}, the background is assumed to be uniform in right ascension (RA). Over long observation periods, the Earth's rotation averages out local detector asymmetries, while the detector acceptance depends primarily on declination δ. The declination dependence of the acceptance is modeled numerically as:

$$B_{\text{space}}(\delta_i) = \frac{1}{2\pi} P_{\text{data}}(\delta_i),$$

where $P_{\text{data}}(\delta_i)$ is determined by a spline fit to the declination distribution of the full experimental dataset.

The **energy term** B_{energy} captures the declination-dependent energy spectrum of the atmospheric background as measured by the detector. It is constructed by binning the experimental data in declination bands and generating a smoothed representation (spline or kernel density estimate) of the reconstructed energy distribution within each band.

Numerical Implementation

For computational efficiency and numerical stability, the likelihood is often expressed using the ratio of signal to background PDFs:

$$W_i(\vec{x}_s, \gamma) = \frac{\mathcal{S}_i(\vec{x}_s, \gamma)}{\mathcal{B}_i}. \tag{11.3}$$

Substituting this into the likelihood definition (Equation (11.1)) yields a simplified form used in the maximization routines:

$$\ln \mathcal{L}(n_s, \gamma) = \sum_{i=1}^{N} \ln \left[\frac{n_s}{N}(W_i - 1) + 1 \right] + \text{const.}$$

This formulation highlights that events with $W_i \gg 1$ (signal-like) favor larger fitted values of n_s, while events with $W_i < 1$ (background-like) favor the null hypothesis.

11.2 Hypothesis Testing and Test Statistic

To quantify the significance of an excess, a hypothesis test is performed comparing the best-fit signal-plus-background hypothesis H_1 against the background-only null hypothesis H_0.

- **Null Hypothesis (H_0):** $n_s = 0$. All observed events are consistent with atmospheric background.
- **Alternative Hypothesis (H_1):** $n_s > 0$. An excess of events exists, best described by a source with parameters $\hat{n}_s$ and $\hat{\gamma}$.

The compatibility of the data with these hypotheses is evaluated using a test statistic (TS) based on the log-likelihood ratio:

$$\mathrm{TS} = -2 \ln \left(\frac{\mathcal{L}(n_s = 0)}{\mathcal{L}(\hat{n}_s, \hat{\gamma})} \right), \tag{11.4}$$

where $\hat{n}_s$ and $\hat{\gamma}$ maximize the likelihood under the signal-plus-background hypothesis. As the ratio is taken, the numerical implementation from Equation (11.3) is naturally justified. Larger TS values indicate that the signal hypothesis provides a better description of the data than the null background hypothesis.

Statistical Interpretation

The significance of an observed TS is determined empirically. By performing thousands of pseudo-experiments on randomized background data (see Section 11.3), a background-only TS distribution is constructed. The p-value is defined as the probability of obtaining a TS at least as large as the observed value under the null hypothesis [48]:

$$p = P(\mathrm{TS}_{\mathrm{bkg}} \geq \mathrm{TS}_{\mathrm{obs}} \mid H_0).$$

According to Wilks' theorem, in the asymptotic limit, the positive tail of the background TS distribution follows a χ^2 distribution, allowing the conversion of p-values into standard deviations σ for reporting significance. The χ^2 distribution has degrees of freedom equal to the number of free parameters (typically 2: n_s and γ). However, because n_s is physically constrained to be non-negative, the full distribution is a mixture of a Dirac delta function at TS = 0 (representing negative statistical fluctuations where the best fit is $n_s = 0$) and a χ^2 tail for positive fluctuations. For finite background samples, this behavior is approximated by fitting the distribution with a χ^2 function scaled by a non-zero fraction η. This framework allows for the characterization of type-I errors (false positives, significance level α) and type-II errors (false negatives, β). The power of the test $(1 - \beta)$ expresses the ability to discover a true source.

When testing multiple hypotheses (such as scanning the entire sky or evaluating a catalog) the significance must be corrected for the look-elsewhere effect to obtain a global p-value. Analytical adjustments like the Bonferroni correction are overly conservative for sky scans because the fine grid of test locations is correlated due to the detector's angular resolution. Instead, IceCube analyses employ an empirical trials correction: the analysis records the maximum TS obtained anywhere in a scrambled background map and compares the unblinded TS against this distribution, inherently accounting for spatial correlations.

11.3 Source Injection, Sensitivity and Discovery Potential

To assess sensitivity and discovery potential as well as to set upper limits, the analysis performance is quantified through large ensembles of pseudo-experiments referred to as trials.

Background Trials

Background-only trials for the H_0 hypothesis are constructed by randomly scrambling the right ascension of real data events, which washes out any physical sources while preserving the detector's declination-dependent background rate and energy response. For each trial, the full likelihood fit is performed to obtain a TS value. Repeating this thousands of times builds the empirical background TS distribution, which incorporates true detector characteristics and systematic effects.

Signal Injection

Signal injection is performed to test the H_1 hypothesis. Simulated astrophysical neutrino events are added to the scrambled background trials:

- **Weighting:** Each injected event is assigned a probability based on its *OneWeight* (see Section 12.2), which accounts for the detector effective area, simulated livetime and the assumed source flux spectrum ($E^{-\gamma}$, see Equation (2.1)).
- **Spatial Distribution:** Injected events are distributed according to the source profile (point-like or extended) and the detector's PSF.
- **Flux Scan:** The mean number of injected signal events n_{inj} is varied to scan different flux levels.

For each n_{inj}, the accumulated trials yield a signal-plus-background TS distribution.

Sensitivity and Discovery Potential Definition

The analysis performance is quantified using two primary metrics derived from these TS distributions.

The first metric is the **sensitivity**, which represents the median[1] 90 % confidence level (C.L.) upper limit on the source flux. It is defined as the signal flux for which 90 % of the signal-injected trials yield a TS value greater than the median of the background-only TS distribution. This threshold characterizes the detector's expected upper-limit performance in the absence of a true signal. The second metric is the **discovery potential**, which quantifies the signal strength required to claim a detection. It is defined as the flux required for 50 % of the signal-injected trials to exceed the 5σ significance threshold of the background distribution. This threshold corresponds to a one-sided p-value of 2.87×10^{-7} and represents a much stricter requirement than the sensitivity, necessitating a significant separation between the signal and background populations.

[1] While the foundational methodology for optimizing sensitivity (model rejection potential) was established in [234] using the average limit, modern analyses typically utilize the median limit to provide robustness against background fluctuations. The confidence intervals are constructed following the unified approach of [235].

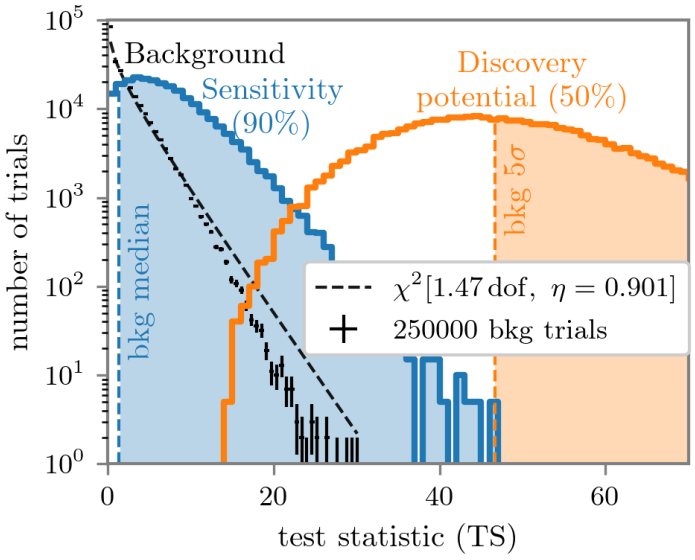


Figure 11.1: Visual representation of the statistical definitions based on data from Chapter 13. (black) Background-only TS distribution derived from scrambled data, overlaid with a χ^2 fit approximation (dashed) used to extrapolate the high-TS tail. (blue) Signal flux injection used to determine the sensitivity, where 90 % of trials exceed the background median (dashed vertical line). (orange) Signal flux injection used to determine the discovery potential, where 50 % of trials exceed the extrapolated background 5σ threshold (dashed vertical line).

As shown in Figure 11.1, the background TS distribution is modeled using the standard csky parameterization: a scaled χ^2 distribution with floating degrees of freedom and a normalization fraction η. While accurately describing the low-TS bulk, this analytic fit overestimates the empirical background high-TS tail (TS $\gtrsim$ 15). This deviation does not affect the sensitivity calculation, as the densely populated background median is read directly from empirical trials. However, identifying the 5σ threshold for the discovery potential requires extrapolating far beyond the typical 1×10^5 pseudo-experiments. Because of this overestimated tail, the resulting 5σ threshold is conservative. Consequently, a larger signal flux is required to claim a detection. Although this conservatism is a byproduct of the analytic fit rather than a deliberate design choice, it increases the resilience against statistical under-fluctuations in the background model.

Both quantities are defined by the physical neutrino flux $\Phi(E)$. Assuming a power law spectrum with a normalization constant Φ_0 at a pivot energy E_0 (typically 1 TeV), the differential flux is:

$$\Phi(E) = \frac{dN}{dE} = \Phi_0 \cdot \left(\frac{E}{E_0}\right)^{-\gamma}. \tag{11.5}$$

The expected number of signal events n_s in the detector is the integral of this flux folded with the detector's effective area A_{eff} (see Section 12.2) over the livetime T_{live}. By substituting the flux model into this integral, the normalization constant can be factored out to define the acceptance $\mathcal{A}$:

$$n_s = T_{\text{live}} \int dE'\, A_{\text{eff}}(E', \delta)\, \Phi(E') = \Phi_0\, T_{\text{live}} \underbrace{\left[\int dE'\, A_{\text{eff}}(E', \delta) \left(\frac{E'}{E_0}\right)^{-\gamma}\right]}_{\mathcal{A}}, \tag{11.6}$$

where δ is the declination of the source and E' is the neutrino energy. The acceptance $\mathcal{A}$ acts as the conversion factor between the physical flux normalization and the expected number of signal events, such that $\Phi_0 = n_s/\mathcal{A}$. Sensitivity and discovery potential are often reported as a differential energy flux in units of $\mathrm{TeV\,cm^{-2}\,s^{-1}}$, with the particle flux scaled by E^2. This E^2-scaling is a standard convention in high-energy astrophysics; it visually flattens a theoretical E^{-2} reference spectrum, allowing the flux to be represented across multiple orders of magnitude without the power law behavior dominating the visualization:

$$E^2\Phi(E) = E^2 \cdot \frac{n_s}{\mathcal{A}} \cdot \left(\frac{E}{E_0}\right)^{-\gamma}. \tag{11.7}$$

However, when quoted as a single benchmark scalar value, the sensitivity is conventionally evaluated at the pivot energy E_0 (often 1 TeV). At this point, the spectral shape term becomes unity and drops out, simplifying the reported scalar value to:

$$\text{Sensitivity}(E_0) = E_0^2 \cdot \Phi(E_0) = E_0^2 \cdot \Phi_0 = E_0^2 \cdot \frac{n_s}{\mathcal{A}}. \tag{11.8}$$

This formulation provides a direct conversion between the required statistical event count threshold n_s obtained from the pseudo-experiments and the corresponding astrophysical flux normalization at the source.

12 Analysis Framework

The performance of a point source analysis depends on the quality of the event selection and reconstruction used to suppress the dominant atmospheric backgrounds. In IceCube, the detector is initially dominated by a flux of atmospheric muons, produced by cosmic ray interactions in the Earth's atmosphere (see Section 2.1.2), triggering the detector at a rate of approximately 2.15 kHz[1] (or $\mathcal{O}(10^{11})$ events per year) [236]. After applying filtering techniques to remove these muons, the dominant remaining background consists of atmospheric neutrinos, observed at rates of several millihertz (or $\mathcal{O}(10^5)$ events per year) at analysis level [106, 237]. In contrast, the astrophysical neutrino signal is rare, occurring at a rate of only $\mathcal{O}(10)$ events per year observed for high-energy starting events (HESE) [238, 239]. Effective background mitigation strategies, such as using the Earth as a shield to filter atmospheric muons as well as applying dedicated event selection and reconstruction techniques, are therefore essential.

This work presents a time-integrated, unbinned likelihood point-source analysis (see Chapter 11) targeting the active galactic nucleus NGC 1068, as motivated in Chapter 10. Section 12.1 provides a brief overview of the data processing chain from raw detector signals to analysis-level datasets and introduces the two complementary event samples used in this study. Thereafter, Section 12.2 summarizes the relevant detector performance metrics and Section 12.3 details the specific analysis configuration.

12.1 Event Samples

The analysis-level datasets used in IceCube analyses are the output of a processing pipeline starting with the detection of photons (originating from direct Cherenkov radiation - Section 3.3, stochastic losses - Section 3.5, or noise - Chapter 9) in the Antarctic ice. When a photon is detected by a PMT, the PMT signal is digitized and time-stamped locally. This raw data is transmitted to the surface for computing, where a data acquisition system implements a series of triggers to identify hit patterns consistent with particle trajectories (e.g., multiplicity triggers). If a global trigger condition is met, the full detector readout within the corresponding time window is stored. Due to bandwidth limitations at the South Pole, an online filtering system processes these events in real-time. Only events passing specific filters are transmitted via satellite to data centers in the Northern Hemisphere. Here, the data undergoes further offline processing ("levels"), where additional algorithms reject background and reconstruct key event parameters such as direction and deposited energy. The final analysis files are stored as lightweight `.npy` arrays containing only the essential reconstructed observables required for the likelihood inference.

The analysis, as introduced in Section 12.3, relies on two distinct event samples: a baseline dataset representing the legacy detector performance (*GRECO*) and a simulation dataset incorporating IceCube Upgrade (*IC93*). The starting point for the following

[1] Although it does not represent a periodic signal, it is standard practice to state event rates in units of Hz instead of 1/s.

analysis is briefly outlined here, while the reader is referred to the provided sources for additional information.

12.1.1 IceCube DeepCore-based Baseline Dataset (GRECO)

The detection of neutrinos in the energy range from a few GeV to tens of TeV is important for studying atmospheric neutrino oscillations and searching for low-energy astrophysical transients. IceCube addresses this regime with the *GRECO* (*GeV reconstructed events with containment for oscillations*) dataset [240]. Originally developed for precision oscillation measurements, this sample was adapted in 2018 into the *GRECO Astronomy* selection to facilitate searches for astrophysical transients in the GeV regime [241]. The GRECO sample includes data from April 26, 2012 to October 28, 2023, corresponding to a total livetime of 11.6 years.[2]

Particle Generation

To train the event selection algorithms, a Monte Carlo (MC) simulation chain is employed. The neutrino signal is simulated using two complementary generators: **GENIE** (generates events for neutrino interaction experiments) [242] is utilized for the low-energy regime ($E_\nu \lesssim 100\,\mathrm{GeV}$), providing a precise treatment of neutrino-nucleus interactions and resonances, while **NuGen** [243, 244] handles the higher energy spectrum and propagation through the Earth. The dominant background of atmospheric muons is modeled using both **CORSIKA** (cosmic ray simulations for kascade) [245] and **MuonGun** [246, 247], specifically optimized for IceCube's depth.

Event Selection and Vetoing

The GRECO selection leverages DeepCore, where the surrounding, sparsely instrumented main optical array acts as an active veto used to reject atmospheric muons [240]. Events are required to have their interaction vertex reconstructed within the DeepCore fiducial volume while depositing minimal charge in PMTs in the outer veto region. This technique reduces the background from penetrating atmospheric muons.

After initial geometric cuts, the final classification is performed using a boosted decision tree (BDT). The BDT is trained on approximately 13 variables characterizing event topology, timing and charge distribution [248]. Events passing the BDT threshold form the final sample, which has a rate of $\sim 4.6\,\mathrm{mHz}$ and consists of roughly 60 % atmospheric neutrinos and 40 % residual muons.

Reconstruction

Events are reconstructed using **Pegleg** [249], a hybrid likelihood-based algorithm. Unlike high-energy reconstructions that assume infinite tracks, Pegleg fits a hypothesis combining a short muon track escaping the vertex and a localized hadronic cascade. The angular uncertainty is estimated per event using a random forest regressor trained on the fit output.

[2] The version utilized in this analysis is version-002-p14.

12.1.2 IceCube Upgrade Simulation (IC93)

To project the physics capabilities of the IceCube detector following the integration of the new IceCube Upgrade strings, a dedicated simulation sample is required. This role is filled by the *IC93* dataset [250] (internally known as *QUESO*: *quick upgrade event selection for oscillations*[3]) which was developed and utilized [106]. This dataset is based on a full simulation of the combined detector array, incorporating multi-PMT optical modules (such as mDOMs and D-Eggs; see Section 3.6). The IC93 selection is a methodological shift from GRECO, replacing traditional likelihood-based methods with a full machine-learning pipeline to handle the complex geometry of the new optical modules.

A caveat regarding the simulated detector geometry must be addressed. While IceCube Upgrade was originally designed to instrument seven new strings (bringing the total array to 93 strings), the final physical deployment only comprises five strings due to technical difficulties. Since the IC93 dataset was generated prior to the final deployment, it reflects the initially planned seven-string geometry. Thus, the projected performance improvements derived from this dataset should be regarded as a slightly optimistic estimate compared with the actual five-string configuration that will be operated.

Simulation and Background Modeling

The dataset generation generally follows the procedure established in Section 12.1.1 but utilizes an updated simulation chain to model IceCube Upgrade's instrumentation.

- **Neutrino Generation:** Neutrino interactions are once again simulated using **GENIE**, following an $E^{-2.5}$ energy spectrum in the energy range 1 GeV to 500 GeV, with vertices distributed in a cylindrical volume surrounding DeepCore, now incorporating the IceCube Upgrade strings.
- **Particle & Photon Propagation:** The propagation of outgoing particles is handled by **PROPOSAL** (propagator with optimal precision and optimized speed for all leptons) [251] for muons and **GEANT4** [169] for hadrons and electromagnetic showers. Subsequently, the Cherenkov photons are propagated through the ice using **CLSim** [252], which models scattering and absorption in the bulk ice[4] and converts photons into PMT pulses.
- **Trigger:** The detector response simulation includes the intrinsic optical module noise model as investigated in Chapter 9. An event is recorded if it satisfies the trigger condition of eight or more local coincidence (LC) pulses within a 1.75 µs time window.
- **Background:** Atmospheric muons are again modeled using **MuonGun**, while pure noise events originating from optical modules (simulated as described in Chapter 9) are mitigated using a graph neural network (GNN)-based noise cleaning algorithm. This algorithm performs a binary classification to predict whether an individual pulse is a physics-induced signal or thermal noise [250]. For IC93, the atmospheric muon fraction is 12.3 %, compared with 44.7 % for GRECO. While the enhanced detection capabilities of IceCube Upgrade's new optical modules will aid event reconstruction, the primary performance gains are attributed to algorithmic advancements. Specifically, the implementation of the GraphNeT architecture enables improved background rejection compared with legacy methods [106, 250].

[3] The specific version utilized in this work is IC93 v29.

[4] The ice model used is spice_bfr-v21 and no hole ice is assumed.

Reconstruction and Classification Pipeline

The processing chain transforms raw triggers into analysis-level events through a multi-stage deep learning pipeline.

Event Filtering: Triggered events first undergo a cleaning stage to discard isolated noise pulses. This is followed by a muon-veto algorithm that rejects events with causal hits in the outer detector consistent with incoming tracks. A series of cuts on variables such as the number of pulses in the fiducial volume is then applied to suppress pure noise triggers before high-level processing.

GNN-based Reconstruction: Surviving events are reconstructed using **DynEdge**, a dynamic GNN within the **GraphNeT** framework [253]. Treating optical module hits as nodes in a graph, the network dynamically learns spatial and temporal correlations to simultaneously reconstruct energy, direction and event topology. This method significantly outperforms traditional algorithms at low energies.

Final Event Selection: The extracted features are fed into two **LightGBM** classifiers (gradient boosting decision trees) trained for strict rejection of atmospheric muons [254]. The resulting sample achieves a neutrino purity of $> 98.5\,\%$ with a rate of $\sim 6.6\,\mathrm{mHz}$ [106].

Systematic Uncertainties

Systematic uncertainties are handled via dedicated MC sets with varied parameters, including optical module efficiencies (treated separately for DOMs and IceCube Upgrade modules) and bulk ice properties (scattering/absorption). Notably, the refrozen *hole ice* for IceCube Upgrade strings is assumed to be clear (bubble-free) due to improved drilling procedures, and thus no hole ice variations are currently simulated [250].

Dataset Limitations

While effective, the IC93 dataset has limitations. The 500 GeV energy cutoff restricts sensitivity to high-energy spectral tails compared with the GRECO selection, which extends to $\sim 10\,\mathrm{TeV}$, thereby omitting higher-energy contributions in direct comparisons. In addition, the mock background sampled in Section 12.3.2 contains only atmospheric neutrinos, as the available atmospheric muon and detector noise simulations lack sufficient statistics to construct reliable background PDFs without introducing oversampling artifacts. Given the high neutrino purity of the event selection, the residual atmospheric muon contamination is treated as negligible. The estimated detector noise rate of 0.2 mHz, together with the atmospheric muon rate of 0.8 mHz, amounts to approximately $15\,\%$ of the atmospheric neutrino rate (6.6 mHz), implying a sensitivity degradation of $\sqrt{1.15} \approx 1.07$ in the background-dominated regime. However, these preliminary rates require validation with in situ data following detector deployment. Finally, the available simulation statistics limit the maximum mock livetime that can be constructed without introducing sampling artifacts to approximately five years.

12.2 Detector Performance Metrics

To evaluate the sensitivity of the likelihood analysis, it is essential to characterize the detector response using metrics derived from Monte Carlo (MC) simulations. In particular,

the effective area determines the signal acceptance, while the angular and energy reconstruction performance directly influences the separation between signal and atmospheric background PDFs (see Chapter 11).

Simulation Weighting

Since IceCube simulations are typically generated using artificial spectra (e.g., power laws) to ensure sufficient statistics across a wide energy range, simulated events must be weighted to reproduce the desired physical flux. This weighting scheme allows for the re-normalization of simulated events to any physical flux model, where two primary quantities are defined:

OneWeight ($w_{\mathbf{one}}$): This weight represents the normalization factor associated with a single simulated event, accounting for the generated phase space and the neutrino interaction probability. It allows for the recovery of the geometric acceptance of the detector, independent of the natural atmospheric flux. Conceptually, w_{one} is proportional to the product of the interaction probability P_{int} and the generation phase space:

$$w_{\text{one}} \approx P_{\text{int}}(E_\nu) \cdot A_{\text{gen}} \cdot \Omega_{\text{gen}} \cdot \Delta E_{\text{gen}}, \tag{12.1}$$

where A_{gen} is the injection area, Ω_{gen} is the solid angle and ΔE_{gen} is the energy range of the simulation.

Atmospheric Weight ($w_{\mathbf{atmo}}$): To compare simulation with physical expectations, the expected event rate (in Hz) is calculated. The atmospheric weight for an event i is derived by combining the normalized OneWeight with the physical total neutrino flux Φ_{total}:

$$w_{\text{atmo},i} = \frac{w_{\text{one},i}}{N_{\text{gen}}} \times \Phi_{\text{total}}(E_i, \theta_i). \tag{12.2}$$

Here, N_{gen} is the total number of generated events (normalized per dataset). The total flux Φ_{total} schematically sums the contributions of electron and muon (anti)neutrinos, weighted by their survival probabilities $P(\nu_\alpha \to \nu_{\text{final}})$ due to oscillations:

$$\Phi_{\text{total}} = \Phi^{\text{atm}}_{\nu_e} P(\nu_e \to \nu_{\text{final}}) + \Phi^{\text{atm}}_{\nu_\mu} P(\nu_\mu \to \nu_{\text{final}}).$$

Effective Area

The effective area quantifies the detector's efficiency in detecting neutrinos. It can be interpreted as the area of an ideal detector with 100 % detection efficiency that would observe the same event rate. Mathematically, it is calculated by summing the normalized OneWeights of all events that pass the selection criteria within a specific bin (defined by energy and declination here).

Figure 12.1 compares the effective area of the GRECO and IC93 datasets. The addition of the IceCube Upgrade strings noticeably increases the low-energy acceptance. Below $\sim 30\,\text{GeV}$, the dense instrumentation of multi-PMT modules allows for the triggering of faint events that remain unidentified in the standard GRECO selection, achieving a gain of over an order of magnitude at $\mathcal{O}(1\,\text{GeV})$. At energies $E_\nu \gtrsim 100\,\text{GeV}$, the IC93 acceptance decreases relative to GRECO. This is an expected consequence of the event selection strategy, as the IC93 cuts are optimized specifically for the low-energy region relevant to oscillation physics, deliberately rejecting higher-energy events that can be better analyzed

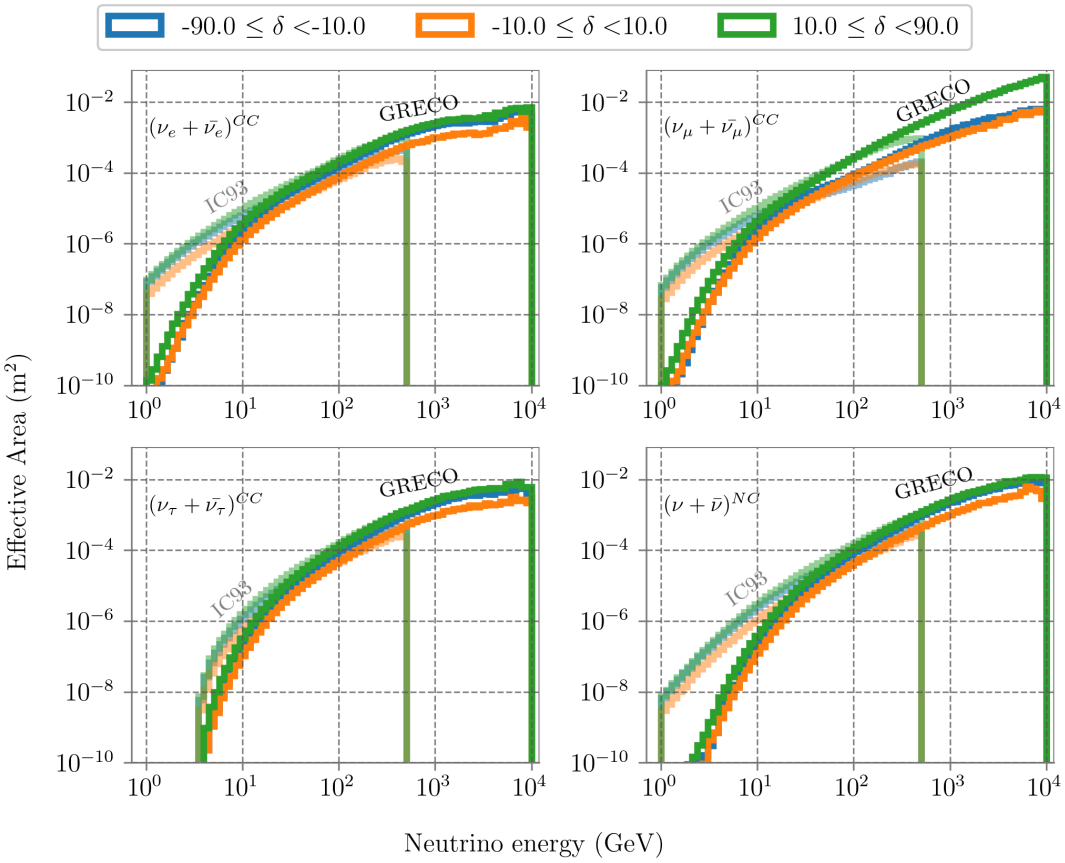


Figure 12.1: Comparison of the effective area from the GRECO (solid) and IC93 (transparent) samples for different declinations (colors) and neutrino interactions (panels). Note that IC93 is simulated only in the energy range 1–500 GeV.

with standard IceCube samples. Note that the values of the effective areas depend on the processing level (see Section 12.1) and are at analysis level in this section.

Typically, detector performance metrics are separated by interaction channel ($(\nu_e + \bar{\nu}_e)^{\mathrm{CC}}$, $(\nu_\mu + \bar{\nu}_\mu)^{\mathrm{CC}}$, $(\nu_\tau + \bar{\nu}_\tau)^{\mathrm{CC}}$ and all neutral current interactions) as done in Figure 12.1. However, in order to streamline this section, the following figures display performance for the combined all-flavor, all-interaction combination $(\nu + \bar{\nu})^{\mathrm{CC+NC}}$.

Angular Resolution

Angular resolution is the primary metric driving point-source sensitivity besides effective area, as it allows for more accurate source localization and background rejection. It is defined as the median angular distance $\Delta\Psi$ between the reconstructed direction $\vec{r}_{\mathrm{reco}}$ and the true neutrino direction $\vec{r}_{\mathrm{true}}$:

$$\Delta\Psi = \arccos\left(\vec{r}_{\mathrm{reco}} \cdot \vec{r}_{\mathrm{true}}\right) = \arccos\left(\sin(\delta_{\mathrm{true}})\sin(\delta_{\mathrm{reco}}) + \cos(\delta_{\mathrm{true}})\cos(\delta_{\mathrm{reco}})\cos(\alpha_{\mathrm{reco}} - \alpha_{\mathrm{true}})\right),$$

where δ and α represent declination and right ascension, respectively. The IC93 dataset provides improved angular resolution, as shown in Figure 12.2 (left). The GNN-based reconstruction consistently outperforms the GRECO dataset across the valid energy range. A crucial feature of low-energy reconstruction is the kinematic limit, indicated by the

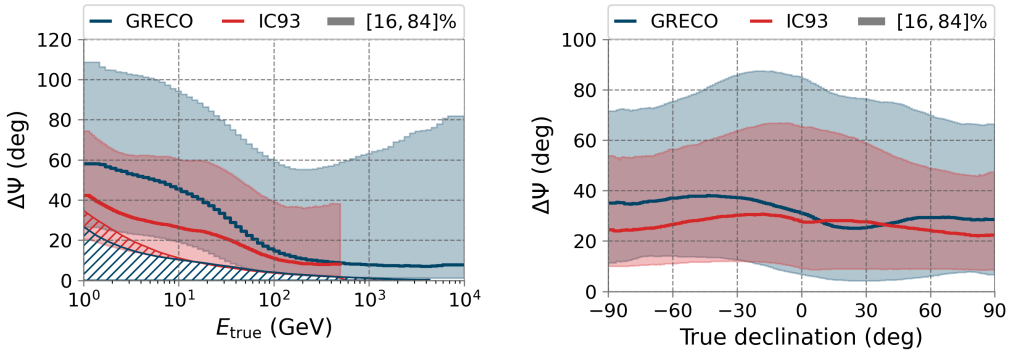


Figure 12.2: Angular resolution evaluated as the median (solid lines) and 68th percentile (bands) of the angular distance between true and reconstructed directions, weighted by atmospheric weights. **(left)** Energy dependence. The hatched regions indicate the kinematic limits evaluated from the retained events. **(right)** Declination dependence.

hatched regions. At energies below $\sim 5\,\text{GeV}$, the physical scattering angle between the incoming neutrino and the outgoing lepton becomes large and represents the main limit to reconstruction accuracy. In this regime, even a perfect detector cannot reconstruct the incoming neutrino direction more accurately than this intrinsic scattering angle. The IC93 reconstruction approaches this fundamental limit near 1 GeV, indicating that the performance of IceCube Upgrade is constrained primarily by neutrino interaction physics rather than instrumentation.

Figure 12.2 (right) illustrates the dependence on declination. Although IC93 generally provides better angular resolution, GRECO exhibits a smaller median angular uncertainty in the 5° to 40° range. This local minimum is an artifact of the differing energy spectra rather than a geometric advantage. Because the GRECO sample includes events up to approximately 10 TeV, it contains a high-energy population of up-going muon tracks from the Northern sky. These high-energy tracks possess excellent pointing resolution, which reduces the median angular error in this specific declination band. When the comparison is restricted to the shared $< 500\,\text{GeV}$ energy range, IC93 consistently achieves superior resolution across the entire sky. This spectral artifact is further amplified by the differing flavor compositions, as muon (anti)neutrinos constitute approximately 77 % of the GRECO sample compared with 65 % for IC93.

A *pull plot* is a diagnostic tool used to assess the calibration of a reconstruction algorithm, typically defined for scalar observables as $(x_{\text{reco}} - x_{\text{true}})/\sigma_{\text{reco}}$. For directional reconstruction, since the angular error is strictly positive, the relative angular error is evaluated via the ratio $\Delta\Psi/\sigma_{\text{reco}}$. A ratio greater than 1 indicates that the algorithm underestimates its true error ($\sigma_{\text{reco}} < \Delta\Psi$), while a ratio less than 1 indicates an overestimation. Figure 12.3 demonstrates the calibration stability of both algorithms. Below $\sim 100\,\text{GeV}$, both reconstructions systematically underestimate their directional error, with the ratio increasing toward the lowest energies due to kinematic scattering limits. In the higher-energy regime ($> 100\,\text{GeV}$), the medians closely track the ideal calibration line. Furthermore, the 1σ quantile bands indicate that the GNN-based IC93 reconstruction maintains a narrower envelope (spanning approximately 0.55 to 1.96 on average) than the GRECO baseline (0.49 to 2.15), demonstrating improved event-by-event stability.

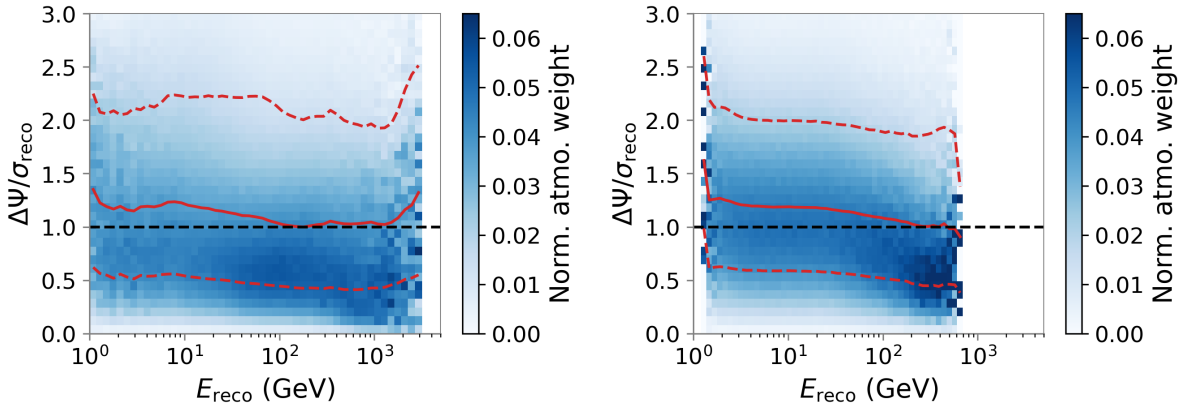


Figure 12.3: Calibration of angular uncertainty estimates. The 2D histograms show the distribution of the relative angular error ($\Delta\Psi/\sigma_{\text{reco}}$) as a function of reconstructed energy. The solid red line indicates the median of the distribution, while the dashed lines represent the 1σ quantile spread. **(left)** GRECO reconstruction. **(right)** IC93 reconstruction.

Energy Reconstruction

The accuracy with which the reconstruction estimates the neutrino energy from the observed Cherenkov light pattern is visualized in Figure 12.4. Ideally, the median reconstruction (solid red line) aligns with the diagonal ($E_{\text{reco}} = E_{\text{true}}$).
The IC93 reconstruction follows the true energy closely, particularly in the region between 10 GeV and 100 GeV. In contrast, the GRECO algorithm exhibits a larger systematic deviation at higher energies, where the median drops far below the diagonal, indicating a tendency to underestimate the deposited energy. To quantify these effects, the fractional energy bias and resolution are shown in Figure 12.5.

The energy bias quantifies whether the reconstruction systematically over- or underestimates the true energy. As shown in Figure 12.5 (left), IC93 maintains a bias close to zero in the 10 GeV to 100 GeV range, whereas both datasets exhibit a negative bias at high energies and a similar positive bias at low energies. The relative energy resolution

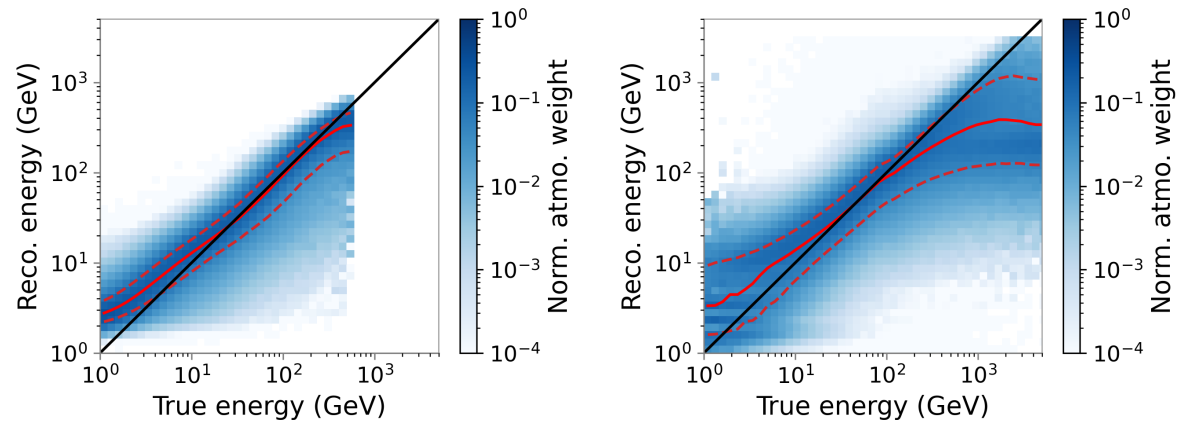


Figure 12.4: Energy reconstruction accuracy. The 2D histograms show the joint distribution of reconstructed against true energy. The solid red line indicates the median reconstructed energy, while the dashed lines represent the 1σ quantile spread. **(left)** IC93 reconstruction. **(right)** GRECO reconstruction.

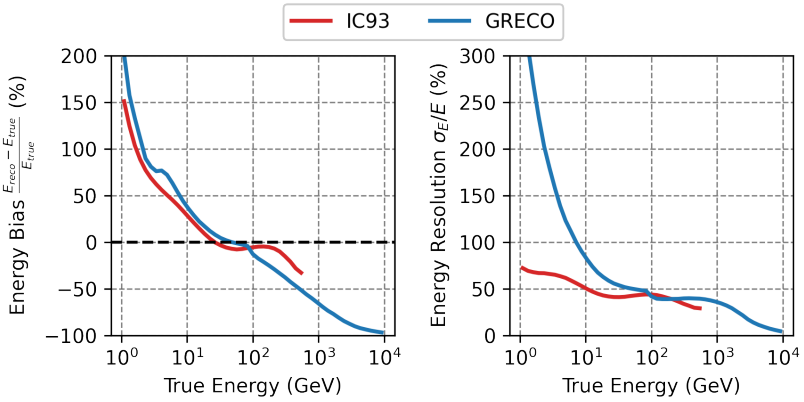


Figure 12.5: Energy reconstruction performance metrics for IC93 (red) and GRECO (blue). **(left)** Fractional energy bias, defined as the median of $(E_{\text{reco}} - E_{\text{true}})/E_{\text{true}}$. **(right)** Energy resolution, defined as half the 68 % interquantile range of the fractional error distribution.

quantifies the precision as the spread of the reconstruction error, defined as:

$$\sigma_E/E = \frac{1}{2}(P_{84} - P_{16}),$$

where P_{84} and P_{16} are the 84th and 16th percentiles of the fractional error distribution $(E_{\text{reco}} - E_{\text{true}})/E_{\text{true}}$. As seen in Figure 12.5 (right), IC93 offers better stability at low energies (<100 GeV), maintaining a resolution of around 50–70 %, whereas GRECO's resolution degrades in this regime. At higher energies (>100 GeV), both algorithms achieve comparable precision.

Event and Background Rates

For experimental data, the observed (background) event rate is determined from the number of selected events passing the *good run list* (GRL)[5] quality criteria divided by the corresponding livetime. For the GRECO selection, this yields an experimentally observed rate of ~ 4.6 mHz, as illustrated in Figure 12.6. This plot also illustrates seasonal variations in the atmospheric component of the background rate, as described in Section 2.1.2.

Since IceCube Upgrade data acquisition has not been integrated at the time of writing this thesis, a direct GRL-based calculation is not possible for IC93. Instead, expected rates are derived from simulation by summing the atmospheric weights. The resulting event rates for both IC93 and GRECO are detailed in Table 12.1. The IC93 selection exhibits a higher atmospheric neutrino event rate (6.6 mHz) compared with GRECO (2.5 mHz). Including atmospheric muons, the total expected background rates are 7.5 mHz and 4.6 mHz, respectively. This increase is primarily driven by the enhanced sensitivity of the new instrumentation to lower energy events and the machine-learning-based algorithms (see Section 12.1.2).

[5]For IceCube, data taking is usually split into 8-hour intervals and validated through a dedicated monitoring procedure. The GRL is a curated catalog of such data taking intervals during which the detector met the stability and quality requirements necessary for physics analysis. This list is specific to each event selection: While a standard high-energy analysis might require the full 86-string array, the GRECO selection utilizes a specialized GRL focused on the stability and noise conditions of the inner DeepCore strings.

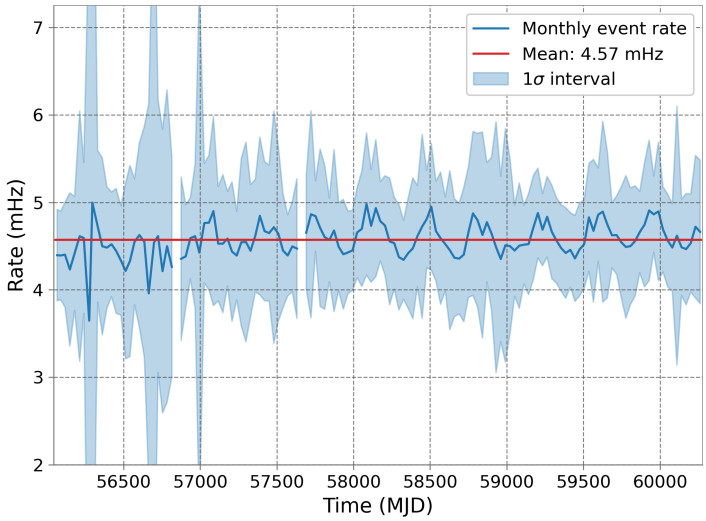


Figure 12.6: Temporal stability of GRECO event rates. The mean rate is 4.57 mHz, comprising contributions from atmospheric neutrinos and the residual atmospheric muon background.

To explore this further, Figure 12.7 presents the distribution of expected events as a function of true energy and declination. These histograms provide insight into the energy reach of the detector and potential directional dependencies. The energy distribution confirms that the IC93 selection extends deeper into the low-energy regime compared with GRECO. IC93 retains a much higher overall rate at energies below 200 GeV. Towards the lower end, electron (anti)neutrinos, which are difficult to reconstruct, constitute a large fraction. The rates for both selections converge for > 300 GeV, indicating similar performance in the regime where the standard IceCube strings dominate. Muon neutrinos constitute the dominant component in both datasets, while tau neutrinos remain rare.

The zenith distribution shows a comparable trend: the angular profiles are morphologically similar, while IC93 exhibits an elevated rate across all angles. Both datasets show a slightly higher background rate for events originating from positive declinations, which is expected because the main optical array is more sensitive to up-going events.

Table 12.1: Comparison of expected atmospheric neutrino rates (in mHz) for IC93 and GRECO selections, broken down by neutrino flavor and interaction type.

IC93

Type	ν (mHz)	$\bar{\nu}$ (mHz)	Total (mHz)
ν_e CC	1.325	0.534	1.859
ν_e NC	0.121	0.033	0.155
ν_μ CC	2.674	1.199	3.873
ν_μ NC	0.341	0.108	0.449
ν_τ CC	0.124	0.053	0.177
ν_τ NC	0.097	0.028	0.125
Total	**4.683**	**1.955**	**6.638**

Note: Expected background rate from MuonGun simulation is 0.82 mHz. Neutrinos thus account for 89.0 % of the total background rate.

GRECO

Type	ν (mHz)	$\bar{\nu}$ (mHz)	Total (mHz)
ν_e CC	0.315	0.132	0.446
ν_e NC	0.028	0.008	0.036
ν_μ CC	1.232	0.536	1.768
ν_μ NC	0.151	0.047	0.198
ν_τ CC	0.051	0.020	0.072
ν_τ NC	0.019	0.005	0.023
Total	**1.796**	**0.747**	**2.543**

Note: The GRL average rate is 4.57 mHz. Neutrinos thus account for 55.6 % of the total rate, implying an atmospheric muon background of ≈ 2.05 mHz.

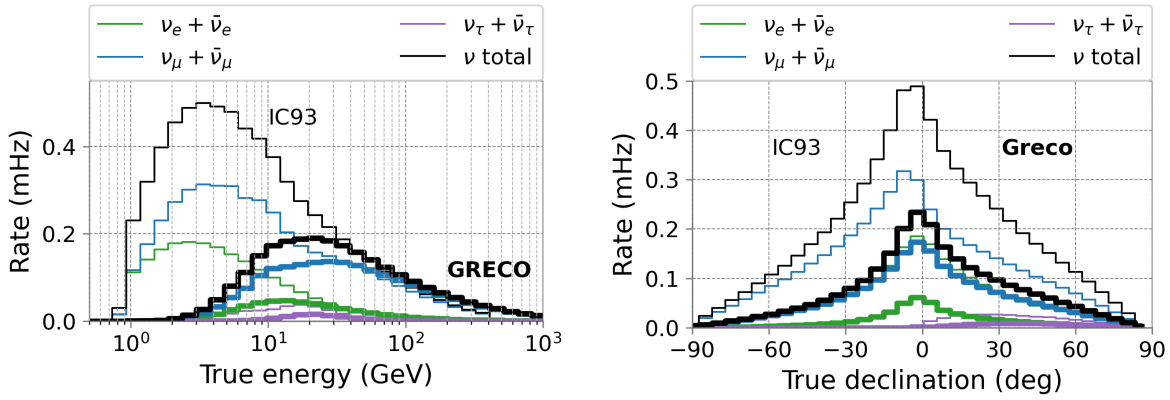


Figure 12.7: Expected event distributions by neutrino flavor. Bold lines show GRECO, while thin lines represent IC93. **(left)** Rate density as a function of true neutrino energy. **(right)** Rate density as a function of true zenith angle (declination).

Probability Density Functions

The final components required for calculating the test statistic are the PDFs that describe the expected distribution of events in spatial coordinates and energy, as introduced in Section 11.1. The spatial PDFs and the energy-dependent signal-over-background (S/B) ratio maps for both datasets are shown in Figure 12.8. The spatial background PDFs (top row) reveal distinct differences in detector acceptance between the two selections. The

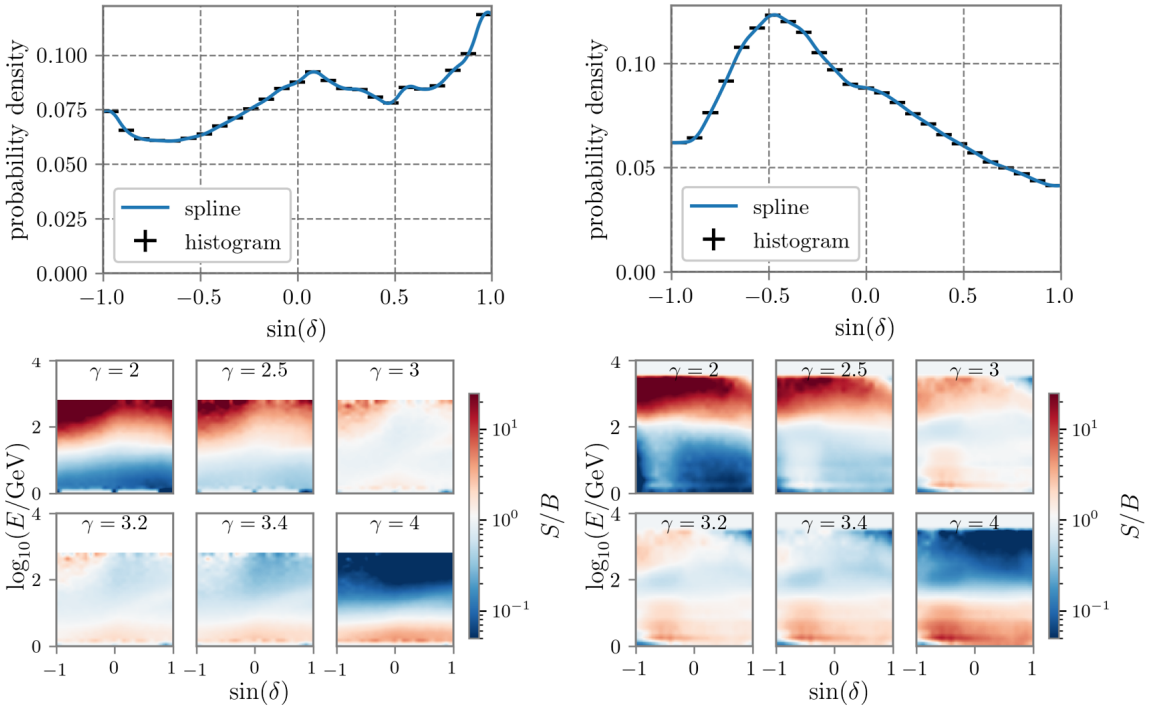


Figure 12.8: Likelihood PDFs for IC93 **(left)** and GRECO **(right)**. **(top)** Spatial background probability density as a function of $\sin(\delta)$. **(bottom)** 2D signal-over-background (S/B) energy weights for different spectral indices γ. The heatmap axes represent declination (x-axis) and reconstructed energy (y-axis), with red indicating regions of higher signal probability.

GRECO selection exhibits a broad acceptance peak in the Southern sky ($\sin\delta \approx -0.5$), which is characteristic of DeepCore selections that utilize the surrounding IceCube strings to veto down-going atmospheric muons. In contrast, the IC93 selection demonstrates a flatter acceptance profile with a notable rise towards the Northern sky ($\sin\delta \to 1$). This indicates that the GNN-based selection maintains a more uniform event acceptance across the sky and is particularly effective at retaining up-going events, which are less contaminated by atmospheric backgrounds.

The S/B energy PDFs (bottom row) visualize the weight assigned to events in the likelihood analysis as a function of reconstructed energy and declination. These maps are dependent on the assumed source spectral index γ. For hard spectra (e.g., $\gamma = 2.0$), the high-weight regions are concentrated at higher energies, where the astrophysical signal dominates over the softer atmospheric background. As the spectrum softens, the weighting becomes more diffuse, reflecting the increased difficulty in distinguishing signal from background at lower energies. For increasingly soft spectra (e.g., $\gamma = 4.0$), high weights shift towards lower energies. Both datasets display distinct high-contrast regions, suggesting effective signal discrimination power in specific areas of the energy-declination phase space.

12.3 Analysis Setup

This section outlines the configuration of the analysis, detailing the construction of the projected observation timelines, the background TS landscape and the final event selection criteria applied to optimize sensitivity.

12.3.1 Overview

This analysis is the first to integrate the IC93 simulation into a csky analysis, the primary objective of which is to quantify the expected sensitivity improvement provided by IceCube Upgrade for the active galactic nucleus NGC 1068 compared with the existing detector capabilities. Furthermore, the aim is to develop a general-purpose event selection and analysis strategy that is applicable across the entire sky, as well as adaptable for long detector livetimes and a variety of studies. To achieve this, the detector performance is projected forward in time to simulate a five-year observation window starting from the assumed IceCube Upgrade commissioning date.

The baseline GRECO dataset contains historical data from 2012 through 2023. IceCube Upgrade deployment has been completed, with integration into the analysis framework initially anticipated around May 2026.[6] As csky is designed to utilize real data, the temporal gap between the existing data and the future analysis period is bridged by a continuous timeline that is constructed by extrapolating the GRECO background rates. Since the dominant background components are approximately steady on these timescales, future years are modeled by resampling the randomized (scrambled) event distributions from the existing GRECO dataset. Two distinct scenarios are defined for the five-year projection period (May 2026 – May 2031), as illustrated in Figure 12.9:

[6]Historically, season transitions have occurred around May. While recent IceCube Upgrade data is saved for offline processing, the timeline for data integration of the various new instruments is currently pending. As this analysis was initialized in 2024, it solely relies on simulation for modeling IceCube Upgrade.

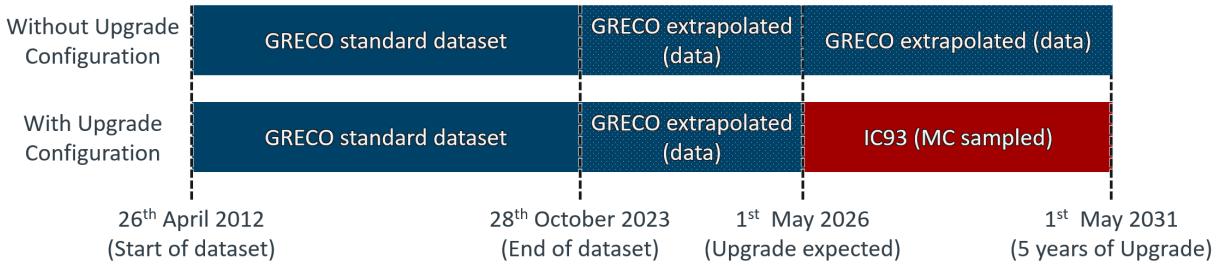


Figure 12.9: Timeline of the sensitivity projection scenarios. Both scenarios use identical historical GRECO data (blue) up to the expected commissioning date in May 2026. The *without Upgrade* scenario (top) projects the continued operation of the legacy detector, while the *with Upgrade* scenario (bottom) transitions to the IC93 simulation (red) for the subsequent five-year observation period.

- **without Upgrade**: Represents the continued operation of the legacy detector configuration. This scenario uses the extrapolated GRECO dataset for the entire future window, resulting in a total livetime of roughly 19 years (2012 – 2031) of legacy detector performance.
- **with Upgrade**: Represents the combined dataset of the legacy detector followed by the inclusion of IceCube Upgrade. The timeline consists of the historical and extrapolated GRECO dataset up to the commissioning cutoff (May 2026), after which it transitions to the IC93 sample for the subsequent five years. This yields a hybrid dataset combining $\sim$ 14 years of GRECO data with five years of IC93 predictions. These datasets are combined into a single csky analysis object, allowing for seamless likelihood maximization across the transition between detector configurations.

The two datasets exhibit complementary kinematic strengths. IC93 provides high purity and angular resolution in the low-energy regime, but is restricted by a simulation ceiling of $\sim$ 500 GeV. While truncating the historical GRECO dataset to match this upper bound might appear methodologically consistent for a direct comparison, doing so would artificially restrict the physics reach of the analysis. Even soft spectrum sources such as NGC 1068 exhibit emission spectra that extend well into the multi-TeV regime (see Figure 10.2). Because the atmospheric background falls off steeply at higher energies, even a sparse number of TeV-scale events carry substantial statistical weight in the likelihood maximization. Therefore, the GRECO dataset is utilized in its entirety to preserve the detector's sensitivity to the high-energy tail of the astrophysical flux. It should be noted, however, that extending the IC93 simulation beyond 500 GeV would similarly enhance the sensitivity of the "with Upgrade" scenario, as explored in Chapter 13.

12.3.2 Dataset Generation and statistical Framework

To construct the continuous background timelines described in Figure 12.9, a hybrid statistical approach is employed for both the historical data and the future simulations.

The experimental GRECO dataset spans the period from April 2012 to October 2023. To cover the remaining timeline until the defined start of IceCube Upgrade (May 1, 2026; MJD 61161), the existing data is extrapolated. The missing livetime ($\Delta T \approx 2.5$ years) is populated by resampling events from the stable 2012–2023 configuration. A random subset of events is drawn from the historical pool to fill the gap. This data-driven method preserves the physical characteristics and correlations of the background but relies on a

single statistical realization. A corresponding entry is appended to the GRL to ensure a continuous timeline for csky.

For IceCube Upgrade (IC93), no experimental data exists at the time of writing. Consequently, "mock data" is generated from Monte Carlo simulations to represent the expected atmospheric neutrino background. The generation procedure employs Poisson sampling of the simulation's expected event counts, which are derived from the simulation's atmospheric weights, to introduce realistic statistical fluctuations:

1. The total livetime T is defined for the projection period (e.g., five years).
2. For each Monte Carlo event i with atmospheric weight $w_{\text{atmo},i}$ (as defined in Section 12.2), the expected number of occurrences is calculated as $\lambda_i = w_{\text{atmo},i} \cdot T$.
3. The number of times n_i the event appears in the pseudo-data is drawn from a Poisson distribution: $n_i \sim \text{Poisson}(\lambda_i)$.
4. The corresponding simulated event is included n_i times in the pseudo-data sample.

This process yields a realization whose statistical properties reproduce the expected atmospheric background while accounting for Poisson variance. A synthetic GRL is created to define this period as a single continuous run starting from MJD 61161, contrasting with the typical segmentation of experimental data into 8-hour runs. The method developed here is applied in [255], co-written by the author, which features two additional analyses utilizing the IC93 dataset. Unless otherwise specified, the sensitivity projections in this analysis assume detector livetimes extending to 2031 (five years after the expected IceCube Upgrade deployment), using 250 000 background trials and 50 000 signal-injection trials.

12.3.3 Analysis Cuts

An event selection defines the balance between retaining astrophysical neutrino candidates and suppressing the overwhelming atmospheric background. As this analysis spans a projected livetime of 19 years (2012–2031), careful event selection is essential. Over such long observation periods, the massive accumulation of background events populates the non-Gaussian heavy tails inherent to low-energy PSFs. The random spatial coincidence of these extreme outliers can mimic genuine astrophysical clustering, generating an excess of high-TS values in the background-only distribution. Such deviations from the asymptotic χ^2 expectation increase the 5σ discovery threshold and reduce sensitivity. Therefore, the selection criteria suppress these PSF tails to ensure the background TS follows the asymptotic χ^2 expectation while maximizing full-sky sensitivity.

For the IC93 dataset, which relies entirely on simulation, the selection is designed to match the valid kinematic range of the underlying Monte Carlo production while rejecting events with poor angular resolution that destabilize the likelihood fit. The following cuts are applied:

- **Angular Uncertainty:** $\sigma_{\text{reco}} < 13^\circ$.
  This quality cut is the primary driver of analysis stability. As illustrated in Figure 12.10, including events with large estimated angular uncertainties leads to increased deviations from the asymptotic χ^2 expectation in the background-only TS distribution. Specifically, the inclusion of these poorly reconstructed events generates an artificially high-TS tail that mimics signal behavior. Applying the stricter 13° cut suppresses this contribution and restores the asymptotic χ^2 expectation, which is required for a valid statistical interpretation of the analysis. The right

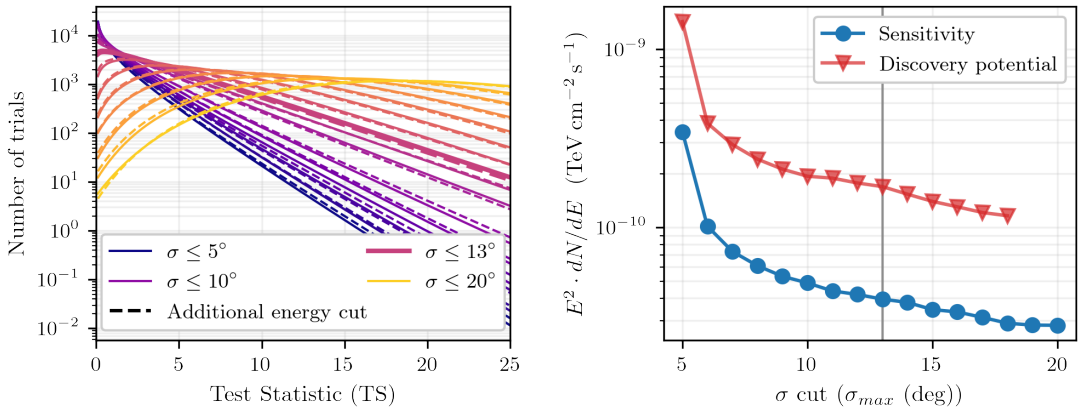


Figure 12.10: Impact of event selection on the statistical performance of the IC93 dataset. **(left)** Evolution of the background-only TS distributions. The curves represent the χ^2 fits for angular uncertainty cuts between 5° to 20°. Solid lines indicate the standard selection, while dashed lines show the effect of the additional energy cut. **(right)** Sensitivity and 5σ discovery potential for the source coordinates of NGC 1068, assuming a spectral index of $\gamma = 3.4$.

panel illustrates the trade-off between signal retention and background rejection as a function of the σ cut, highlighting the optimal sensitivity range for this source.

- **Reconstructed Energy:** $0.30 < \log_{10}(E_{\text{reco}}/\text{GeV}) < 2.69$.
  This range ($\approx 2\,\text{GeV}$ to $500\,\text{GeV}$) corresponds approximately to the reconstructed-energy range covered by the GENIE simulation. It removes outliers where the GNN reconstruction (trained on this specific range) becomes unreliable and statistics are limited, resulting in distinct hotspots in the signal-over-background energy PDFs shown in Figure 12.8.

As mentioned before, to avoid oversampling artifacts during the generation of pseudo-data, the sub-dominant atmospheric muon component (simulated via MuonGun [246]) is excluded. The background model for IceCube Upgrade's operation is thus constructed solely from the high-purity neutrino set. This selection retains approximately 20.3 % of the nominal signal sample, corresponding to 8.6 % when weighted according to an $E^{-3.4}$ spectrum, while retaining 17.6 % of the atmospheric background. The cuts are designed to maintain sensitivity across the accessible energy range while avoiding contamination and instability from poorly reconstructed or high-background events.

For the GRECO dataset, which spans a broader energy range, the same approach is taken. While this study initially employed a loose angular uncertainty cut ($\sigma_{\text{reco}} < 25°$) producing reliable sensitivities, a stricter selection from [256] was adopted to improve the stability of the background description across the sky. The final cuts are:

- **Angular Uncertainty ($\sigma_{\text{reco}} < 7°$):** Serving the same primary stabilizing function as in the IC93 dataset, this cut truncates the heavy tails of the point spread function.
- **Topology Identification (`pidDeltaLLH` < -10):** This parameter evaluates the log-likelihood difference between the hybrid track-plus-cascade reconstruction (Pegleg) and a reconstruction assuming a cascade-only hypothesis (Monopod). Enforcing a negative value preferentially selects track-like events, which generally provide better directional reconstruction than approximately spherical hadronic or electromagnetic cascades.
- **Direct Light Charge (`qdir` > 2.5):** This metric quantifies the total charge

recorded from "direct hits": photons arriving at the optical modules with little to no scattering. Requiring a high direct-charge fraction ensures the directional reconstruction is driven by prompt light, mitigating the degrading effects of photon scattering in the bulk ice.

The angular uncertainty cut is the dominant selection criterion, rejecting approximately 90 % of the initial sample while serving the same purpose as for the IC93 dataset. The supplementary quality cuts provide additional topological and timing-based selection criteria, removing a further $\sim$ 5 % of the remaining events. While the total reduction is large, it primarily removes events with little to no pointing information that contribute negligible weight to the point-source likelihood. As demonstrated in [256], this selection achieves sensitivity comparable to looser selections while reducing the occurrence of background-induced hotspots, particularly at southern declinations where atmospheric muon contamination is highest. The final retention rates are 14.4 % for signal (3.6 % when weighted to an $E^{-3.4}$ spectrum) and 8.9 % for background.

Figure 12.11 demonstrates that the optimized event selections produce a stable background landscape across all declinations. The TS distributions remain within the expected range for background fluctuations, indicating that the transition between the two detector configurations does not introduce systematic artifacts. This validates the setup for the final sensitivity projections.

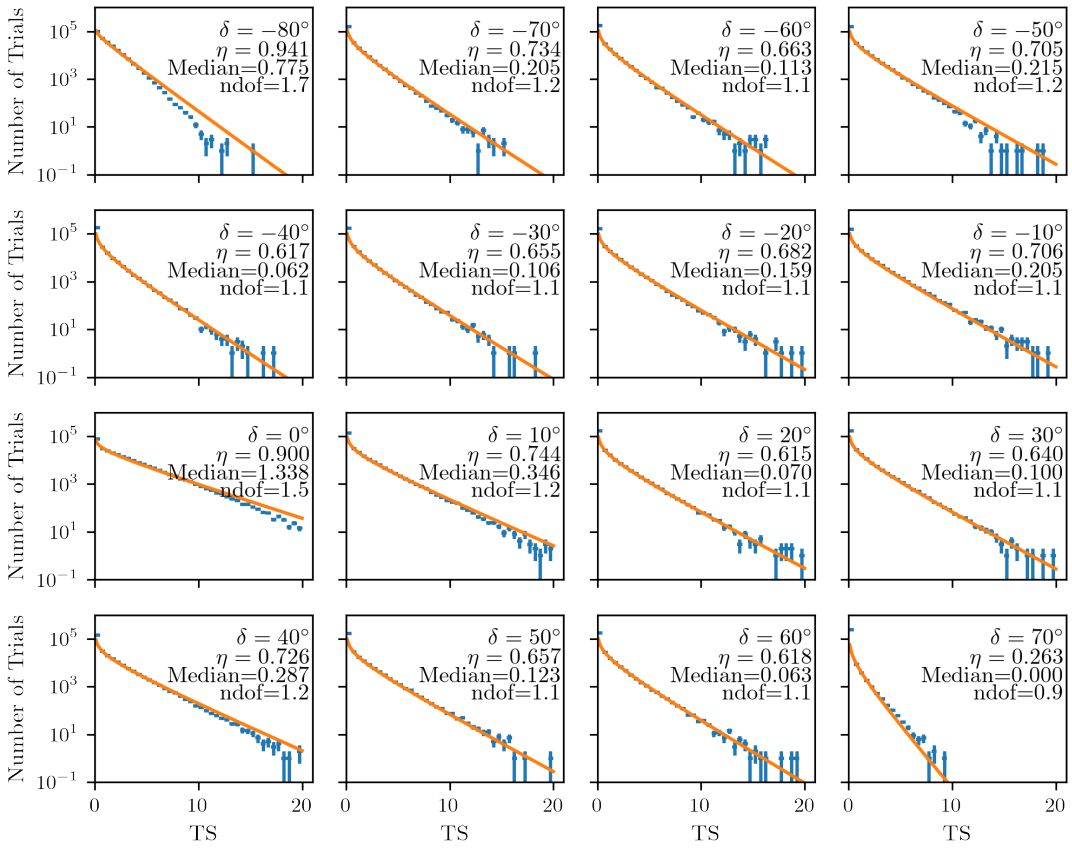


Figure 12.11: All-sky background TS distribution stability for the "with Upgrade" scenario (19 years). Each panel shows the TS distribution obtained from background-only pseudo-experiments for different declinations.

Sample Verification

To ensure that the dataset generation from Section 12.3.2 introduces no statistical bias, the sampling methods are validated by comparing the generated pseudo-data against their unsampled dataset distributions. This verification is performed for both the extrapolated GRECO baseline and the IC93 simulation.

IC93 Validation: The IC93 pseudo-data is generated by Poisson-sampling the atmospheric weights of the Monte Carlo simulation. Figure 12.12 (left) demonstrates that the resulting energy distribution perfectly reproduces the shape of the weighted MC expectation. This confirms that the Poisson fluctuations are correctly implemented and that the pseudo-data preserves the spectral features of the underlying physics model.

GRECO Validation: The GRECO extrapolation relies on resampling events from the 2012–2023 period. As shown in Figure 12.12 (center), the resampled future dataset extends the available historical sample while preserving its underlying distributions. This demonstrates that the bootstrapping procedure successfully conserves the background kinematics without introducing artificial distortions.

Scenario Comparison: The impact of IceCube Upgrade is visualized in Figure 12.12 (right). The "without Upgrade" scenario, relying solely on GRECO, exhibits a spectral cutoff below $\sim 10\,\mathrm{GeV}$ after applying the final cuts from Section 12.3.3. In contrast, the "with Upgrade" scenario shows a significant increase in event rate in the $2\,\mathrm{GeV}$–$20\,\mathrm{GeV}$ range. This explicitly visualizes the gain discussed in Section 12.2: IceCube Upgrade opens a new observation window at low energies that is inaccessible to the current detector configuration.

While Figure 12.12 displays the reconstructed energy, this validation was performed for all primary analysis observables (e.g., declination or angular uncertainty). In all cases, the sampled distributions were found to be consistent with their parent distributions, validating the datasets for use in sensitivity projections.

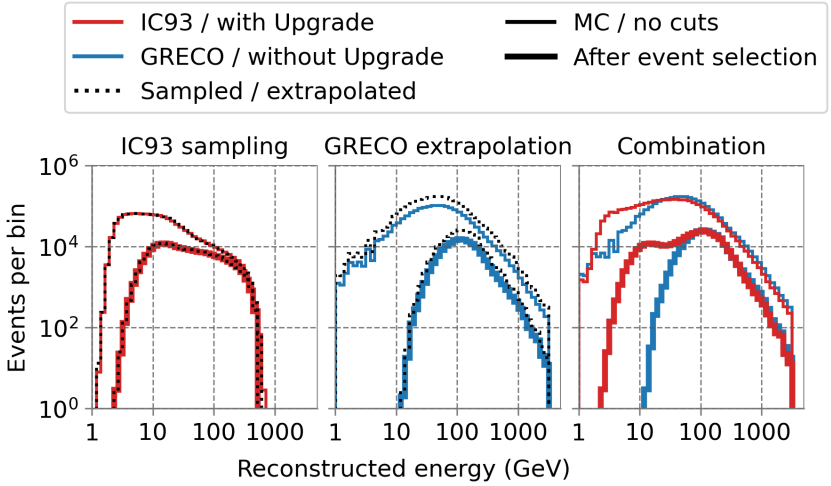


Figure 12.12: Validation of dataset generation using the reconstructed energy distribution. All distributions are shown before and after applying the final cuts (bold). **(left) IC93:** Comparison between the weighted Monte Carlo expectation (solid) and the Poisson-sampled pseudo-data (dotted). **(center) GRECO:** Comparison between the historical experimental data (solid) and the bootstrapped extrapolation (dotted). **(right) Scenario Comparison:** Combined energy spectra for the "without Upgrade" (blue) and "with Upgrade" (red) analysis scenarios.

Dataset Detector Performance Metrics after Analysis Cuts

The event selection described in Section 12.3.3 alters the statistical properties of the analysis samples. By prioritizing reconstruction quality and background stability, a significant fraction of the initial event stream is discarded, favoring the density of high-confidence signal candidates over raw event rates. Exemplary resulting performance metrics, which represent the actual input to the likelihood analysis, are summarized in Figure 12.13.

The impact of the selection is most prominent in effective area (Figure 12.13, left), for which the GRECO dataset shows suppressed sensitivity below $\sim 10\,\mathrm{GeV}$. In contrast, the IC93 selection retains robust acceptance down to $\sim 2\,\mathrm{GeV}$. This reinforces that the primary advantage of IceCube Upgrade in this analysis is the enhanced sensitivity in the low-energy regime, where the baseline configuration has limited acceptance.

While the inherent distributions (Section 12.2) showed a significant resolution lead for the IC93 GNN-based reconstruction, the post-selection angular resolution (Figure 12.13, right) exhibits convergence between the two samples. Especially in the 10 GeV to 500 GeV range, the resolutions are much closer than in the initial comparison. This is a direct consequence of the selection process: by filtering out cascade-like topologies and poorly reconstructed tracks, both datasets are reduced to samples dominated by well-pointing events. Notably, the GRECO resolution appears marginally better in several energy bins. This difference is partly attributed to the stricter angular cut ($\sigma < 7°$ vs. $13°$ for IC93), but it also reflects a composition bias, as the surviving GRECO sample maintains a higher fraction of high-energy muon (anti)neutrinos, which possess inherently sharper pointing signatures. However, the IC93 dataset remains the better choice for low-energy point-source searches, as it maintains this resolution while providing an effective area over an order of magnitude larger than the baseline in the overlapping regime.

The trade-off for such high-purity selections is a reduced total event rate. The expected atmospheric neutrino rates drop to 0.378 mHz for IC93 and 0.124 mHz for GRECO. The final samples are dominated by ν_μ CC interactions (82.4 % for IC93 and 99.1 % for GRECO, comprising approximately 60 % neutrinos and 40 % antineutrinos in both cases), ensuring that the likelihood maximization is driven by events with strong directional discrimination.

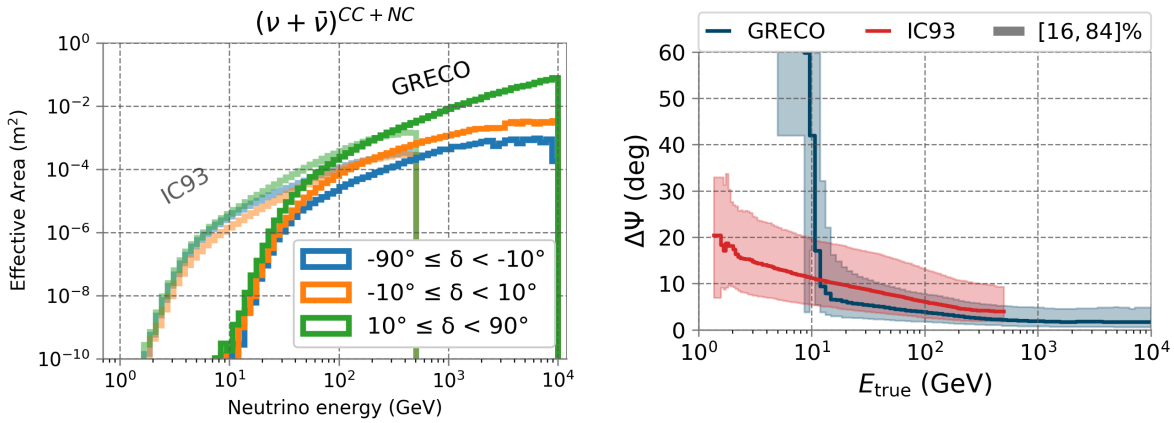


Figure 12.13: Key performance metrics for the final analysis samples after applying cuts. **(left)** Effective area. **(right)** Angular resolution as a function of energy.

13 Results

Building on the analysis methods and framework introduced in the previous chapters, this chapter evaluates the expected impact of IceCube Upgrade on astrophysical neutrino searches. The projected performance is quantified using point-source sensitivity as the primary performance metric. Unless stated otherwise, the reported sensitivities correspond to the median 90 % confidence level (C.L.) upper limit (per-flavor neutrino and anti-neutrino averaged), assuming a terrestrial flavor equipartition ($\nu_e : \nu_\mu : \nu_\tau \approx (1:1:1)_\oplus$) consistent with standard neutrino oscillations over cosmic distances (see Section 2.3.2).

The results presented here contribute to the publication [255], co-authored with Y. Kobayashi and B. Schlüter. While the collaborative work also encompasses time-dependent (transient) and Galactic Plane steady-state searches, this thesis focuses exclusively on the author's primary contributions: the full-sky steady-state analysis and the targeted sensitivity study of NGC 1068.

For consistency, all figures follow a unified graphical scheme: blue lines denote the baseline (without Upgrade) scenario, while red lines represent the detector performance after the inclusion of IceCube Upgrade (with Upgrade). Unless stated otherwise, scenarios are projected to May 2031, corresponding to a five-year observation period following the assumed start of IceCube Upgrade operation in May 2026 (see Figure 12.9). While current simulations rely on datasets incorporating solely IceCube and DeepCore, the future integration of experimental IceCube Upgrade data, together with dedicated simulations, is expected to further refine these performance estimates. The term *legacy detector* or *baseline* strictly refers to the IceCube + DeepCore configuration for the 2012–2026 period, noting that this terminology may evolve with the deployment of IceCube-Gen2 (see Section 3.7).

13.1 Projected Upgrade Performance across the Sky

This work establishes the first point-source sensitivity benchmark for IceCube Upgrade. By optimizing for stability across all declinations (Section 12.3.3), it provides a foundation for future catalog searches once experimental Upgrade data becomes available to complement the IC93 simulations.

To probe the impact of IceCube Upgrade, the sensitivity is evaluated across the entire sky in declination bands of 10°, ranging from −80° to 80°. In each band, the background is estimated using right ascension scrambling. Point source emission is modeled with power law spectral indices of $\gamma \in \{2.0, 2.5, 3.0, 3.2, 3.4, 4.0\}$. A critical technical consideration is the transition between the IceCube Upgrade simulation sample, which is limited to $E_{\text{true}} < 500\,\text{GeV}$, and the GRECO dataset, which extends to approximately 5 TeV.[1] To assess the impact of the 500 GeV energy cutoff inherent to the IC93 Monte Carlo, two configurations are tested:

[1] In the context of this chapter, "high-energy" refers to the upper end of GRECO events, distinct from the PeV scale typically associated with IceCube.

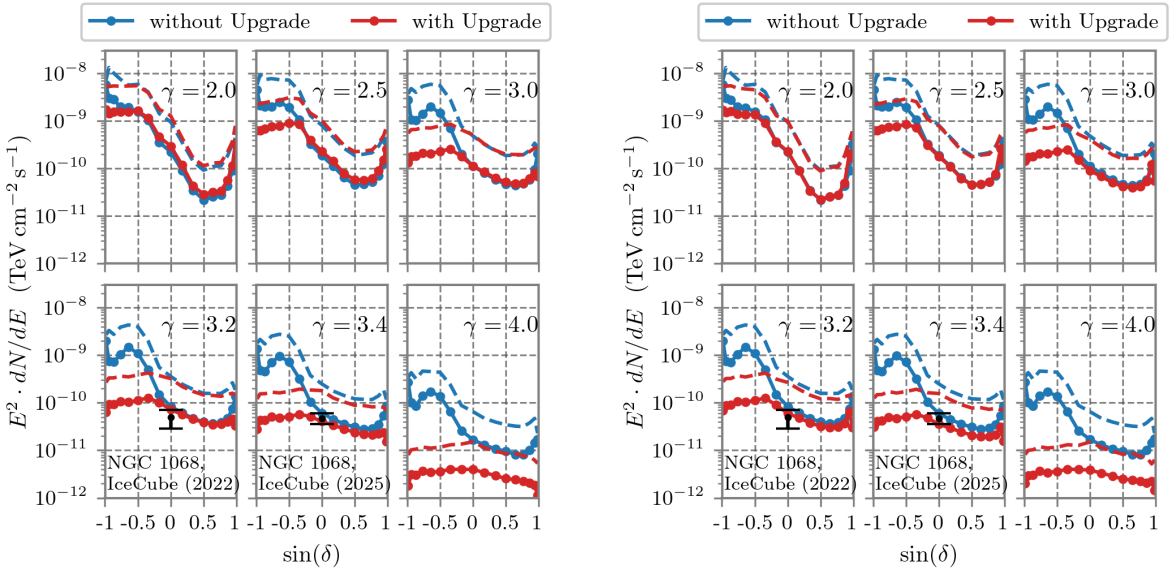
1. A conservative configuration that respects the dataset boundaries (with IC93 being truncated at 500 GeV).
2. An extended configuration that additionally extrapolates the GRECO MC events for $E_{\text{true}} > 500\,\text{GeV}$ to resolve the energy discrepancy.

The resulting sensitivities, calculated via the likelihood method described in Section 11.3, are summarized in Figure 13.1. A clear difference in performance is observed depending on the spectral hardness of the source:

Northern sky ($\sin\delta > 0$): In the Northern Hemisphere, the Earth acts as a natural shield against atmospheric muons, allowing for a cleaner neutrino selection. Here, the impact of IceCube Upgrade is most pronounced for soft spectra ($\gamma \gtrsim 3.2$). The decreased energy threshold recovers the soft tail of the signal flux that would otherwise fall below the legacy detector's detection threshold.

Southern sky ($\sin\delta < 0$): In the Southern Hemisphere, the background is dominated by down-going atmospheric muons. Despite this high-background environment, IceCube Upgrade demonstrates sensitivity improvements even for harder spectra, which are driven by three key factors:

- Denser Instrumentation: The reduced inter-string spacing allows for better resolution of low-energy events, together with enhanced transverse sampling.
- Multi-PMT Modules: The increased directional information from segmented optical modules improves photon detection uniformity and reconstruction performance.



(a) Standard dataset cutoff

(b) Extrapolated high-energy continuation

Figure 13.1: Projected 90 % C.L. sensitivity (solid lines) and 5σ discovery potential (dashed lines) as a function of source declination δ for different spectral indices γ, derived from 250 000 background and 50 000 signal trials. The black error bars indicate the best-fit flux reported for NGC 1068 [6, 226]. **(a)** Sensitivity derived using the standard definition, where the IC93 dataset is truncated at 500 GeV. **(b)** Sensitivity derived by extrapolating the GRECO dataset to extend the upper-energy threshold stemming from IC93 beyond 500 GeV after IceCube Upgrade inclusion.

- Advanced Reconstruction: By leveraging IceCube Upgrade's denser instrumentation and multi-PMT optical modules, advanced reconstruction techniques (see Section 12.1.2) can be applied to improve event resolution and atmospheric muon vetoing.

While the Southern sky enhancement is important for Galactic searches (e.g., the Galactic Center at $\delta = -29°$), it is equally applicable for extragalactic sources located in the south, such as the Circinus Galaxy or NGC 7582 [257]. For southern sources, improvements of up to an order of magnitude are projected depending on declination and spectral index. For a source near the celestial equator such as NGC 1068 ($\delta \approx -0.01°$), the sensitivity gain is evident for the observed spectral index of $\gamma = 3.4$ and becomes larger for any softer flux component.

The importance of correctly modeling the dataset transition is highlighted in Figure 13.1. In the left panel (500 GeV *with Upgrade* cutoff), the Northern sky sensitivity for hard spectra ($\gamma = 2.0, 2.5$) appears to degrade with IceCube Upgrade. This is an artifact of the simulation constraints: the IC93 dataset limits injected neutrinos to $E < 500\,\mathrm{GeV}$, effectively removing the contribution of high-energy events that carry higher statistical weight for hard spectra. Since the background remains constant while the signal is truncated, the sensitivity worsens relative to the baseline scenario. While this conservative approach was adopted for the published results in [255] to adhere to validation protocols, it represents a lower bound on performance.

Conversely, the right panel illustrates a more accurate scenario where the GRECO MC (using true neutrino energy information to avoid double counting) is extrapolated to cover the energy range. With this, the unexpected degradation is no longer apparent and the sensitivity curves for both scenarios at $\gamma = 2.0$ and 2.5 become consistent. Based on this validation, the extrapolated methodology is applied to the remainder of this chapter unless specified otherwise.

13.2 Sensitivity Improvements for NGC 1068

As outlined in Chapter 10, a central focus of this analysis is the active galactic nucleus NGC 1068, the current benchmark source for extragalactic neutrino astronomy. While initial studies characterized the source with a spectral index of $\gamma = 3.2$ [6], the most recent follow-up analysis indicates a softer spectrum of $\gamma = 3.4$ [226]. Given IceCube Upgrade's enhanced effective area at low energies, this spectral softening makes the source an increasingly favorable target. Thus, the results presented here utilize the current $\gamma = 3.4$ best-fit spectral index.

13.2.1 Temporal Evolution and accumulated Gain

The projected sensitivity as a function of livetime is shown in Figure 13.2. The comparison between the *without Upgrade* and *with Upgrade* scenarios exhibits a steeper slope for the latter, indicating that IceCube Upgrade accumulates statistical significance at a faster rate.

A consequence of this steeper slope is the reduction in observation time required to probe the source flux. As illustrated, the *with Upgrade* sensitivity reaches below the

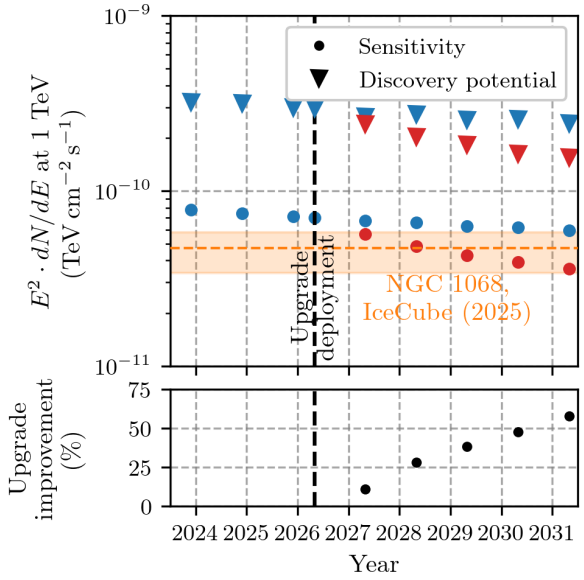


Figure 13.2: Temporal evolution of the projected sensitivity for NGC 1068, assuming a spectral index of $\gamma = 3.4$, derived from 250 000 background and 50 000 signal trials. The sensitivity is compared with the best-fit flux reported in [226].

NGC 1068 best-fit flux faster than the *without Upgrade* scenario. While the legacy detector configuration only approaches the upper bounds of the measured flux for most of the projection time period, the detector crosses this threshold quickly after the inclusion of IceCube Upgrade. It must be noted that this projection assumes the NGC 1068 flux follows a single, unbroken power law behavior down to the lower energies probed by IceCube Upgrade, extending well below the $\mathcal{O}$(TeV) to $\mathcal{O}$(PeV) energy range where the flux has been experimentally constrained by the legacy array. Under this assumption, the improved sensitivity transitions the analysis from a borderline detection regime into a high-statistics measurement regime, enabling earlier access to the signal regime where spectral characterization becomes possible than what is achievable with the legacy detector. Quantitatively, the inclusion of Upgrade data corresponds to an average annual sensitivity improvement of approximately 12.2 %. Over the five-year post-Upgrade analysis period, this accumulates to a total relative improvement of roughly 58 %.

This temporal behavior is driven by the increasing portion of IceCube Upgrade data within the total sample. In the first year (2026–2027), the analysis is dominated by 14 years of historical data (2012–2026), restricting the impact of the new instrumentation. As operations continue, the IceCube Upgrade dataset constitutes a progressively larger fraction of the total livetime, causing the relative improvement to rise from $\sim$ 10.8 % in the first year to $\sim$ 57.8 % after five years[2]. Even under the most conservative assumption (truncating the dataset at 500 GeV), the improvement remains consistent at $\sim$ 7.2 % per year, confirming that the gain is stable and not solely an artifact of high-energy modeling. It should be noted that sensitivity projections become increasingly less accurate the further a given dataset is extrapolated into the future, for which statistical fluctuations

[2]Essentially, IceCube Upgrade "dilutes" the limitations of the legacy dataset. While the absolute sensitivity improves as $\propto 1/\sqrt{T}$ for both scenarios (background-limited regime), the *with Upgrade* scenario's prefactor is smaller. As T increases, the combined dataset converges closer to the pure performance including IceCube Upgrade, widening the gap relative to the baseline.

are expected, especially in the discovery potential (requiring a p value of 2.87×10^{-7}).

It is important to clarify why IceCube Upgrade's sensitivity limit involves detecting *more* signal events ($n_s = 97.74$) than the legacy detector ($n_s = 38.15$). This is not an artificial inflation of performance, but a consequence of the low-energy signal-to-noise environment. Because IceCube Upgrade is sensitive to lower energies where the atmospheric background rate is higher, a relaxed angular uncertainty cut ($\sigma < 13°$) is utilized to increase the effective area. As a result, a larger absolute number of signal events is statistically required to reject the background hypothesis. However, due to this increase, IceCube Upgrade accumulates these events at a lower physical source flux than the legacy detector accumulates its ~ 38 events. This illustrates a qualitative change in the analysis regime at the sensitivity threshold: IceCube Upgrade balances a lower energy threshold and larger effective area with a higher background environment, requiring the likelihood framework to extract signal information from a larger event sample.

13.2.2 Sensitive Energy Range

To physically characterize these improvements, the kinematic regime driving the sensitivity is identified. This *sensitive energy range* is defined as the central 90 % interval of the expected signal distribution, derived from the Monte Carlo simulation by constructing the cumulative distribution function (CDF) of the true neutrino energy. The expected number of detector events (the absolute weight w_i) contributed by a single simulated neutrino over the analysis livetime T_{live} is defined as:

$$w_i = w_{\text{one},i} \cdot \frac{T_{\text{live}}}{N_{\text{gen}}} \cdot \Phi_0 \cdot \left(\frac{E_{\text{true},i}}{E_0} \right)^{-\gamma},$$

where N_{gen} is the total number of generated simulation events. However, because constructing a CDF involves normalizing the distribution by the total sum of all weights, constant scalar factors (namely T_{live}, N_{gen} and the absolute flux normalization Φ_0) cancel

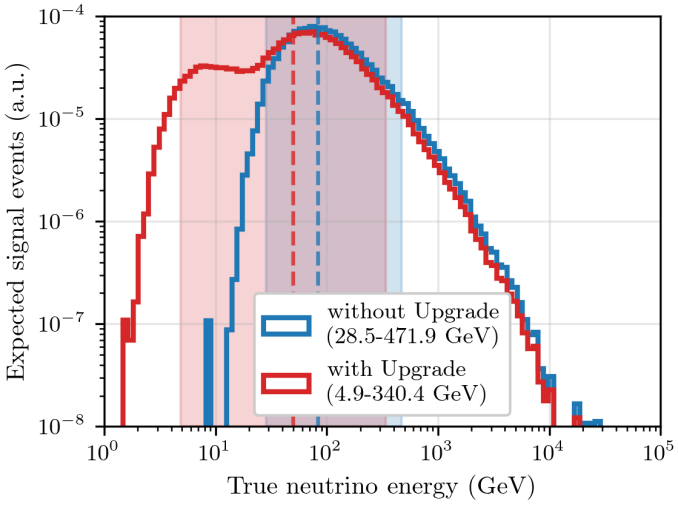


Figure 13.3: Distribution of expected neutrino signal events as a function of true neutrino energy for NGC 1068 ($\gamma = 3.4$). The shaded regions indicate the central 90 % sensitive energy range for each configuration.

out entirely. Physically, this weighting procedure folds the intrinsic geometric acceptance and interaction probability of the detector (captured by $w_{\text{one},i}$ from Equation (12.1)) with the assumed physical flux distribution of the astrophysical source, being a power law with index γ. Consequently, when restricting to events that pass the final selection criteria (see Section 12.3.3), the resulting CDF represents the true energy spectrum of the signal events that the analysis is actually expected to observe.

Figure 13.3 compares the signal distributions, where the most prominent feature is the reduction of the lower energy threshold. For the *without Upgrade* scenario, the sensitive range begins at 28.5 GeV. Below this energy, the effective area combined with the selection cuts is insufficient to retain significant signal. The addition of IceCube Upgrade results in a distinct peak centered near 8 GeV, lowering the sensitive range threshold to 4.9 GeV. This shift confirms that the sensitivity gains reported above are driven by IceCube Upgrade's ability to resolve neutrinos in the 5 GeV to 30 GeV window. While the values quoted here are specific to the applied event selection and combined dataset livetime, the physical conclusion that a regime largely inaccessible to the legacy detector becomes available holds true.

Note on the high-energy Cutoff: The upper bound of the sensitive range decreases from 472 GeV (without Upgrade) to 340 GeV (with Upgrade). This is primarily due to the influx of low-energy events shifting the percentile, but also stems from the 500 GeV simulation cutoff described in Section 13.1. The GRECO dataset is not extrapolated for $E_{\text{true}} > 500\,\text{GeV}$ in this specific figure to facilitate the comparison. Had the additional GRECO MC extrapolation been applied, both right tails would align closer (as similarly seen in Figure 13.1). While this represents a conservative limitation, its impact for soft sources like NGC 1068 is minor, as the signal naturally tapers off at these energies.

13.2.3 Physics Reach: spectral Constraints

Figure 13.4 synthesizes the projected physics reach of IceCube Upgrade in the context of the evolving experimental picture of NGC 1068. The figure compares the five-year sensitivity projections within the sensitive range from Figure 13.3 against the best-fit fluxes reported in the initial 2022 discovery ($\gamma = 3.2$) [6] and the updated 2025 follow-up ($\gamma = 3.4$) [226]. The two distributions shown represent different quantities: the measured best-fit fluxes represent the static, intrinsic brightness of the source, whereas the sensitivity limits represent the detector's statistical capability, which improves (lowers) as data accumulates over time. The fact that IceCube Upgrade sensitivity drops below the most recently measured flux earlier than the legacy configuration demonstrates the statistical margin that will become available to perform precision spectral measurements. The comparison highlights two advantages provided by IceCube Upgrade:

1. Extending the Energy Horizon: The *with Upgrade* scenario maintains a limit down to 4.9 GeV, extending deeper into the low-energy flux space than the baseline scenario with 28.5 GeV. Effectively, IceCube Upgrade opens an observational window that is currently inaccessible. This extended energy lever arm is crucial for constraining how far the observed power law spectrum can be extrapolated to lower energies. Since a continuous, unbroken power law is expected to turn over or break at some minimum energy to avoid divergent energetics, IceCube Upgrade's lowered threshold enables the testing of broken power law models and the identification of low-energy spectral cutoffs.

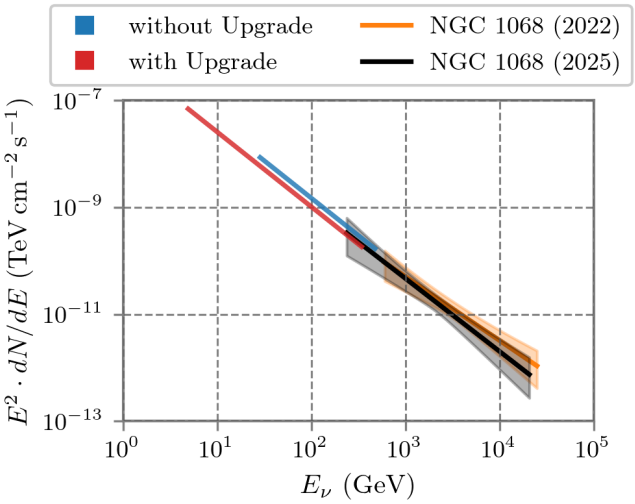


Figure 13.4: Comparison of projected IceCube Upgrade sensitivities (red) and existing detector sensitivities (blue) with measured neutrino fluxes from NGC 1068. The sensitivities represent the projected integrated performance in 2031 from Figure 13.2, plotted over the sensitive energy range from Figure 13.3 and weighted to an $E^{-3.4}$ spectrum. The shaded bands indicate the 1σ uncertainties of the published best-fit fluxes from 2022 (orange, $\gamma = 3.2$) [6] and 2025 (black, $\gamma = 3.4$) [226].

2. Relevance to Spectral Softening: The importance of this low-energy reach is amplified by the shift in the measured properties of NGC 1068. As shown in Figure 13.4, the 2025 result favors a softer spectral index compared with 2022. A softer spectrum implies that a larger fraction of the neutrino flux is emitted at lower energies. Therefore, IceCube Upgrade's increased acceptance in the GeV range is well-suited to constrain the bulk of the signal flux favored by the most recent best-fit.

Within five years of operation, IceCube Upgrade is expected to provide sufficient statistical power not only to detect the source but also to begin constraining its spectral shape at GeV energies, complementing the current decade-long measurements at higher energies. These results are likely conservative, as the analysis does not yet model the expected reduction in systematic uncertainties resulting from IceCube Upgrade's calibration campaign, which would further enhance sensitivity beyond the estimates presented here. Artificially improving reconstruction accuracy results in a relative gain of 1 % in sensitivity for every 2 % improvement in angular uncertainty, as shown in [255] for the Galactic Plane.

13.3 Applicability to transient Searches

While the primary focus of this chapter is the 19-year steady-state emission, IceCube Upgrade offers even higher performance enhancements for time-dependent analyses. Transient astrophysical phenomena (see Section 2.2.3) constitute a particularly promising target class, as the short duration of the signal window naturally suppresses the accumulation of atmospheric muon background, allowing for the detection of lower-energy events that would be indistinguishable from background in a steady-state search. Candidates for GeV-scale neutrino bursts include gamma-ray bursts and supernovae, both of which have

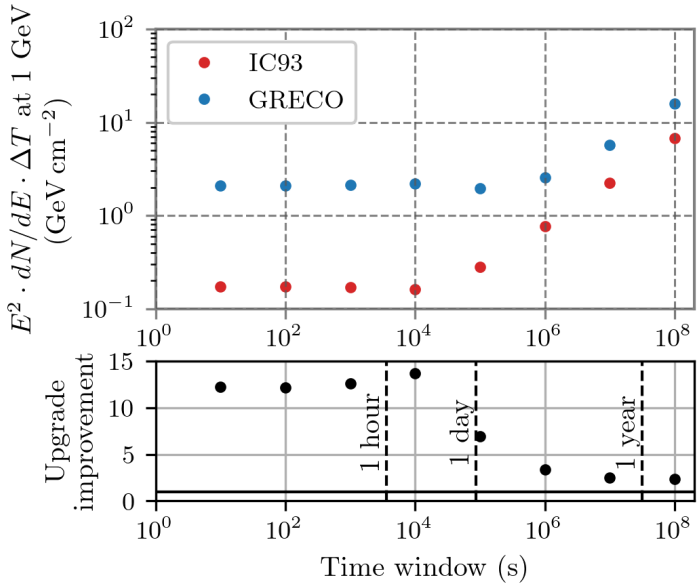


Figure 13.5: Projected sensitivity and relative IceCube Upgrade improvement as a function of emission time window. The target is a box-shaped flare from NGC 1068 with a spectral index of $\gamma = 3.4$. The y-axis represents the flux sensitivity ($E^2 dN/dE \cdot \Delta T$) evaluated at 1 GeV. The bottom panel shows the improvement factor, defined as the ratio of the GRECO sensitivity to the IC93 sensitivity. Vertical dashed lines mark reference durations.

been proposed as potential sources in this energy regime [52, 258]. To evaluate IceCube Upgrade's potential for such searches, this section presents the projected sensitivity to short-timescale events.

Methodological Note: Unlike previous steady-state projections that combined future IceCube Upgrade data with historical data, this evaluation directly compares the IC93 dataset against the GRECO dataset. This isolates the specific performance gain provided by the new instrumentation.

To quantify the performance gain, the steady-state event selection is applied to box-shaped flares from the direction of NGC 1068 ($\gamma = 3.4$) with durations ranging from 10 s to 10^8 s (≈ 3.2 years), evaluated using the extended likelihood formulation from Section 11.1. The results are summarized in Figure 13.5. The analysis reveals two distinct regimes governed by background accumulation:

1. **Signal-Limited Regime ($t \lesssim 10^4$ s):** For durations shorter than a few hours, the expected atmospheric background is negligible. In this regime, sensitivity is limited only by the Poisson detection threshold ($n_s \approx 2.3$) and the detector's effective area. As shown in the bottom panel of Figure 13.5, IceCube Upgrade provides an improvement factor of over an order of magnitude in this regime. Because the background is practically zero, the analysis can fully exploit IceCube Upgrade's larger effective area at low energies (1 GeV to 50 GeV) without needing to apply strict cuts to reject muons. IceCube Upgrade is therefore particularly well suited for short transient searches.
2. **Background-Limited Regime ($t \gtrsim 10^5$ s):** As the time window extends beyond one day, atmospheric background events begin to accumulate. To maintain a constant (low) false-alarm rate, the analysis requires a higher signal flux, causing the

sensitivity values to degrade (rise). As the background increases, the relative advantage of IceCube Upgrade diminishes, dropping from a factor of ~ 12 to a factor of ~ 2.3. In this high-statistics background regime, an analysis must either leverage higher-energy events or employ quality cuts to distinguish signal from atmospheric background, where the relative effective area advantage of IceCube Upgrade is less pronounced than at the energy threshold.

This contextualizes the results of previous sections. The factor ~ 1.58 improvement reported for the 19-year steady-state analysis approaches the asymptotic limit on the far right of Figure 13.5. In contrast, for short-term transient phenomena where the low-energy capabilities are fully utilized, the performance gains are significantly higher, further demonstrated in [255].

13.4 Additional Validations

To complete the characterization of this analysis, this section presents supplementary validations for parameter estimation, specifically examining the parameter fit bias and demonstrating the impact of IceCube Upgrade on signal localization.

13.4.1 Fit Bias

To confirm the reliability of an unbinned likelihood analysis as described in Section 11.1, a fit bias test is performed. This test quantifies the accuracy with which the likelihood maximizer can reconstruct the physical parameters of a source, specifically the number of signal events n_s and the spectral index γ, under signal-injection conditions. The verification involves injecting a known number of signal events n_{inj} from a point source at the location of NGC 1068 with a spectral index of $\gamma = 3.4$ and running the likelihood maximization on these pseudo-experiments.

As shown in Figure 13.6, both scenarios exhibit a systematic reconstruction bias: the analysis tends to overestimate the number of signal events and reconstruct a harder spectral index than injected. This is known behavior of unbinned likelihood analyses over very long livetimes (19 years), where the maximizer can drive the fit towards *larger and harder* signal hypotheses to explain minor background fluctuations.

To address such biases, particularly those arising from non-Gaussian long tails of the PSF inherent to low-energy reconstruction, a related analysis [256] investigated the use of the King function[3] as an alternative spatial probability density function. While this modified PSF modeling was found to reduce the reconstruction bias, its impact on the final sensitivity and discovery potential is negligible. This result is important for the interpretation of this work: it demonstrates that while the standard likelihood estimator may be biased in its parameter recovery, the statistical power to distinguish signal from background (the sensitivity) remains robust and is not artificially inflated by these reconstruction systematics.

[3]IceCube analyses typically model the PSF using a Kent distribution, as explained in Section 11.1. However, low-energy reconstructions often exhibit non-Gaussian tails (a surplus of events with large angular errors) which are exponentially suppressed by the Kent model. The King function accommodates these tails via a power law fall-off, providing a better statistical description of outliers without compromising the core resolution.

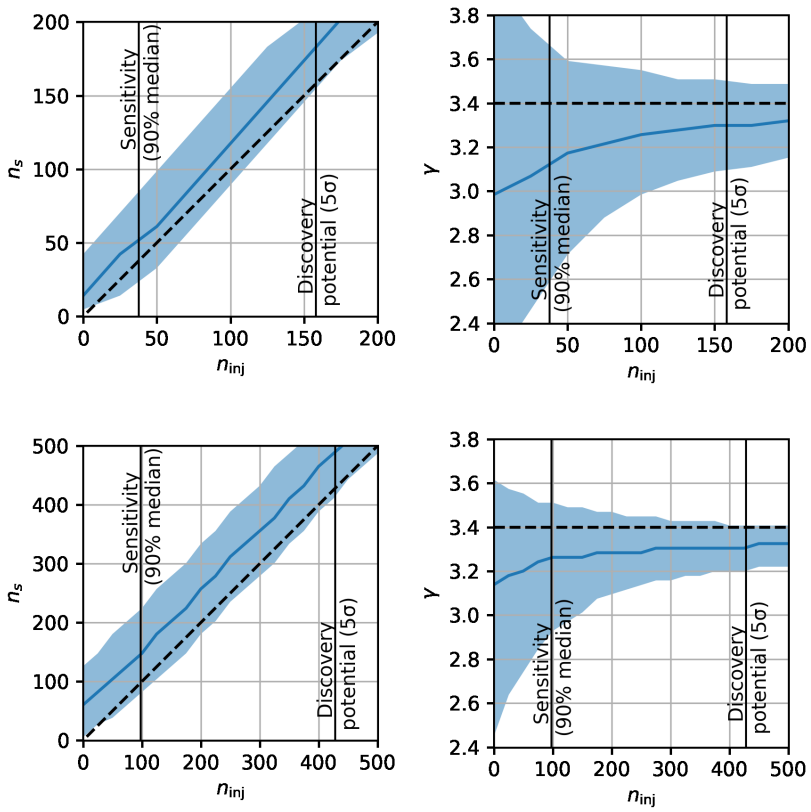


Figure 13.6: Reconstruction bias for the number of signal events (n_s, left) and spectral index (γ, right) for the baseline (top) and IceCube Upgrade (bottom) configurations from the direction of NGC 1068. The blue solid line indicates the median reconstructed value, while the shaded bands represent the 68 % central quantile range. The dashed black lines indicate the true injected values. The analysis systematically overestimates the signal strength (slope > 1), a known effect in long-livetime integrated searches.

Crucially, the presence of this bias does **not** invalidate the sensitivity results presented in this chapter. Sensitivity is established using a frequentist construction based on TS distributions. The sensitivity threshold is defined relative to the background-only TS distribution, which is subject to the same reconstruction characteristics as the signal trials. Because the sensitivity is determined by comparing signal-injected and background-only TS distributions, the absolute value of the fitted parameter $\hat{n}_s$ does not directly determine the sensitivity limit. Furthermore, this study is limited to sensitivity projections (a quantification of detector capability) and does not involve *unblinding*[4] the data. A flux measurement would require correcting this bias to recover the true number of neutrinos. However, for determining the threshold at which a signal becomes statistically distinguishable from background (the sensitivity), the uncorrected estimator remains a valid and reliable metric.

[4]In IceCube, unblinding refers to the process of examining actual detector data after the analysis procedure has been finalized. To avoid experimenter bias, analyses are developed on "blinded" data (scrambled in right ascension). Once the cuts and methods are established, the real data is revealed to search for astrophysical sources. Since this thesis presents a sensitivity study for a future detector, no unblinding is performed and therefore no physical measurement of a source flux is reported.

13.4.2 Demonstration of spatial Resolution

While the sensitivity projections presented in previous sections quantify the statistical power of the analysis, they do not directly illustrate an important driver of this improvement: the spatial localization of neutrino events. To visualize the impact of IceCube Upgrade's improved angular resolution at analysis level, a spatial demonstration using a fixed signal injection is performed. To isolate reconstruction improvements from statistical fluctuations in signal count, 275 000 pseudo-trials are generated for each scenario with a fixed injection of $n_{\text{inj}} = 79$ events (the physical excess observed for NGC 1068) at the source coordinates. For each trial, the spatial spread of the reconstructed events is quantified by calculating an approximate area of a variance ellipse:

$$\mathcal{A} \approx \pi \cdot \frac{\sigma_{\text{dec}}}{2} \cdot \frac{\sigma_{\text{ra}}}{2}, \tag{13.1}$$

where σ_{dec} and σ_{ra} denote the standard deviation of the reconstructed event coordinates in declination and right ascension, respectively. By considering the spread of the events rather than the angular error of a single event, this metric serves as a proxy for the signal's angular extent on the sky. For this specific comparison, the cut on the angular uncertainty from Section 12.3.3 is set to 7° for both datasets to compare the pointing performance under identical selection criteria rather than to maximize sensitivity. The results, visualized in Figure 13.7, reveal two key insights:

The histogram in the left panel shows the distribution of variance ellipse area as per Equation (13.1) across all trials. A systematic reduction is observed: the distribution of the *with Upgrade* scenario is shifted to lower values relative to the *without Upgrade* distribution, with a mean reduction in signal area of approximately 7 deg^2. This demonstrates that the improvement arises from detector performance rather than statistical fluctuations. For the scenario including IceCube Upgrade, 42.4 % of the 79 generated signal

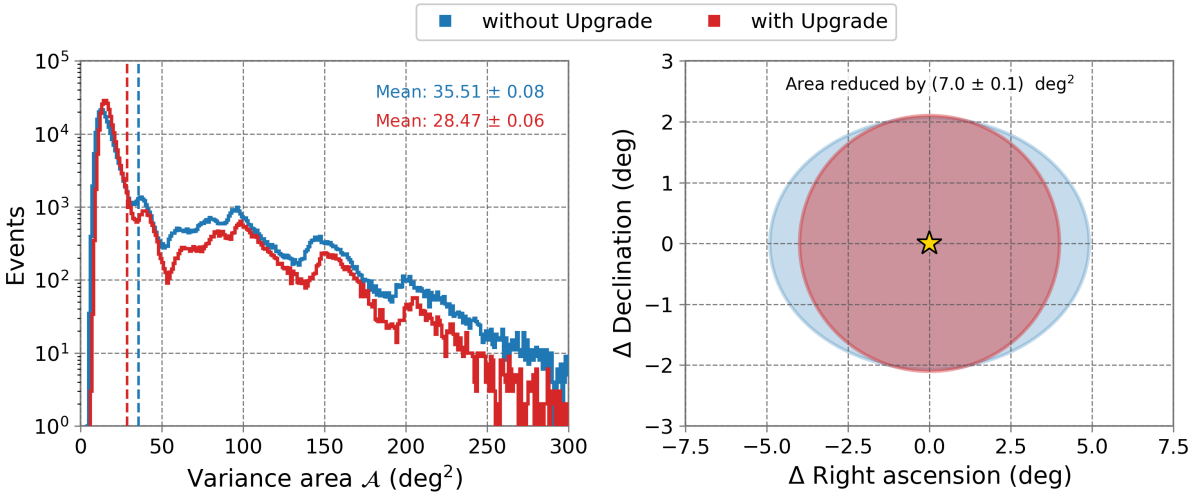


Figure 13.7: Quantification of the spatial resolution improvement using fixed-signal pseudo-experiments ($n_{\text{inj}} = 79$) from 275 000 trials. **(left)** Distribution of the signal spread area quantified by variance ellipse area. **(right)** Average 1σ variance ellipses, computed from the standard deviations of the reconstructed coordinates (σ_{RA}, σ_{dec}) for each pseudo-experiment. The axes represent the offset centered at the true source position (0,0), defined as Δ = reconstructed − true.

events stem from the IC93 dataset, despite the dataset only accounting for 26.9 % of the total livetime. To visualize this effect, the mean variance ellipses are plotted relative to the true source position in the right panel. Geometrically, the improvement in resolution occurs predominantly in right ascension. This reflects the specific geometry of the overall detector: while the legacy detector possesses sufficient vertical segmentation for resolving declination, the additional denser inter-string spacing from IceCube Upgrade provides more accurate transverse sampling to improve the resolution of right ascension.

It is worth emphasizing that while the background estimation relies on right ascension scrambling (creating a uniform background distribution in RA), the signal sensitivity is driven by the contrast between the signal cluster and this background. By reducing the signal variance ellipse (driven primarily by the improved RA resolution) IceCube Upgrade confines the event cluster to a smaller solid angle. This reduces the integrated background contamination within the signal region, thereby increasing the statistical significance of the source.

13.5 Deployment Constraints and the IC91 Approximation

While the analytical framework developed for this thesis was built upon the IC93 simulation, the physical deployment of IceCube Upgrade encountered hardware constraints during the 2025–2026 polar season. These real-world constraints necessitated a re-evaluation of the detector's realized state during the final phase of this thesis. With respect to the initially planned seven-string layout, string 87 was not deployed, string 92 suffered a total communication failure post-installation and communications have not been established with approximately 10 % of the newly installed optical modules (specifically mDOMs and D-Eggs) at the time of writing. To assess the impact of this reduced geometry, the analysis was reevaluated using the IC91 configuration, which serves as a heuristic approximation that masks the hits associated with non-functional components within the existing IC93 simulation. It is emphasized that IC91 is an imperfect approximation. Because the machine learning architectures, such as the GraphNeT reconstruction and background rejection models detailed in Section 12.1.2, were optimized for the full IC93 footprint, their

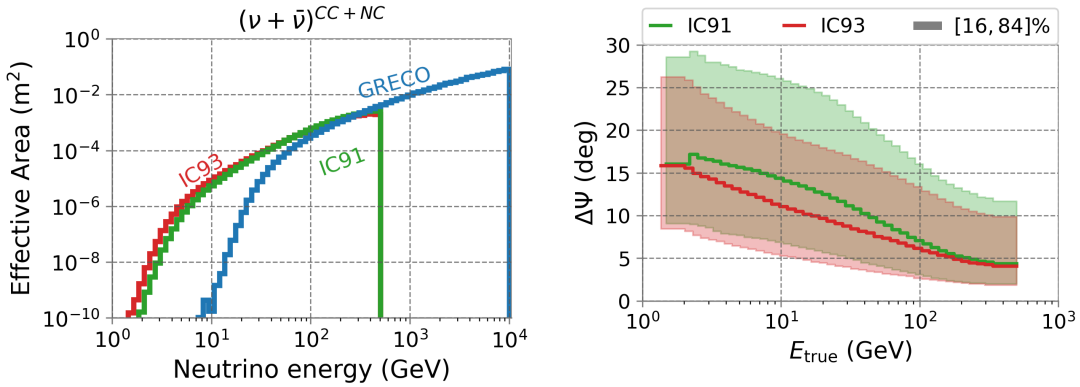


Figure 13.8: Comparison between the IC93 (red) and IC91 (green) configurations after applying the cuts as described in Section 12.3.3. **(left)** Effective area, including GRECO (blue). **(right)** Angular resolution.

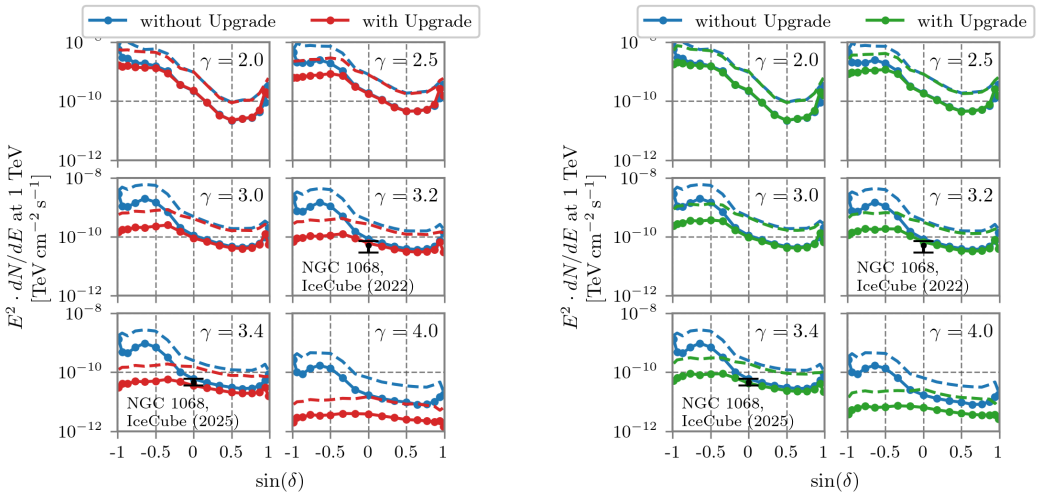


Figure 13.9: Sensitivity limits as a function of source declination for different spectral indices. **(left)** IC93 configuration from Figure 13.1b. **(right)** IC91 configuration.

application to the IC91 geometry is inherently sub-optimal. Furthermore, this selection does not include dedicated systematic re-simulations or updated cuts. Thus, the results presented here likely represent a conservative lower bound on the true performance of the realized array.

As illustrated in Figure 13.8 (left), the reduction in effective area is non-uniform across the energy range. At the lowest energies ($< 5\,\text{GeV}$), the ratio of IC93 to IC91 effective area is between 2.4 and 10.3 for the specific cuts applied in Section 12.3.3. The effective areas become comparable around $\sim 60\,\text{GeV}$, above which IC91 even exhibits slightly larger values toward the 500 GeV threshold. This change in effective area, combined with the worsened angular resolution shown in Figure 13.8 (right), impacts the point-source sensitivity. The application of the IC91 configuration demonstrates that the underlying event selection framework remains statistically robust. Systematic verification of the background TS distributions across the sky, together with the optimization of angular uncertainty cut values, shows that the analysis reproduces the baseline performance, exhibiting the same trends as in Figure 12.10 and Figure 12.11. However, while these heuristic checks justify the current projections, a comprehensive re-evaluation will be necessary once a dedicated simulation set (modeled specifically on the final deployed hardware footprint) is available to be validated against the first acquisition of experimental detector data.

The gain threshold at which IceCube Upgrade begins to improve upon the legacy detector has thus shifted (see Figure 13.9). While the idealized IC93 case showed improvements for any spectrum softer than $\gamma = 3.2$ in the Northern sky, the realized IC91 array requires even softer spectra ($\gamma \geq 3.4$) to maintain that dominance. The Southern sky advantage remains qualitatively consistent with the trends discussed in Section 13.1, though at degraded absolute sensitivities (a reduction of 23 % in the Northern and 66 % in the Southern sky at $\gamma = 3.4$ relative to the idealized IC93 array).

The temporal impact on the primary target, NGC 1068, is quantified in Figure 13.10 (left). While IceCube Upgrade still accelerates the path to a higher-significance measurement of NGC 1068 ($\gamma = 3.4$), the slope of the sensitivity evolution is shallower in the IC91 configuration. The projected sensitivity improvement after five years of integrated

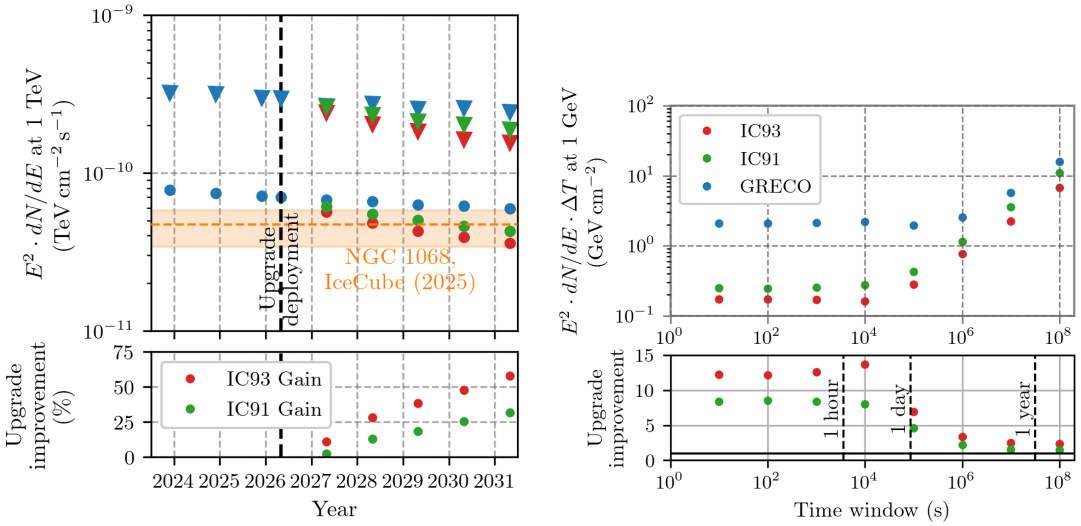


Figure 13.10: Comparison of IC91 and IC93 expectations. **(left)** Temporal sensitivity evolution from Figure 13.2. **(right)** Sensitivity as a function of time window from Figure 13.5.

livetime drops from 57.83 % to 31.73 %. For transient searches (Figure 13.10, right), the annualized improvement factor remains significant, albeit at a reduced level. For short time windows ($\Delta T < 1000\,\mathrm{s}$), the improvement decreases from twelve-fold to roughly eight-fold. For year-scale observations, the improvement stabilizes at approximately 1.4, confirming a lower asymptotic limit compared with what is discussed in Section 13.3.

Although deployment constraints have reduced the anticipated physics reach, these results demonstrate that the realized five-string IceCube Upgrade array nevertheless enhances the detector's sensitivity at low energies. Furthermore, the successful deployment and operation of novel calibration devices (see Section 3.6) will improve the characterization of the optical properties of both the bulk ice and the refrozen hole ice, reducing the dominant systematic uncertainties of the observatory. These improvements are expected to enhance event reconstruction for both future and archival data across the entire optical array, extending the scientific impact of IceCube Upgrade beyond the low-energy sensitivity gains presented here.

13.6 Summary

This chapter demonstrates the expected impact of IceCube Upgrade on neutrino point-source astronomy. By lowering the energy threshold from $\sim 30\,\mathrm{GeV}$ to $\sim 5\,\mathrm{GeV}$ in this specific event selection, IceCube Upgrade opens a new observational window into the soft-spectrum universe, which was previously inaccessible to the legacy detector (Section 13.2.2). The sensitivity to NGC 1068 is projected to improve by $\sim 58\,\%$ over five years, enabling the detector to probe the flux levels indicated by recent measurements with statistical significance. Crucially, the measured $\gamma = 3.4$ flux from [226] implies a high expected flux extending down to sub-TeV energies. This provides motivation for GeV observations, as measuring the spectral shape in this regime is required to break the degeneracies between competing astrophysical production models and constrain low-energy cutoffs.

Even when accounting for the hardware constraints modeled by the IC91 configuration (Section 13.5), an improvement of 31.73 % is retained, accelerating the path toward precision spectral characterization (Section 13.2.3). More significant gains are expected for sources located in the Southern sky due to improved background rejection. Furthermore, the largest improvements are found for short-term transient phenomena, where an eight-fold improvement in sensitivity allows for the probing of faint, GeV-scale neutrino bursts (Section 13.3). Although the realized five-string array tempers the idealized IC93 baseline projections, IceCube Upgrade remains a necessary enhancement that ensures the detector's continued relevance in multi-messenger astronomy.

Ultimately, these results establish that IceCube Upgrade does not merely extend the current detector's lifespan but significantly broadens its scientific capabilities, shifting the detector from the discovery of sparse high-energy events toward precision spectral characterization in the GeV regime.

14 Summary and Outlook

The planned IceCube-Gen2 expansion will considerably increase the instrumented volume, requiring the mass production and deployment of thousands of next-generation optical modules. This thesis presented comprehensive hardware developments, characterization efforts and simulation benchmarks crucial for finalizing the IceCube-Gen2 optical module. These developments are intrinsically linked to the successful precursor modules, the LOM16 and LOM18, which have recently been deployed in IceCube Upgrade.

Extensive PMT characterization efforts directly influenced the development across multiple design iterations (Chapter 6), demonstrating that both evaluated manufacturer models meet the IceCube-Gen2 specifications. While high quantum efficiency variants remain a viable option for increasing optical performance, their integration involves a cost-benefit trade-off compared with deploying a larger number of standard modules, particularly due to concerns regarding increased background rates.
To meet the demand for producing nearly 10 000 optical modules for the IceCube-Gen2 optical array, scalable mechanical assembly procedures were pioneered (Chapter 7). By transitioning to fully pre-cast silicone gel pads coupled using automated pushing mechanisms, such as the presented pneumatically driven scissor lifts, the time-consuming liquid backfilling inside the pressure vessel during LOM production could be eliminated. Building on the lessons learned from the LOM16 and LOM18, the groundwork for the Gen2-OM is now established, with ongoing efforts focused on finalizing the module layout utilizing a newly optimized pressure vessel shape.

To evaluate the performance of these hardware designs, a Geant4 simulation framework was extended to encompass all IceCube optical modules (Chapter 8). By integrating detailed CAD geometries to quantify mechanical shadowing and implementing empirical spatial PMT maps, the simulation now provides a realistic description of the directional acceptance of all optical modules. Evaluating the directional acceptance against the local photon incidence angle revealed complex internal optics (such as structural off-axis peaks and photon traversal within the module) unique to the multi-PMT geometries. When weighted against the physical propagation spectrum of Cherenkov light in glacial ice, the LOM16 baseline achieves a factor of 3.8 to 4.3 improvement in effective area relative to the DOM. These results strongly validate the LOM architecture for the Gen2-OM, further demonstrating that it outperforms the two primary optical module designs currently deployed in IceCube Upgrade.

Additionally, studies of the module background (Chapter 9) establish that the high scintillation noise observed in Vitrovex borosilicate glasses originates from the silica sand used in the glass production process, as well as from the potassium nitrate concentration. By remeasuring the scintillation spectrum to verify the temperature-independent parameterization, the background simulation was expanded to encompass all multi-PMT optical modules. The results are in excellent agreement with available measurements of integrated PMT background rates across all optical modules. As a result, the OMSim framework is now positioned better than ever to produce simulation data to be integrated into IceTray, enabling comprehensive background modeling and the generation of directional sensitivity tables.

The final part of this thesis establishes the first point-source sensitivity benchmark for IceCube Upgrade, providing a quantitative roadmap for the detector's early physics phase. By developing an unbinned maximum likelihood framework optimized for pointing accuracy across the sky, this work indicates that IceCube Upgrade will extend IceCube's capabilities from a high-energy discovery instrument toward precision spectral measurements at GeV energies. The analysis shows that IceCube Upgrade provides its largest gains in regimes previously dominated by atmospheric backgrounds. In the Southern sky, the denser instrumentation and multi-PMT optical modules enable improved atmospheric muon vetoing, yielding substantial sensitivity improvements for soft-spectrum sources. This enhanced background rejection extends into the time domain, as for short-term transient phenomena ($\Delta T < 1000\,\mathrm{s}$), IceCube Upgrade provides up to an order of magnitude improvement in sensitivity, allowing the detector to probe faint, GeV-scale neutrino bursts from sources such as GRBs and supernovae that were previously obscured. For the neutrino reference source NGC 1068, the inclusion of IceCube Upgrade data is projected to improve sensitivity by 57.83 % over five years, accelerating the prospects for characterizing the hidden core's acceleration mechanisms. Crucially, this work also addressed the deployment constraints, providing the first heuristic projections for the realized five-string array. Even under these constrained hardware conditions, the analysis identified a five-year integrated sensitivity gain of 31.73 % for NGC 1068, affirming IceCube Upgrade's scientific potential. These projections are likely conservative, as they do not account for the improved understanding of the Antarctic ice expected from IceCube Upgrade's dedicated calibration instrumentation.

Ultimately, these results bridge the gap between idealized simulation and experimental reality, providing the statistical foundation necessary to move toward precision neutrino astronomy once the first unblinded data becomes available. In particular, future analyses incorporating IceCube Upgrade should move beyond single power law descriptions and test physically motivated spectral models, including broken power laws, exponential cutoffs and hadronic interaction models predicting low-energy spectral features. Such models may exploit IceCube Upgrade's enhanced GeV sensitivity differently than the standard power law model used in this thesis.

Overall, this work spans the complete chain from detector development to physics performance. The prototyping methods and PMT characterization constitute the hardware basis for the next generation of IceCube optical modules. The optical performance and background studies provide essential inputs to IceCube's simulation framework, with the developed tools being readily applicable once the Gen2-OM geometry is finalized. Building on this foundation, this thesis presents the first point-source sensitivity study for IceCube Upgrade incorporating its multi-PMT optical modules, showing how their improved detector performance translates directly into enhanced astrophysical sensitivity. As such, this work contributes to both realizing next-generation IceCube instrumentation and establishing a quantitative benchmark for the scientific capabilities enabled by this detector technology.

Bibliography

[1] B. T. Cleveland, et al., *Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector*, ApJ 496 (1) (1998) 505. URL https://doi.org/10.1086/305343

[2] BAIKAL Collaboration, *The Baikal underwater neutrino telescope: Design, performance and first results*, Astroparticle Physics 7 (3) (1997) 263–282. URL https://www.sciencedirect.com/science/article/pii/S0927650597000224

[3] IceCube Collaboration, *The IceCube Neutrino Observatory: instrumentation and online systems*, JINST 12 (03) (2017) P03012. URL https://doi.org/10.1088/1748-0221/12/03/P03012

[4] IceCube Collaboration, *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*, Science 342 (6161) (2013) 1242856. URL https://www.science.org/doi/abs/10.1126/science.1242856

[5] IceCube Collaboration, *Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert*, Science 361 (2018) 147–151. URL https://science.sciencemag.org/content/361/6398/147

[6] IceCube Collaboration, *Evidence for neutrino emission from the nearby active galaxy NGC 1068*, Science 378 (6619) (2022) 538–543. URL http://dx.doi.org/10.1126/science.abg3395

[7] IceCube Collaboration, *Observation of high-energy neutrinos from the Galactic plane*, Science 380 (6652) (2023) 1338–1343. URL http://dx.doi.org/10.1126/science.adc9818

[8] A. Ishihara, *The IceCube Upgrade - Design and Science Goals*, PoS ICRC2019 (2019) 1031. URL https://pos.sissa.it/358/1031/pdf

[9] IceCube-Gen2 Collaboration, *IceCube-Gen2: A Vision for the Future of Neutrino Astronomy in Antarctica* (2014). URL https://arxiv.org/abs/1412.5106

[10] IceCube-Gen2 Collaboration, *IceCube-Gen2 Technical Design Report*, https://icecube-gen2.wisc.edu/science/publications/tdr/, accessed: 2026-06-20 (2024).

[11] T. K. Gaisser, R. Engel, E. Resconi, Cosmic Rays and Particle Physics, 2nd Edition, Cambridge University Press, 2016, ISBN: 9781139192194. URL https://doi.org/10.1017/CBO9781139192194

[12] P. Chen, Y.-K. Liu, *Plasma wakefield: from accelerators to black holes*, Rev. Mod. Plasma Phys. 10 (2026) 7. URL http://dx.doi.org/10.1007/s41614-026-00215-z

[13] J. Blümer, R. Engel, J. R. Hörandel, *Cosmic rays from the knee to the highest energies*, Progress in Particle and Nuclear Physics 63 (2) (2009) 293–338. URL http://dx.doi.org/10.1016/j.ppnp.2009.05.002

[14] K. Greisen, *End to the Cosmic-Ray Spectrum?*, Phys. Rev. Lett. 16 (1966) 748–750. URL https://link.aps.org/doi/10.1103/PhysRevLett.16.748

[15] G. T. Zatsepin, V. A. Kuz'min, *Upper Limit of the Spectrum of Cosmic Rays*, Soviet Journal of Experimental and Theoretical Physics Letters 4 (1966) 78.

[16] The Pierre Auger Collaboration, *Features of the Energy Spectrum of Cosmic Rays above* 2.5×10^{18} eV *Using the Pierre Auger Observatory*, Phys. Rev. Lett. 125 (2020) 121106. URL https://link.aps.org/doi/10.1103/PhysRevLett.125.121106

[17] Pierre Auger Collaboration, *Astroparticle and particle physics at ultra-high energy: Results from the Pierre Auger Observatory*, SciPost Phys. Proc. (2023) 034. URL https://scipost.org/10.21468/SciPostPhysProc.13.034

[18] F. Halzen, A. Kheirandish, *IceCube and High-Energy Cosmic Neutrinos* (2022). URL https://arxiv.org/abs/2202.00694

[19] R. Aloisio, V. Berezinsky, P. Blasi, *Ultra high energy cosmic rays: implications of Auger data for source spectra and chemical composition*, JCAPBP 2014 (10) (2014) 020. URL https://doi.org/10.1088/1475-7516/2014/10/020

[20] N. Globus, D. Allard, E. Parizot, *A complete model of the cosmic ray spectrum and composition across the Galactic to extragalactic transition*, Phys. Rev. D 92 (2015) 021302(R). URL https://link.aps.org/doi/10.1103/PhysRevD.92.021302

[21] E. Fermi, *On the origin of the cosmic radiation*, Phys. Rev. 75 (1949) 1169–1174. URL https://link.aps.org/doi/10.1103/PhysRev.75.1169

[22] A. M. Hillas, *The Origin of Ultra-High-Energy Cosmic Rays*, ARA&A 22 (1) (1984) 425–444. URL http://dx.doi.org/10.1146/annurev.aa.22.090184.002233

[23] IceCube Collaboration, *Astrophysical neutrinos and cosmic rays observed by IceCube*, Advances in Space Research 62 (10) (2018) 2902–2930. URL http://dx.doi.org/10.1016/j.asr.2017.05.030

[24] IceCube Collaboration, *Lowering IceCube's Energy Threshold for Point Source Searches in the Southern Sky*, ApJL 824 (2) (2016) L28. URL https://doi.org/10.3847/2041-8205/824/2/L28

[25] M. Ahlers, *Neutrino Sources from a Multi-Messenger Perspective*, EPJ Web of Conferences 209 (2019) 01013. URL http://dx.doi.org/10.1051/epjconf/201920901013

[26] Fermi/LAT Collaboration, *The Large Area Telescope on the Fermi Gamma-ray Space Telescope mission*, ApJ 697 (2) (2009) 1071–1102. URL http://dx.doi.org/10.1088/0004-637X/697/2/1071

[27] N. Tomassetti, *AMS-02 in space: physics results, overview, and challenges*, Nuclear and Particle Physics Proceedings 265–266 (2015) 245–247. URL http://dx.doi.org/10.1016/j.nuclphysbps.2015.06.063

[28] HAWC Collaboration, *HAWC: A next generation all-sky gamma-ray telescope*, Journal of Physics: Conference Series 60 (1) (2007) 131. URL https://doi.org/10.1088/1742-6596/60/1/024

[29] Pierre Auger Collaboration, *The Pierre Auger Cosmic Ray Observatory*, NIM-A 798 (2015) 172–213. URL http://dx.doi.org/10.1016/j.nima.2015.06.058

[30] T. K. Gaisser, *Atmospheric neutrinos* (2019). URL https://arxiv.org/abs/1910.08851

[31] P. Heix, S. Tilav, C. Wiebusch, et al., *Seasonal Variation of Atmospheric Neutrinos in IceCube*, PoS ICRC2019 (2019) 465. URL https://pos.sissa.it/358/0465

[32] M. Honda, et al., *Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data*, Phys. Rev. D 75 (4). URL http://dx.doi.org/10.1103/PhysRevD.75.043006

[33] R. Enberg, M. H. Reno, I. Sarcevic, *Prompt neutrino fluxes from atmospheric charm*, Phys. Rev. D 78 (4). URL http://dx.doi.org/10.1103/PhysRevD.78.043005

[34] A. Franceschini, G. Rodighiero, *The extragalactic background light revisited and the cosmic photon-photon opacity*, Astronomy and Astrophysics 603 (2017) A34. URL http://dx.doi.org/10.1051/0004-6361/201629684

[35] KATRIN Collaboration, *Direct neutrino-mass measurement based on 259 days of KATRIN data*, Science 388 (6743) (2025) 180–185. URL http://dx.doi.org/10.1126/science.adq9592

[36] C. Spiering, *Towards high-energy neutrino astronomy*, EPJ H 37 (3) (2012) 515–565. URL https://doi.org/10.1140/epjh/e2012-30014-2

[37] IceCube Collaboration, *Differential limit on the extremely-high-energy cosmic neutrino flux in the presence of astrophysical background from nine years of IceCube data*, Phys. Rev. D 98 (6). URL http://dx.doi.org/10.1103/PhysRevD.98.062003

[38] S. Hümmer, M. Maltoni, W. Winter, C. Yaguna, *Energy dependent neutrino flavor ratios from cosmic accelerators on the Hillas plot*, Astroparticle Physics 34 (4) (2010) 205–224. URL http://dx.doi.org/10.1016/j.astropartphys.2010.07.003

[39] P. Schneider, Extragalactic Astronomy and Cosmology: An Introduction, 2nd Edition, Springer, 2015, ISBN: 978-3-642-54083-7.

[40] R. D. Blandford, D. G. Payne, *Hydromagnetic flows from accretion disks and the production of radio jets*, MNRAS 199 (1982) 883–903. URL https://doi.org/10.1093/mnras/199.4.883

[41] R. Blandford, D. Meier, A. Readhead, *Relativistic Jets from Active Galactic Nuclei*, ARA&A 57 (2019) 467–509. URL https://www.annualreviews.org/content/journals/10.1146/annurev-astro-081817-051948

[42] C. M. Urry, P. Padovani, *Unified Schemes for Radio-Loud Active Galactic Nuclei*, PASP 107 (1995) 803. URL http://dx.doi.org/10.1086/133630

[43] ASTRO-H Science Working Group, *ASTRO-H White Paper - AGN Reflection* (2014). URL https://arxiv.org/abs/1412.1177

[44] A. Khatee Zathul, M. Moulai, K. Fang, F. Halzen, *An NGC 1068-informed Understanding of Neutrino Emission of the Active Galactic Nucleus TXS 0506+056*, ApJ 984 (1) (2025) 54. URL http://dx.doi.org/10.3847/1538-4357/adc44d

[45] K. Murase, S. S. Kimura, P. Mészáros, *Hidden Cores of Active Galactic Nuclei as the Origin of Medium-Energy Neutrinos: Critical Tests with the MeV Gamma-Ray Connection*, PRL 125 (1). URL http://dx.doi.org/10.1103/PhysRevLett.125.011101

[46] J. Frank, A. King, D. Raine, Accretion Power in Astrophysics, 3rd Edition, Cambridge University Press, 2002, ISBN: 9781139164245.

[47] MAGIC Collaboration, *The MAGIC Telescope* (2005). URL https://arxiv.org/abs/astro-ph/0512184

[48] IceCube Collaboration, *Time-Integrated Neutrino Source Searches with 10 Years of IceCube Data*, PRL 124 (5). URL http://dx.doi.org/10.1103/PhysRevLett.124.051103

[49] C. M. Telesco, R. Decher, *The Infrared Structure and the Origin of the Starburst Disk in NGC 1068*, ApJ 334 (1988) 573. URL dx.doi.org/10.1086/166861

[50] A. Loeb, E. Waxman, *The cumulative background of high energy neutrinos from starburst galaxies*, JCAP 2006 (05) (2006) 003–003. URL http://dx.doi.org/10.1088/1475-7516/2006/05/003

[51] P. Kumar, B. Zhang, *The physics of gamma-ray bursts & relativistic jets*, Physics Reports 561 (2015) 1–109. URL https://doi.org/10.1016/j.physrep.2014.09.008

[52] K. Murase, K. Kashiyama, P. Mészáros, *Subphotospheric Neutrinos from Gamma-Ray Bursts: The Role of Neutrons*, Phys. Rev. Lett. 111 (2013) 131102. URL https://link.aps.org/doi/10.1103/PhysRevLett.111.131102

[53] T. Piran, *Gamma-ray bursts and the fireball model*, Physics Reports 314 (6) (1999) 575–667. URL http://dx.doi.org/10.1016/S0370-1573(98)00127-6

[54] IceCube Collaboration, *An All-Sky Search for Three Flavors of Neutrinos from Gamma-Ray Bursts with the IceCube Neutrino Observatory*, ApJ 824 (2) (2016) 115. URL https://doi.org/10.3847/0004-637X/824/2/115

[55] K. Murase, et al., *High-energy Neutrino and Gamma-Ray Emission from Tidal Disruption Events*, ApJ 902 (2) (2020) 108. URL http://dx.doi.org/10.3847/1538-4357/abb3c0

[56] P.-W. Chang, B. Zhou, K. Murase, M. Kamionkowski, *High-energy neutrinos from choked-jet supernovae: Searches and implications*, Phys. Rev. D 109 (2024) 103041. URL https://link.aps.org/doi/10.1103/PhysRevD.109.103041

[57] B. Pontecorvo, *Mesonium and anti-mesonium*, Sov. Phys. JETP 6 (1957) 429, [Zh. Eksp. Teor. Fiz. 33, 549 (1957)].

[58] Z. Maki, M. Nakagawa, S. Sakata, *Remarks on the Unified Model of Elementary Particles*, Progress of Theoretical Physics 28 (5) (1962) 870–880. URL https://doi.org/10.1143/PTP.28.870

[59] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, Sov. Phys. JETP 26 (1968) 984–988, [Zh. Eksp. Teor. Fiz. 53, 1717 (1967)].

[60] N. J. Benoit, et al., *Determination of Majorana-Type Phases from the Time Evolution of Lepton Numbers*, PTEP 2025 (4) (2025) 043B02. URL https://doi.org/10.1093/ptep/ptaf039

[61] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, T. Schwetz, *NuFit-6.0: updated global analysis of three-flavor neutrino oscillations*, JHEP 2024 (12). URL http://dx.doi.org/10.1007/JHEP12(2024)216

[62] N. Savla, *Classical Reconstruction of the PMNS Matrix Using a Mechanical Neutrino Oscillator* (2025). URL https://arxiv.org/abs/2512.01851

[63] A. V. Ivanchik, O. A. Kurichin, V. Y. Yurchenko, *Neutrino at Different Epochs of the Friedmann Universe*, Universe 10 (4). URL https://www.mdpi.com/2218-1997/10/4/169

[64] W. Rodejohann, *Neutrinoless double-beta decay and neutrino physics*, Physics G 39 (12) (2012) 124008. URL https://doi.org/10.1088/0954-3899/39/12/124008

[65] V. Barger, D. Marfatia, K. Whisnant, *Breaking eightfold degeneracies in neutrino CP violation, mixing, and mass hierarchy*, Phys. Rev. D 65 (7). URL http://dx.doi.org/10.1103/PhysRevD.65.073023

[66] L. Visinelli, P. Gondolo, *Neutrino Oscillations and Decoherence* (2008). URL https://arxiv.org/abs/0810.4132

[67] H. Athar, M. Jeżabek, O. Yasuda, *Effects of neutrino mixing on high-energy cosmic neutrino flux*, Phys. Rev. D 62 (2000) 103007. URL https://link.aps.org/doi/10.1103/PhysRevD.62.103007

[68] IceCube Collaboration, *Characterization of the Three-Flavor Composition of Cosmic Neutrinos with IceCube* (2025). URL https://arxiv.org/abs/2510.24957

[69] M. Ackermann, et al., *High-energy and ultra-high-energy neutrinos: A Snowmass white paper*, JHEAP 36 (2022) 55–110, Preprint 2203.08096. URL https://doi.org/10.1016/j.jheap.2022.08.001

[70] P. Lipari, M. Lusignoli, D. Meloni, *Flavor composition and energy spectrum of astrophysical neutrinos*, Phys. Rev. D 75 (2007) 123005. URL https://link.aps.org/doi/10.1103/PhysRevD.75.123005

[71] M. Bustamante, J. F. Beacom, W. Winter, *Theoretically Palatable Flavor Combinations of Astrophysical Neutrinos*, PRL 115 (16). URL http://dx.doi.org/10.1103/PhysRevLett.115.161302

[72] J. A. Formaggio, G. P. Zeller, *From eV to EeV: Neutrino cross sections across energy scales*, RMP 84 (3) (2012) 1307–1341. URL http://dx.doi.org/10.1103/RevModPhys.84.1307

[73] M. A. Unland Elorrieta, *Development, simulation, and characterisation of a novel multi-PMT optical module for IceCube Upgrade with emphasis on detailed understanding of photomultiplier performance parameters*, Ph.D. thesis, Universität Münster (2023). URL https://zenodo.org/doi/10.5281/zenodo.8121321

[74] C. G. Callan, D. J. Gross, *Bjorken Scaling in Quantum Field Theory*, Phys. Rev. D 8 (1973) 4383–4394. URL https://link.aps.org/doi/10.1103/PhysRevD.8.4383

[75] IceCube Collaboration, *Measurement of the inelasticity distribution of neutrino-nucleon interactions for* 80 GeV $< E_\nu <$ 560 GeV *with IceCube DeepCore*, Phys. Rev. D 111 (2025) 112001. URL https://link.aps.org/doi/10.1103/PhysRevD.111.112001

[76] R. Gandhi, C. Quigg, M. H. Reno, I. Sarcevic, *Neutrino interactions at ultrahigh energies*, Phys. Rev. D 58 (9). URL http://dx.doi.org/10.1103/PhysRevD.58.093009

[77] A. Cooper-Sarkar, P. Mertsch, S. Sarkar, *The high energy neutrino cross-section in the Standard Model and its uncertainty*, JHEP 2011 (8). URL http://dx.doi.org/10.1007/JHEP08(2011)042

[78] G.-y. Huang, M. Lindner, N. Volmer, *Inferring astrophysical neutrino sources from the Glashow resonance*, JHEP 2023 (11). URL http://dx.doi.org/10.1007/JHEP11(2023)164

[79] IceCube Collaboration, *Detection of a particle shower at the Glashow resonance with IceCube*, Nature 591 (7849) (2021) 220–224. URL http://dx.doi.org/10.1038/s41586-021-03256-1

[80] A. M. Dziewonski, D. L. Anderson, *Preliminary reference Earth model*, Physics of the Earth and Planetary Interiors 25 (4) (1981) 297–356. URL https://www.sciencedirect.com/science/article/pii/0031920181900467

[81] IceCube Collaboration, *Measurement of the multi-TeV neutrino interaction cross-section with IceCube using Earth absorption*, Nature 551 (7682) (2017) 596–600. URL http://dx.doi.org/10.1038/nature24459

[82] L. Wolfenstein, *Neutrino oscillations in matter*, Phys. Rev. D 17 (1978) 2369–2374. URL https://link.aps.org/doi/10.1103/PhysRevD.17.2369

[83] A. Y. Smirnov, *The Mikheyev-Smirnov-Wolfenstein (MSW) Effect* (2019). URL https://arxiv.org/abs/1901.11473

[84] E. K. Akhmedov, *Parametric resonance in neutrino oscillations in matter*, Pramana 54 (2000) 47–63. URL https://doi.org/10.1007/s12043-000-0006-4

[85] J. Krishnamoorthi, A. K. Upadhyay, A. Kumar, S. K. Agarwalla, *Probing Earth's core using atmospheric neutrino oscillations in the presence of NSI at INO-ICAL*, JHEP 10 (2025) 213, Preprint 2507.02167. URL https://doi.org/10.1007/JHEP10(2025)213

[86] P. B. Denton, R. Pestes, *Neutrino oscillations through the Earth's core*, Phys. Rev. D 104 (2021) 113007. URL https://link.aps.org/doi/10.1103/PhysRevD.104.113007

[87] Particle Data Group Collaboration, *Review of Particle Physics*, Phys. Rev. D 110 (2024) 030001. URL https://link.aps.org/doi/10.1103/PhysRevD.110.030001

[88] I. Safa, et al., *Observing EeV neutrinos through Earth: GZK and the anomalous ANITA events*, JCAP 2020 (01) (2020) 012–012. URL http://dx.doi.org/10.1088/1475-7516/2020/01/012

[89] R. E. JENNINGS, *Čerenkov radiation*, Science Progress (1933-) 50 (199) (1962) 364–375. URL http://www.jstor.org/stable/43425324

[90] D. Groom, S. Klein, *Passage of particles through matter*, EPJ C 15 (2000) 163–173. URL https://doi.org/10.1007/BF02683419

[91] M. Unland, *Studies on dark rates induced by radioactive decays of the multi-PMT digital optical module for future IceCube extensions*, Master's thesis, Universität Münster (2017). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/1712-ma_munland.pdf

[92] I. M. Frank, I. E. Tamm, *Coherent visible radiation of fast electrons passing through matter*, Compt. Rend. Acad. Sci. URSS 14 (3) (1937) 109–114.

[93] N. W. H. Kolanoski, Teilchendetektoren - Grundlagen und Anwendungen, Springer Spektrum, Berlin, Heidelberg, Germany, 2016, ISBN: 978-3-662-45349-0.

[94] IceCube Collaboration, *The design and performance of IceCube DeepCore*, Astroparticle Physics 35 (10) (2012) 615–624. URL http://dx.doi.org/10.1016/j.astropartphys.2012.01.004

[95] IceCube Collaboration, *Optical properties of deep glacial ice at the South Pole*, Journal of Geophysical Research: Atmospheres 111 (D13). URL http://dx.doi.org/10.1029/2005JD006687

[96] IceCube Collaboration, *Measurement of South Pole ice transparency with the IceCube LED calibration system*, NIM-A 711 (2013) 73–89. URL https://www.sciencedirect.com/science/article/pii/S0168900213001460

[97] IceCube Collaboration, *The IceCube Neutrino Observatory Part VI: Ice Properties, Reconstruction and Future Developments* (2013). URL https://arxiv.org/abs/1309.7010

[98] IceCube Collaboration, *IceTop: The surface component of IceCube*, NIM-A 700 (2013) 188–220. URL http://dx.doi.org/10.1016/j.nima.2012.10.067

[99] IceCube Collaboration, *A muon-track reconstruction exploiting stochastic losses for large-scale Cherenkov detectors*, JINST 16 (08) (2021) P08034. URL https://doi.org/10.1088/1748-0221/16/08/P08034

[100] IceCube Collaboration, *Energy reconstruction methods in the IceCube neutrino telescope*, JINST 9 (03) (2014) P03009. URL https://doi.org/10.1088/1748-0221/9/03/P03009

[101] IceCube Collaboration, *A convolutional neural network based cascade reconstruction for the IceCube Neutrino Observatory*, JINST 16 (07) (2021) P07041. URL https://doi.org/10.1088/1748-0221/16/07/P07041

[102] IceCube Collaboration, *Improved modeling of in-ice particle showers for IceCube event reconstruction*, JINST 19 (06) (2024) P06026. URL https://doi.org/10.1088/1748-0221/19/06/P06026

[103] IceCube Collaboration, *Observation of Seven Astrophysical Tau Neutrino Candidates with IceCube*, Phys. Rev. Lett. 132 (2024) 151001. URL https://link.aps.org/doi/10.1103/PhysRevLett.132.151001

[104] Anna Franckowiak, *Searching for High-energy Neutrinos from Supernovae with IceCube and an Optical Follow-up Program*, Ph.D. thesis, Rheinische Friedrich-Wilhelms-Universität Bonn (Oct. 2011). URL https://hdl.handle.net/20.500.11811/5047

[105] S. Iyer Dutta, M. H. Reno, I. Sarcevic, D. Seckel, *Propagation of muons and taus at high energies*, Phys. Rev. D 63 (9). URL http://dx.doi.org/10.1103/PhysRevD.63.094020

[106] IceCube Collaboration, *Physics potential of the IceCube Upgrade for atmospheric neutrino oscillations*, Phys. Rev. D 113 (2026) 072009. URL https://journals.aps.org/prd/abstract/10.1103/nnjw-jp1n

[107] G. Jungman, M. Kamionkowski, K. Griest, *Supersymmetric dark matter*, Physics Reports 267 (5) (1996) 195–373. URL https://doi.org/10.1016/0370-1573(95)00058-5

[108] D. Heinen, Shefali, R. Turcotte, L. S. Weinstock, C. Wiebusch, S. Zierke, *An Acoustic Calibration System for the IceCube Upgrade* (2019). URL https://arxiv.org/abs/1909.02047

[109] W. Kang, J. Lee, S. Rodan, C. Rott, C. Tönnis, *The Camera System for the IceCube Upgrade*, Physical Sciences Forum 8 (1). URL https://www.mdpi.com/2673-9984/8/1/49

[110] M. Rongen, D. Chirkin, *Advances in IceCube ice modelling & what to expect from the Upgrade*, JINST 16 (09) (2021) C09014. URL https://doi.org/10.1088/1748-0221/16/09/C09014

[111] IceCube Collaboration, *Mapping the ice stratigraphy in IceCube using camera deployment footage*, PoS ICRC2025 (2025) 1035. URL https://pos.sissa.it/501/1035/

[112] IceCube Collaboration, *Performance Studies of the Acoustic Module for the IceCube Upgrade*, PoS ICRC2023 (2023) 241. URL https://pos.sissa.it/444/241/

[113] N. Khera, F. Henningsen, *POCAM in the IceCube Upgrade*, PoS ICRC2021 (2021) 1049. URL https://pos.sissa.it/395/1049/

[114] C. J. Lozano Mariscal, *Ice calibration studies for IceCube Upgrade with multi-PMT optical modules, and studies of the sensitivity of future IceCube extensions to astrophysical neutrinos in the MeV and GeV range*, PhD thesis, Universität Münster (2023). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/doktorarbeiten/phd_definitivo.pdf

[115] P. B. Price, et al., *Temperature profile for glacial ice at the South Pole: Implications for life in a nearby subglacial lake*, Proceedings of the National Academy of Sciences 99 (12) (2002) 7844–7847. URL https://doi.org/10.1073/pnas.082238999

[116] L. Classen, *The mDOM - a multi-PMT digital optical module for the Icecube-Gen2 neutrino telescope*, PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg (2017). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/doktorarbeiten/1702-phd_lclassen.pdf

[117] IceCube Collaboration, *Design and performance of the multi-PMT optical module for IceCube Upgrade*, PoS ICRC2023 (2021) 1070. URL https://pos.sissa.it/395/1070/pdf

[118] IceCube Collaboration, *D-Egg: a dual PMT optical module for IceCube*, JINST 18 (04) (2023) P04014. URL http://dx.doi.org/10.1088/1748-0221/18/04/P04014

[119] A. Ishihara, A. Stoessl, S. Shimizu, et al., *Overview and performance of the D-Egg optical sensor for IceCube-Gen2*, PoS ICRC2017 (2017) 1051. URL https://pos.sissa.it/301/1051/

[120] IceCube Collaboration, *Camera Calibration for the IceCube Upgrade and Gen2*, PoS ICRC2021 (2021) 1064. URL https://pos.sissa.it/395/1064/

[121] A. Ishihara, A. Kiriki, *Calibration LEDs in the IceCube Upgrade DEgg modules*, PoS ICRC2019 (2019) 923. URL https://pos.sissa.it/358/0923/

[122] KM3NeT Collaboration, *KM3NeT Conceptual Design for a Deep-Sea Research Infrastructure Incorporating a Very Large Volume Neutrino Telescope in the Mediterranean Sea* (2008). URL https://www.km3net.org/wp-content/uploads/2023/05/KM3NeT_DS_2008_002-CDR.pdf

[123] J. Rack-Helleis, A. Pollmann, M. Rongen, *The Wavelength-shifting Optical Module (WOM) for the IceCube Upgrade*, PoS ICRC2021 (2021) 1038. URL https://pos.sissa.it/395/1038/pdf

[124] P. Sandstrom (IceCube-PINGU Collaboration Collaboration), *Digital optical module design for PINGU*, AIP Conference Proceedings 1630 (1) (2014) 180–183. URL https://doi.org/10.1063/1.4902801

[125] T. Benson, et al., *IceCube Enhanced Hot Water Drill functional description*, Annals of Glaciology 55 (68) (2014) 105–114. URL https://doi.org/10.3189/2014AoG68A032

[126] IceCube-Gen2 Collaboration, *The IceCube-Gen2 project*, PoS ICRC2025 (2025) 1080. URL https://pos.sissa.it/501/1080/

[127] IceCube Collaboration, *Compact Optical Module with Enhanced Detection Efficiency for the Next-Generation Neutrino Telescope at the South Pole*, in preparation, JINST.

[128] F. G. Schröder, *Cosmic-Ray Physics in the PeV to EeV Energy Range with the IceCube-Gen2 Surface Array*, PoS ICRC2025 (2025) 387. URL https://pos.sissa.it/501/387/pdf

[129] C. Glaser, *Probing ultra-high-energy neutrinos with the IceCube-Gen2 in-ice radio array*, PoS ICRC2025 (2025) 1045. URL https://pos.sissa.it/501/1045/

[130] P. W. Gorham, et al., *Observations of the Askaryan Effect in Ice*, PRL 99 (17). URL http://dx.doi.org/10.1103/PhysRevLett.99.171101

[131] J. Alvarez-Muñiz, P. M. Hansen, A. Romero-Wolf, E. Zas, *Askaryan radiation from neutrino-induced showers in ice*, Phys. Rev. D 101 (8). URL http://dx.doi.org/10.1103/PhysRevD.101.083005

[132] S. Verpoest, C. Merx, B. Flaggs, *A prototype station of the IceCube-Gen2 Surface Array at the Pierre Auger Observatory*, PoS ICRC2025 (2025) 428. URL https://pos.sissa.it/501/428

[133] S.-O. Flyckt, C. Marmonier, Photomultiplier tubes - Principles and applications, Photonis, 2002.

[134] Hamamatsu Photonics K.K., Photomultiplier tubes - Basics and Applications, 3rd Edition, Hamamatsu Photonics K.K., 2007. URL https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/PMT_handbook_v4E.pdf

[135] A. Wright, The Photomultiplier Handbook, Oxford University Press, 2017, ISBN: 9780199565092.

[136] J. Averbeck, *Nachbau eines Photomultipliers in COMSOL Multiphysics und Untersuchung der Elektronentrajektorien zwischen Photokathode und erster Dynode*, Bachelor's thesis, Universität Münster (2022). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/bachelor_janis.pdf

[137] W. Achtermann, *Replication of a photomultiplier in COMSOL Multiphysics to study PMT performance parameters for the mDOM in IceCube*, Master's thesis, Universität Münster (2023). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/masterarbeit_willem_achtermann.pdf

[138] A. Puke, *Modellierung und Untersuchung der Sekundärelektronenemission in einem Photomultiplier mit COMSOL Multiphysics*, Bachelor's thesis, Universität Münster (2024). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/ba_puke.pdf

[139] North Night Vision Technology, private communication.

[140] J. Rademacker, *An exact formula to describe the amplification process in a photomultiplier tube*, NIM-A 484 (1-3) (2002) 432–443. URL https://doi.org/10.1016/s0168-9002(01)02055-1

[141] W. E. Spicer, *Photoemissive, Photoconductive, and Optical Absorption Studies of Alkali-Antimony Compounds*, Phys. Rev. 112 (1958) 114–122. URL https://link.aps.org/doi/10.1103/PhysRev.112.114

[142] W. E. Spicer, A. Herrera-Gómez, *Modern theory and applications of photocathodes*, Tech. rep., SLAC, Stanford, CA (1993). URL https://cds.cern.ch/record/257133

[143] M. Suyama, K. Nakamura, *Recent Progress of Photocathodes for PMTs*, PoS International Workshop on New Photon Detectors (2009) 013. URL https://pos.sissa.it/090/013/pdf

[144] E. H. Bellamy, et al., *Absolute calibration and monitoring of a spectrometric channel using a photomultiplier*, NIM-A 339 (1994) 468–476. URL https://doi.org/10.1016/0168-9002(94)90183-X

[145] S. Dushman, *Electron Emission from Metals as a Function of Temperature*, Phys. Rev. 21 (1923) 623–636. URL https://link.aps.org/doi/10.1103/PhysRev.21.623

[146] B. Lubsandorzhiev, R. Vasiliev, Y. Vyatchin, B. Shaibonov, *Photoelectron backscattering in vacuum phototubes*, NIM-A 567 (1) (2006) 12–16. URL https://www.sciencedirect.com/science/article/pii/S0168900206008606

[147] L. Classen, *The mDOM – A multi-PMT Digital Optical Module for the IceCube-Gen2 neutrino telescope*, PoS ICRC2017 (2017) 1047. URL https://pos.sissa.it/301/1047/

[148] Y. Makino, *Mechanical design of the optical modules intended for IceCube-Gen2*, PoS ICRC2023 (2023) 979. URL https://pos.sissa.it/444/979/pdf

[149] S. Griffin, *Electronics Design of the IceCube-Gen2 Optical Module Prototype*, PoS ICRC2023 (2023) 1039. URL https://pos.sissa.it/444/1039/pdf

[150] A. Eimer, M. Rongen, *A co-deployed dust-logging instrument for the IceCube Upgrade and IceCube-Gen2*, PoS ICRC2025 (2025) 1034. URL https://pos.sissa.it/501/1034/pdf

[151] A. Kappes, *The Optical Sensor for IceCube-Gen2*, PoS ICRC2025 (2025) 1072. URL https://pos.sissa.it/501/1072/

[152] M. Dittmer, A. Kappes, *Performance studies on new 4″ photomultiplier types intended for IceCube-Gen2 optical modules*, PoS ICRC2023 (2023) 1083. URL https://pos.sissa.it/444/985/

[153] A. S. Singh, A. G. Wright, *The Determination of Photomultiplier Temperature Coefficients for Gain and Spectral Sensitivity Using the Photon Counting Technique*, IEEE Transactions on Nuclear Science 34 (1) (1987) 434–437. URL https://doi.org/10.1109/TNS.1987.4337379

[154] Aya Ishihara, private communication.

[155] L. Y. Halve, *Validation of optical sensors for the mDOM of the IceCube Upgrade*, Ph.D. thesis, RWTH Aachen University (2024). URL https://publications.rwth-aachen.de/record/986971

[156] R. Busse, *Setup and commissioning of a test stand for detailed investigations of quantum efficiency characteristics of photomultiplier tubes, and initial studies for IceCube-Gen2*, Master's thesis, Universität Münster (2017). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/1703-ma_rbusse.pdf

[157] F. Krafft, *Studies on the Temperature Characteristics of the Detection Efficiency of Hamamatsu R15458-02-Type Photomultipliers*, Bachelor's thesis, Universität Münster (2024). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/ba_krafft.pdf

[158] T. DeYoung, *IceTray: A software framework for IceCube*, in: 14th International Conference on Computing in High-Energy and Nuclear Physics, 2005, pp. 463–466. URL https://cds.cern.ch/record/865626/files/p463.pdf

[159] IceCube Collaboration, *In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes*, JINST 15 (06) (2020) P06032–P06032. URL http://dx.doi.org/10.1088/1748-0221/15/06/P06032

[160] Yuya Makino, private communication.

[161] J. Haser, et al., *Afterpulse measurements of R7081 photomultipliers for the Double Chooz experiment*, JINST 8 (04) (2013) P04029–P04029. URL http://dx.doi.org/10.1088/1748-0221/8/04/P04029

[162] M. Unland Elorrieta, R. Busse, L. Classen, A. Kappes, *Homogeneity of the photocathode in the Hamamatsu R15458-02 Photomultiplier Tube*, JINST 16 (11) (2021) P11038. URL https://doi.org/10.1088/1748-0221/16/11/P11038

[163] K. Ueberholz, *Setup and commission of a magnetic fieldtest stand and studies on the influence of magnetic fields on the perfomance of the Hamamatsu R12199-02*, Master's thesis, Universität Münster (2020). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/ueberholz_master.pdf

[164] C. Hemming, *Aufbau und Inbetriebnahme eines Messstandes zur Untersuchung der wellenlängenabhängigen, ortsaufgelösten Quanteneffizienz von Photomultipliern*, Bachelor's thesis, Universität Münster (2024). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/2402-ba_chemming.pdf

[165] A. Chulliat, et al., *The US/UK World Magnetic Model for 2025-2030: Technical Report*, NOAA NCEI. URL https://repository.library.noaa.gov/view/noaa/71569

[166] I. Bahloul, *Aufbau und Inbetriebnahme eines Messstandes zur Untersuchung der wellenlängenabhängigen optischen Eigenschaften von Komponenten des mDOM-Sensors*, Bachelor's thesis, Universität Münster (2021). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/2109-ba_ibahloul.pdf

[167] P. Peiffer, *Overview and Performance of the Wavelength-shifting Optical Module (WOM)*, PoS ICRC2017 (2017) 1052. URL https://pos.sissa.it/301/1052/

[168] Steffen Pausch, Nautilus Marine Service GmbH, private communication.

[169] S. Agostinelli, et al., *Geant4—a simulation toolkit*, NIM-A 506 (3) (2003) 250 – 303. URL https://doi.org/10.1016/S0168-9002(03)01368-8

[170] J. Allison, et al., *Recent developments in Geant4*, NIM-A 835 (2016) 186–225. URL https://doi.org/10.1016/j.nima.2016.06.125

[171] IceCube Collaboration, *OMSim: Geant4 for IceCube optical module studies*, accessed: 2026-02-22. URL https://icecube.github.io/OMSim/index.html

[172] B. Herold, *Simulation and measurement of optical background in the deep sea using a multi-PMT optical module*, PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) (2017).

[173] D. Suarez, *Modeling of a multi-PMT optical sensor for IceCube-Gen2 in GEANT4 and Monte Carlo studies to optimize its sensitivity*, Master's thesis, Universität Münster (2021). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/masterthesis_danielsuarez.pdf

[174] F. Vara, *Gel pad sensitivity studies and muon reconstruction in a water basin in Geant4 for IceCube multi-PMT optical modules*, Master's thesis, Universität Münster (2022). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/master_thesis_vara_carbonell_506864.pdf

[175] C. M. Poole, et al., *A CAD interface for GEANT4*, Australasian Physical & Engineering Science in Medicine. URL https://doi.org/10.1007/s13246-012-0159-8

[176] N. Krybus, *Thin Layer Optics in Photomultipliers: A Geant4 Simulation and Comparison to Ellipsometric Measurements*, Master thesis, Universität Münster (2024). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/ma_krybus.pdf

[177] T. Brand, *Measurement of the optical properties of PMT components and investigation of their influence on the simulated PMT response*, Bachelor's thesis, Universität Münster (2025). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/2509_ba_tbrand.pdf

[178] K. M. Gorski, et al., *HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere*, ApJ 622 (2) (2005) 759–771. URL https://doi.org/10.1086/427976

[179] M. Schlechtriem, *Simulationsstudien zum Einfluss der Modulaufhängung auf die Photon-Sensitivität des mDOM-Sensors im IceCube Upgrade*, Bachelor's thesis, Universität Münster (2021). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/ba_moritz_schlechtriem.pdf

[180] IceCube Collaboration, *Measurement of South Pole ice transparency with the IceCube LED calibration system*, NIM-A 711 (2013) 73–89. URL http://dx.doi.org/10.1016/j.nima.2013.01.054

[181] IceCube Collaboration, *DOUMEKI: Digital Optical Universal Module Emulation KIt - Geant4 based DOM full simulator*, accessed: 2025-02-23. URL https://github.com/icecube/doumeki/tree/main

[182] L. Barthelmeß, *Training and application of a neural network to calculate PMT hit probabilities for incident photons in mDOMs for IceCube Upgrade*, Bachelor's thesis, Universität Münster (2025). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/bachelorarbeiten/2509_ba_lbarthelmess.pdf

[183] M. Dittmer, *Characterisation of scintillation light induced by radioactive excitation in the mDOM glass pressure vessel*, Master's thesis, Universität Münster (2020). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/dittmer.pdf

[184] J. B. Birks, The Theory and Practice of Scintillation Counting, Pergomon press, 1964. URL https://doi.org/10.1016/c2013-0-01791-4

[185] C. R. Ronda (Ed.), Luminescence: From Theory to Applications, Wiley-VCH, 2007, ISBN: 9783527314024.

[186] A. Jha, Inorganic Glasses for Photonics: Fundamentals, Engineering, and Applications (Wiley Series in Materials for Electronic & Optoelectronic Applications), Wiley, 2016, ISBN: 9780470741702.

[187] C. Ronda, *Emission and Excitation Mechanisms of Phosphors*, Luminescence: From Theory to Applications (2007) 1–34. URL https://doi.org/10.1002/9783527621064.ch1

[188] R. Salh, *Defect Related Luminescence in Silicon Dioxide Network: A Review*, in: Crystalline Silicon - Properties and Uses, InTech, 2011. URL https://doi.org/10.5772/22607

[189] Y. Shu, B. S. Fales, W.-T. Peng, B. G. Levine, *Understanding Nonradiative Recombination through Defect-Induced Conical Intersections*, J. Phys. Chem. Lett. 8 (17) (2017) 4091–4099. URL https://doi.org/10.1021/acs.jpclett.7b01707

[190] M. S. Akselrod, N. A. Larsen, V. Whitley, S. W. S. McKeever, *Thermal quenching of F-center luminescence in Al2O3:C*, J. Appl. Phys. 84 (6) (1998) 3364–3373. URL https://doi.org/10.1063/1.368450

[191] W. W. Moses, et al., *Scintillator non-proportionality: Present understanding and future challenges*, IEEE Transactions on Nuclear Science 55 (3) (2008) 1049–1053. URL https://doi.org/10.1109/TNS.2008.922802

[192] G. Knoll, Radiation Detection and Measurement (4th ed.), John Wiley, Hoboken, NJ, 2010, ISBN: 978-0-470-13148-0.

[193] G. R. Gilmore, Practical Gamma-Ray Spectrometry, John Wiley & Sons, Ltd, 2008, ISBN: 9780470861981. URL https://doi.org/10.1002/9780470861981.fmatter

[194] G. Heusser, *Low-Radioactivity Background Techniques*, Annu. Rev. Nucl. Part. Sci. 45 (2003) 543–590. URL https://doi.org/10.1146/annurev.ns.45.120195.002551

[195] National Nuclear Data Center, *NuDat 3 database*, https://www.nndc.bnl.gov/nudat3/, accessed: 2026-04-23 (2023).

[196] International Atomic Energy Agency (IAEA), *Live chart of nuclides*, https://www-nds.iaea.org/relnsd/vcharthtml/VChartHTML.html, accessed: 2026-04-23 (2023).

[197] J. Magill, J. Galy, Radioactivity Radionuclides Radiation, Springer, 2005, ISBN: 3-540-21116-0. URL https://doi.org/10.1007/b138236

[198] Martin Unland, private communication.

[199] H. Pfaender, Schott Guide to Glass, Springer Netherlands, 1996, ISBN: 0412719606, 9780412719608.

[200] A. Talimian, V. M. Sglavo, *Ion-exchange strengthening of borosilicate glass: Influence of salt impurities and treatment temperature*, Journal of Non-Crystalline Solids 456 (2017) 12–21. URL https://www.sciencedirect.com/science/article/pii/S0022309316304732

[201] K. Weise, et al., Determination of the detection limit and decision threshold for ionizing-radiation measurements: fundamentals and particular applications, TÜV Media GmbH, 2005.

[202] M. Naumenko-Dèzes, K. Mezger, T. Nägler, I. Villa, *High precision determination of the terrestrial 40K abundance*, Geochimica et Cosmochimica Acta 122 (2013) 353–362. URL https://doi.org/10.1016/j.gca.2013.08.019

[203] W. Laue, M. Thiemann, E. Scheibler, K. W. Wiegand, *Nitrates and nitrites*, in: Ullmann's Encyclopedia of Industrial Chemistry, Wiley-VCH Verlag GmbH & Co. KGaA, 2000. URL https://doi.org/10.1002/14356007.a17_265

[204] E. Aprile, et al., *Lowering the radioactivity of the photomultiplier tubes for the XENON1T dark matter experiment*, EPJ C 75 (11). URL http://dx.doi.org/10.1140/epjc/s10052-015-3657-5

[205] R. A. Rahimi, A. Ahmadi, S. Kakooei, S. Sadrnezhaad, *Corrosion behavior of ZrO2–SiO2–Al2O3 refractories in lead silicate glass melts*, Journal of the European Ceramic Society 31 (5) (2011) 715–721. URL https://www.sciencedirect.com/science/article/pii/S0955221910005418

[206] E. Hecht, Optics, 5th Edition, Pearson Education, 2016, ISBN: 978-0133977226.

[207] X. Yang, et al., *Temperature dependence of an optical narrow-bandpass filter at 1.5μm*, Appl. Opt. 54 (1) (2015) 96–100. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-54-1-96

[208] Newport, *Stabilife coating, technical notes*, https://www.newport.com.cn/n/stabilife-coating, accessed: 2026-04-23.

[209] A. C. DeFranzo, B. G. Pazol, *Index of refraction measurement on sapphire at low temperatures and visible wavelengths*, Appl. Opt. 32 (13) (1993) 2224–2234. URL https://doi.org/10.1364/AO.32.002224

[210] J. A. R. Samson, Techniques of Vacuum Ultraviolet Spectroscopy, John Wiley & Sons, New York, 1967, ISBN: 978-0471751199.

[211] Hamamatsu Photonics K.K., private communication.

[212] A. Lyashenko, et al., *Measurement of the absolute Quantum Efficiency of Hamamatsu model R11410-10 photomultiplier tubes at low temperatures down to liquid xenon boiling point*, JINST 9 (11) (2014) P11021–P11021. URL http://dx.doi.org/10.1088/1748-0221/9/11/P11021

[213] I. Renge, *Impurity spectroscopy in glasses and disordered crystals: Inhomogeneous broadening and electron phonon coupling*, Journal of Luminescence 128 (3) (2008) 413–420. URL https://doi.org/10.1016/j.jlumin.2007.09.009

[214] J. Mooney, P. Kambhampati, *Get the basics right: Jacobian conversion of wavelength and energy scales for quantitative analysis of emission spectra*, J. Phys. Chem. 4 (19) (2013) 3316–3318. URL https://doi.org/10.1021/jz401508t

[215] N. M. Feigl, *Noise characterization and performance evaluation of new optical sensors for the IceCube Upgrade*, Ph.D. thesis, DESY/HU (2026). URL https://doi.org/10.18452/37266

[216] Martin Unland, *OMSim documentation for Okamoto yield*, accessed: 2026-06-18. URL https://github.com/icecube/OMSim/tree/master/documentation/notebooks/scintillation/OkamotoGlass/Yield

[217] Colton Hill, private communication.

[218] A.-S. Tenbruck, *Local Coincidence Studies for a Multi-PMT Optical Module for IceCube Upgrade*, Master's thesis, Universität Münster (2025). URL https://www.uni-muenster.de/imperia/md/content/physik_kp/agkappes/abschlussarbeiten/masterarbeiten/2512-ma_astenbruck.pdf

[219] NASA/JPL-Caltech, *Hidden Lair at the Heart of Galaxy NGC 1068*, https://www.jpl.nasa.gov/images/pia20058-hidden-lair-at-the-heart-of-galaxy-ngc-1068/, accessed: 2026-01-08.

[220] García-Burillo, et al., *ALMA Resolves the Torus of NGC 1068: Continuum and Molecular Line Emission*, ApJL 823 (1) (2016) L12. URL https://doi.org/10.3847/2041-8205/823/1/L12

[221] B. Eichmann, et al., *Solving the Multimessenger Puzzle of the AGN-starburst Composite Galaxy NGC 1068*, ApJ 939 (1) (2022) 43. URL http://dx.doi.org/10.3847/1538-4357/ac9588

[222] T. Michiyama, Y. Inoue, A. Doi, D. Khangulyan, *ALMA Detection of Parsec-scale Blobs at the Head of a Kiloparsec-scale Jet in the Nearby Seyfert Galaxy NGC 1068*, ApJL 936 (1) (2022) L1. URL https://doi.org/10.3847/2041-8213/ac8935

[223] IXPE Collaboration, *X-ray polarization measurement of the gold standard of radio-quiet active galactic nuclei: NGC 1068*, A&A 689 (2024) A238. URL https://doi.org/10.1051/0004-6361/202449760

[224] C. Blanco, D. Hooper, T. Linden, E. Pinetti, *Neutrino and gamma-ray emissions from NGC 1068*, Phys. Rev. D 112 (2025) 123016. URL https://link.aps.org/doi/10.1103/wnjh-7nwp

[225] A. Marinucci, et al., *NuSTAR catches the unveiling nucleus of NGC 1068*, Monthly Notices of the Royal Astronomical Society: Letters 456 (1) (2015) L94–L98. URL http://dx.doi.org/10.1093/mnrasl/slv178

[226] IceCube Collaboration, *Evidence for Neutrino Emission from X-Ray-bright Active Galactic Nuclei with IceCube*, ApJL 1000 (1) (2026) L26. URL https://doi.org/10.3847/2041-8213/ae4aad

[227] The IceCube Collaboration and Fermi-LAT and MAGIC and AGILE and ASAS-SN and HAWC and H.E.S.S. and INTEGRAL and Kanata and Kiso and Kapteyn and Liverpool Telescope and Subaru and Swift/NuSTAR and VERITAS and VLA/17B-403 teams, *Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A*, Science 361 (6398) (2018) eaat1378. URL https://www.science.org/doi/abs/10.1126/science.aat1378

[228] IceCube Collaboration, *Characterization of the astrophysical diffuse neutrino flux using starting track events in IceCube*, Phys. Rev. D 110 (2). URL http://dx.doi.org/10.1103/PhysRevD.110.022001

[229] O. Gueta, *The Cherenkov Telescope Array: layout, design and performance*, PoS ICRC2021 (2021) 885. URL https://pos.sissa.it/395/0885

[230] A. Asif, et al., *Design of the ULTRASAT UV camera*, in: O. H. Siegmund (Ed.), UV, X-Ray, and Gamma-Ray Space Instrumentation for Astronomy XXII, SPIE, 2021, p. 21. URL http://dx.doi.org/10.1117/12.2594253

[231] J. Braun, et al., *Methods for point source analysis in high energy neutrino telescopes*, Astroparticle Physics 29 (4) (2008) 299–305. URL https://www.sciencedirect.com/science/article/pii/S0927650508000303

[232] IceCube Collaboration, *`csky`: An unbinned likelihood maximization library optimized for the sorts of analyses commonly performed by the IceCube Neutrino Sources working group*, https://github.com/icecube/csky, internal repository, accessed 2026.

[233] IceCube Collaboration, *Search for 10–1000 GeV Neutrinos from Gamma-Ray Bursts with IceCube*, ApJ 964 (2) (2024) 126. URL https://doi.org/10.3847/1538-4357/ad220b

[234] G. C. Hill, K. Rawlins, *Unbiased cut selection for optimal upper limits in neutrino detectors: the model rejection potential technique*, Astroparticle Physics 19 (3) (2003) 393–402. URL http://dx.doi.org/10.1016/S0927-6505(02)00240-2

[235] G. J. Feldman, R. D. Cousins, *Unified approach to the classical statistical analysis of small signals*, Phys. Rev. D 57 (7) (1998) 3873–3889. URL http://dx.doi.org/10.1103/PhysRevD.57.3873

[236] D. Soldin, *Atmospheric Muons Measured with IceCube*, EPJ Web of Conferences 208 (2019) 08007. URL http://dx.doi.org/10.1051/epjconf/201920808007

[237] IceCube Collaboration, *Measurement of the Atmospheric ν_e Spectrum with IceCube*, Phys. Rev. D 91 (12). URL http://dx.doi.org/10.1103/PhysRevD.91.122004

[238] IceCube Collaboration, *IceCube high-energy starting event sample: Description and flux characterization with 7.5 years of data*, Phys. Rev. D 104 (2). URL http://dx.doi.org/10.1103/PhysRevD.104.022002

[239] S. R. Klein, *Astrophysical Neutrinos with IceCube*, Thirteenth Conference on the Intersections of Particle and Nuclear Physics. URL https://arxiv.org/abs/1809.07873

[240] IceCube Collaboration, *Measurement of atmospheric tau neutrino appearance with IceCube DeepCore*, Phys. Rev. D 99 (3). URL http://dx.doi.org/10.1103/PhysRevD.99.032007

[241] M. Larson, *GRECO Online Dataset*, https://drive.google.com/file/d/11TtexQeWUvd2qefIDH-MxTT5XRcEFSJD/view, Internal presentation in the Nu-Sources call, 26 July 2021.

[242] C. Andreopoulos, et al., *The GENIE neutrino Monte Carlo generator*, NIM-A 614 (1) (2010) 87–104. URL https://www.sciencedirect.com/science/article/pii/S0168900209023043

[243] A. Gazizov, M. Kowalski, *ANIS: High energy neutrino generator for neutrino telescopes*, Computer Physics Communications 172 (3) (2005) 203–213. URL https://www.sciencedirect.com/science/article/pii/S0010465505004194

[244] K.Hoshina, *NuGen III*, https://user-web.icecube.wisc.edu/~hoshina/myslides/2013/NuGen3_Madison.pdf, presentation at the IceCube Collaboration meeting, 10 May 2013.

[245] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, T. Thouw, *CORSIKA: A Monte Carlo code to simulate extensive air showers*, Wissenschaftliche Berichte FZKA 6019. URL https://publikationen.bibliothek.kit.edu/270043064/3813660

[246] J. van Santen, *Neutrino Interactions in IceCube above 1 TeV: Constraints on Atmospheric Charmed-Meson Production and Investigation of the Astrophysical Neutrino Flux with 2 Years of IceCube Data taken 2010–2012*, Ph.D. thesis, Wisconsin U., Madison (11 2014). URL https://asset.library.wisc.edu/1711.dl/GX7HPABPUJ5YO84/R/file-dddd0.pdf

[247] K. DeHolton, *oscNext Muon Background Simulation: Optimizing MuonGun MC*, https://drive.google.com/drive/folders/17kaGplxDrf-V9sVZblTqcSjPeT_H94EY, Internal presentation, 20 August 2018.

[248] M. Larson, *GRECO Processing*, https://drive.google.com/file/d/10thJ7InaMofy3X-vbtqtA5i5f-7Eiqdh/view, Internal presentation in the Low Energy Astronomy call, 27 February 2025.

[249] M. Leuermann, *Testing the neutrino mass ordering with IceCube DeepCore*, Ph.D. thesis, RWTH Aachen University (2018). URL https://publications.rwth-aachen.de/record/751704

[250] P. Eller, K. L. DeHolton, J. Weldert, R. Ørsøe, *Sensitivity of the IceCube Upgrade to Atmospheric Neutrino Oscillations*, PoS ICRC2023 (2023) 1036. URL https://pos.sissa.it/444/1036/

[251] J.-H. Koehne, et al., *PROPOSAL: A tool for propagation of charged leptons*, Computer Physics Communications 184 (9) (2013) 2070–2090. URL https://www.sciencedirect.com/science/article/pii/S0010465513001355

[252] D. Chirkin, et al., *Photon Propagation using GPUs by the IceCube Neutrino Observatory*, in: 2019 15th International Conference on eScience (eScience), IEEE, 2019, pp. 388–393. URL https://ieeexplore.ieee.org/document/9041727

[253] IceCube Collaboration, *Graph Neural Networks for low-energy event classification & reconstruction in IceCube*, JINST 17 (11) (2022) P11003. URL http://dx.doi.org/10.1088/1748-0221/17/11/P11003

[254] J. P. González, *Event Selection for the IceCube Upgrade: Machine Learning-Driven event selection for neutrino physics analysis using with the IceCube Upgrade*, Master's thesis, University of Copenhagen (2023). URL https://nbi.ku.dk/english/theses/masters-theses/jorge-prado-gonzalez/Jorge_Prado_Gonz_lez.pdf

[255] IceCube Collaboration, *Astrophysical Sensitivity Projections for the IceCube Upgrade*, in preparation, PRD.

[256] IceCube Collaboration, *Time-Integrated Searches for Sub-TeV Neutrino Sources with IceCube DeepCore*, in preparation, ApJL.

[257] S. Yu, A. Kheirandish, Q. Liu, H. Niederhausen, *Search for TeV Neutrinos from Seyfert Galaxies in the Southern Sky using Starting Track Events in IceCube*, PoS ICRC2023 (2023) 1533. URL https://pos.sissa.it/444/1533/pdf

[258] J. A. Carpio, et al., *Quasithermal GeV neutrinos from neutron-loaded magnetized outflows in core-collapse supernovae: Spectra and light curves*, Phys. Rev. D 110 (2024) 083012. URL https://link.aps.org/doi/10.1103/PhysRevD.110.083012

[259] S. Mechbal, N. Feigl, *Characterization and testing of the IceCube Upgrade mDOM*, PoS ICRC2023 (2023) 1183. URL https://pos.sissa.it/444/1183/

A Scaling towards Mass Production

While Chapter 7 highlights the primary mechanical developments undertaken during the active Gen2-OM design phase, the full scope of the R&D effort also extended to the logistical optimization required for the IceCube-Gen2 array. The final production and integration of the IceCube Upgrade precursor modules (the LOM16 and LOM18) were carried out by dedicated teams at the University of Wisconsin-Madison and Chiba University, respectively. However, the foundational planning and design phases were shaped by the research presented in this thesis. The component validation, optical simulations and mechanical prototyping detailed in Chapters 6 to 8 directly influenced the module development. Furthermore, even before the assembly of these precursor modules, the conceptual design review (CDR) phase required a defined logistical strategy to scale production from localized manual prototyping to industrial mass production, devised by the author.

At the time of writing, the exact number of global assembly sites remains under review and will likely range between three and five facilities. The IceCube-Gen2 array requires approximately 10 000 optical modules to be manufactured over a seven-year period, with a projected peak production rate of 2000 modules in a single year. Assuming the workload is distributed evenly across three sites, a single facility must reliably produce up to approximately 670 modules annually, corresponding to roughly 14 modules per week. To achieve this sustained throughput without introducing logistical bottlenecks, particularly those associated with the extended curing times of the optical gel, all integration processes must be parallelized into dedicated sequential assembly stations.

PMT Verification and Preparation

Upon retrieval from a dedicated warehouse, incoming PMTs must first undergo dimensional inspection to accommodate tight mechanical tolerances during assembly. A dedicated measurement rig verifies the length, neck diameter and photocathode diameter against manufacturing tolerances, rejecting any outliers. Additional performance characterization is then performed following the procedures described in Chapter 6 using a dedicated multi-PMT test setup, similar to that employed during the mDOM production [155]. Following PMT selection, rotational keying elements (introduced in Section 7.1.2) may be precisely bonded to the PMT necks, depending on the finalized internal layout. Alternatively, the pre-soldered wuBase may serve as the mechanical reference for rotational alignment.

Gel Pad Casting

To maximize throughput, gel pad casting is envisioned as a track-based assembly line. Subassembly trays containing 16 (or 18) thermoformed molds are processed sequentially. A specialized automated dispenser, such as the *MeterMix PAR 30*, injects a metered volume of pre-degassed liquid Shin-Etsu gel into each mold. The prepared PMTs would then be lowered into the molds and mechanically secured in position. The entire tray is subsequently transferred into a large-volume vacuum chamber, where a pre-programmed degassing cycle is executed. During this cycle, the ambient pressure is gradually reduced to 5 mbar, maintained for one hour to remove any air pockets introduced during PMT placement, and finally returned to atmospheric pressure at a controlled rate.

The fully loaded trays are then transferred to dedicated heating cabinets. Curing the silicone at an elevated temperature of 40 °C under controlled humidity not only accelerates polymerization but also ensures reproducible optical and mechanical properties by eliminating seasonal variations in the production environment.

Holding Structure Stocking

In parallel with the PMT preparation, the internal holding structures are pre-assembled in a dust-free environment. Depending on the spatial clearances required for the subsequent integration steps (e.g., reserving space for the removable pneumatic scissor lifts proposed in Section 7.1.2), the structures are populated with as many of the internal electronic components as possible at this stage. These include the flasher LED units, Fanout boards and mini-mainboards.

Module Assembly

The populated holding structures are mounted in dedicated assembly fixtures that hold the pressure vessels in place. The cured PMT subassemblies (potentially incorporating fully pre-cast gel pads) are inserted into the holding structure with the pushing mechanism fully retracted. Depending on the finalized optical coupling strategy, either a conformal liquid gel layer is applied (potentially via an automated dip-coating process similar to that used for WOM production [167]) or the gel pads are prepared for dry fitting, thereby eliminating the subsequent degassing and curing stage. The populated holding structure is then lowered into the glass pressure vessel hemisphere (or vice versa) using a dedicated guiding system. Once the vacuum lid has been secured, the pushing mechanism is actuated to press the PMTs radially outward with a pre-calibrated and uniformly distributed force. After all gel pads have established sufficient contact with the inner glass surface, the loaded hemisphere undergoes a secondary degassing cycle.

If additional optical gel has been applied, the complete assembly fixture is again transferred to a heating cabinet for controlled curing. This makes the assembly station available for the next production batch while maintaining continuous contact pressure at the optical interfaces throughout the curing process.

Joining Hemispheres

Once curing has been completed, the temporary integration tooling (such as the scissor lifts and vacuum lid) is removed and any remaining internal components are installed. The two completed hemispheres are handled using aligned vacuum suction cups that permit controlled vertical motion. The central support shafts (if utilized) and the inter-hemisphere ribbon cables are connected. The internal volume is flushed with dry nitrogen to displace residual moisture. The hemispheres are then joined at the equatorial seam, immediately sealed using *Terostat* butyl tape and secured with structural exterior tape. Finally, a large vacuum chamber is lowered over the assembled module and the internal pressure is evacuated to the nominal 500 mbar.

Final Acceptance Testing and Storage

The completed Gen2-OM is packaged in a specialized transport container designed to leave the PCA accessible. The module then undergoes final acceptance testing in dedicated test setups housed within freezer containers similar to those used for the mDOM [259]. Upon successfully passing all operational and optical acceptance criteria, the final packaging is sealed and the module is transferred to the warehouse for eventual deployment to the South Pole.

B Supplementary optical Module Performance Comparison

The primary effective area evaluations for all optical modules are presented in Section 8.2. A comprehensive set of additional figures is provided in this appendix for completeness. Figures B.1 and B.2 display the total module effective area, evaluated at a wavelength of 400 nm. While these simulations were generated across the full wavelength range (from 270 nm to 700 nm), the visualizations presented here are restricted to 400 nm for conciseness.

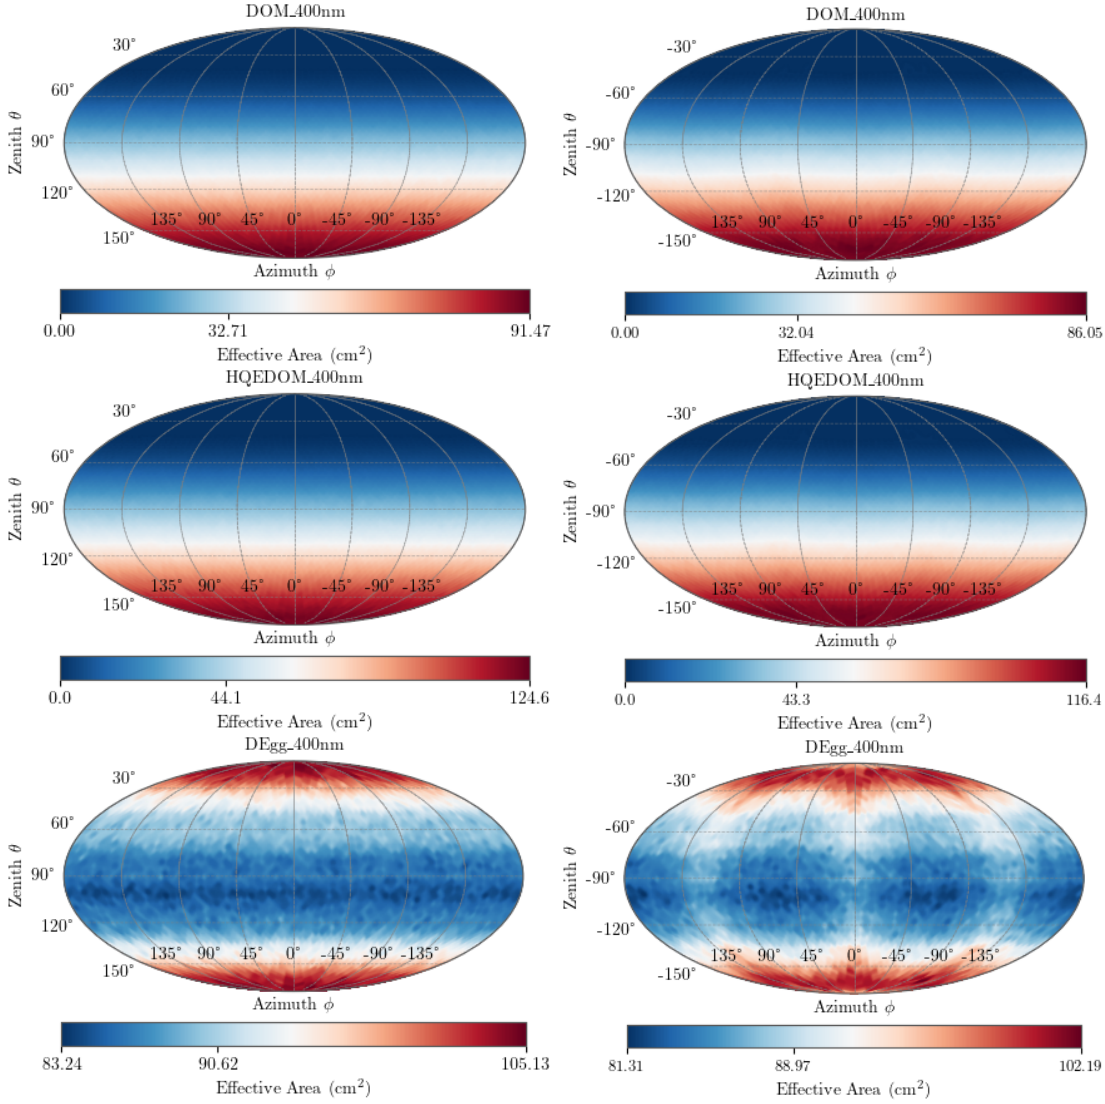


Figure B.1: Total module effective area (summing all PMTs) at 400 nm for the DOM (top), HQE-DOM (center) and D-Egg (bottom). The left column shows simulations without the mechanical harness, while the right column includes the harness. Note that the colorbar limits are adjusted individually for each plot to optimally visualize the relative angular acceptance geometry.

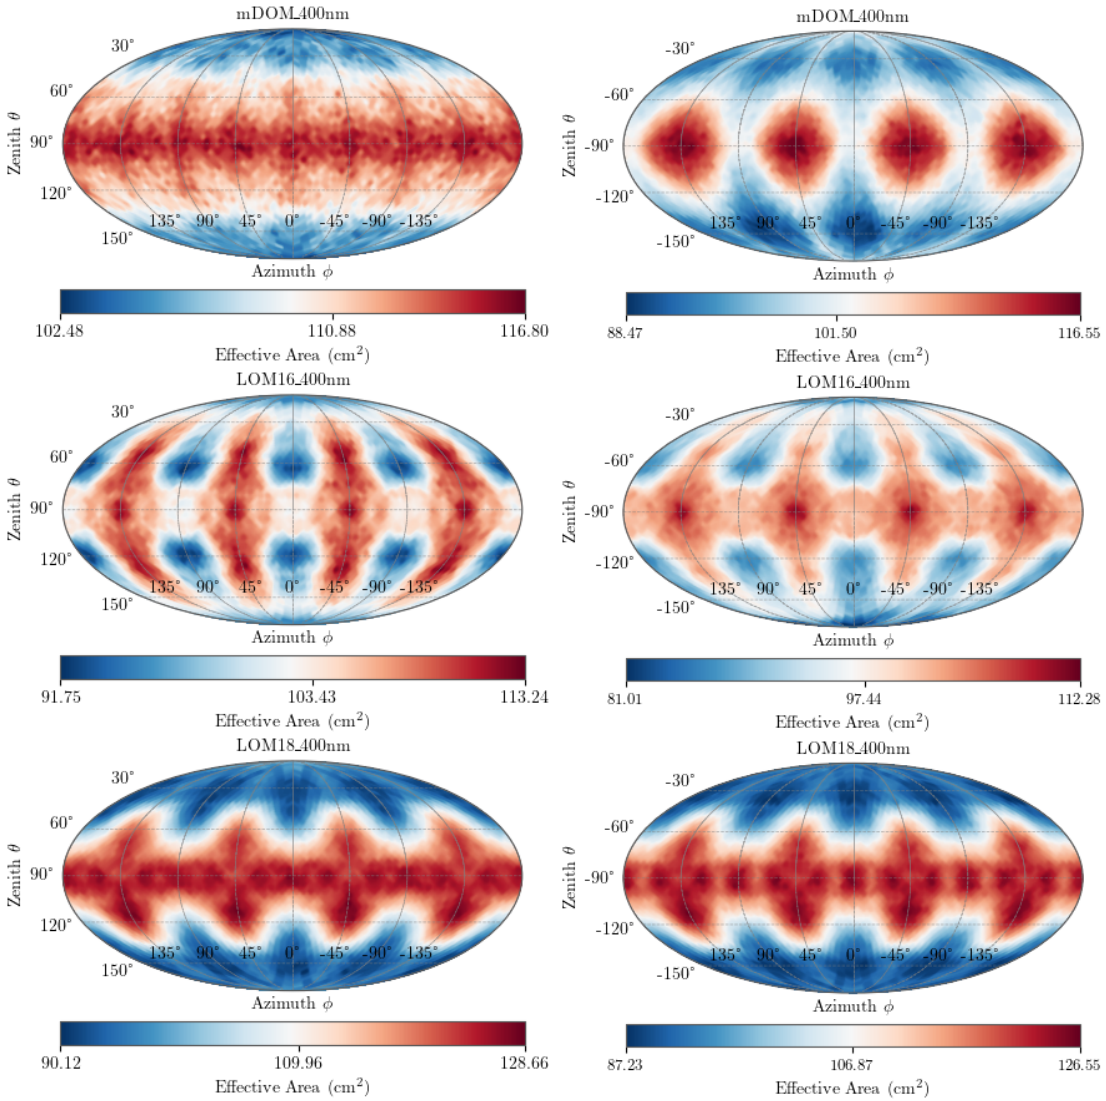


Figure B.2: Total module effective area (summing over all PMTs) at 400 nm for the mDOM (top), LOM16 (center) and LOM18 (bottom). The left column shows simulations without the mechanical harness, while the right column includes the harness. The LOM simulations assume the conservative Hamamatsu PMT quantum efficiencies. Colorbar limits are scaled individually for each plot.

While the previous figures illustrate the summed effective area of the entire module, Figure B.3 isolates the directional acceptance of individual PMTs. For single-PMT modules such as the DOM and HQE-DOM, the individual PMT response is identical to the total effective area shown previously. For the multi-PMT architectures, representative PMTs from different module hemispheres and latitudes are selected: one D-Egg PMT, the polar and equatorial PMTs of the mDOM, LOM16 and LOM18, as well as the "layer 3" vertical PMTs of the LOM18. All subplots assume a wavelength of 400 nm, include the deployment harness and apply the Hamamatsu quantum efficiencies for the LOMs.

To further examine the local PMT optics, Figure B.4 presents the effective area as a function of the local photon incidence angle relative to the PMT axis (evaluated via the dot product) across multiple wavelengths.

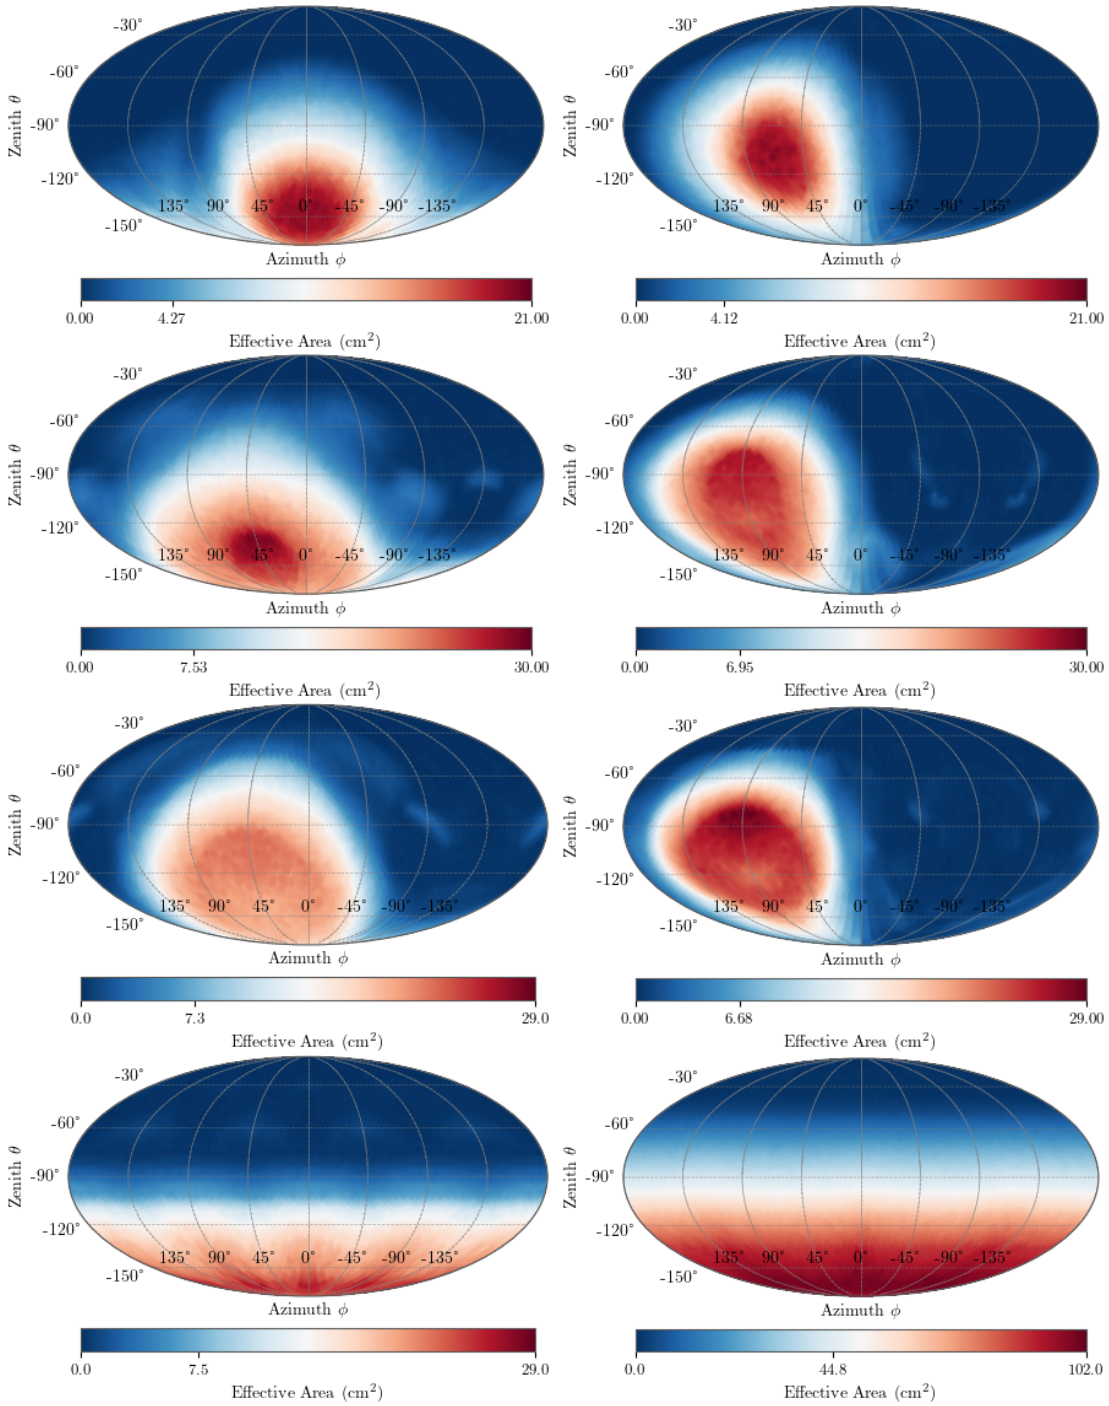


Figure B.3: Individual PMT effective area for representative PMT types within various optical modules at 400 nm. Shown are a polar and equatorial PMTs of the mDOM (first row), LOM16 (second row) and LOM18 (third row). The bottom row displays a vertical "layer 3" PMT of the LOM18 (left) and a single D-Egg PMT (right). Colorbar limits are adjusted individually for each plot.

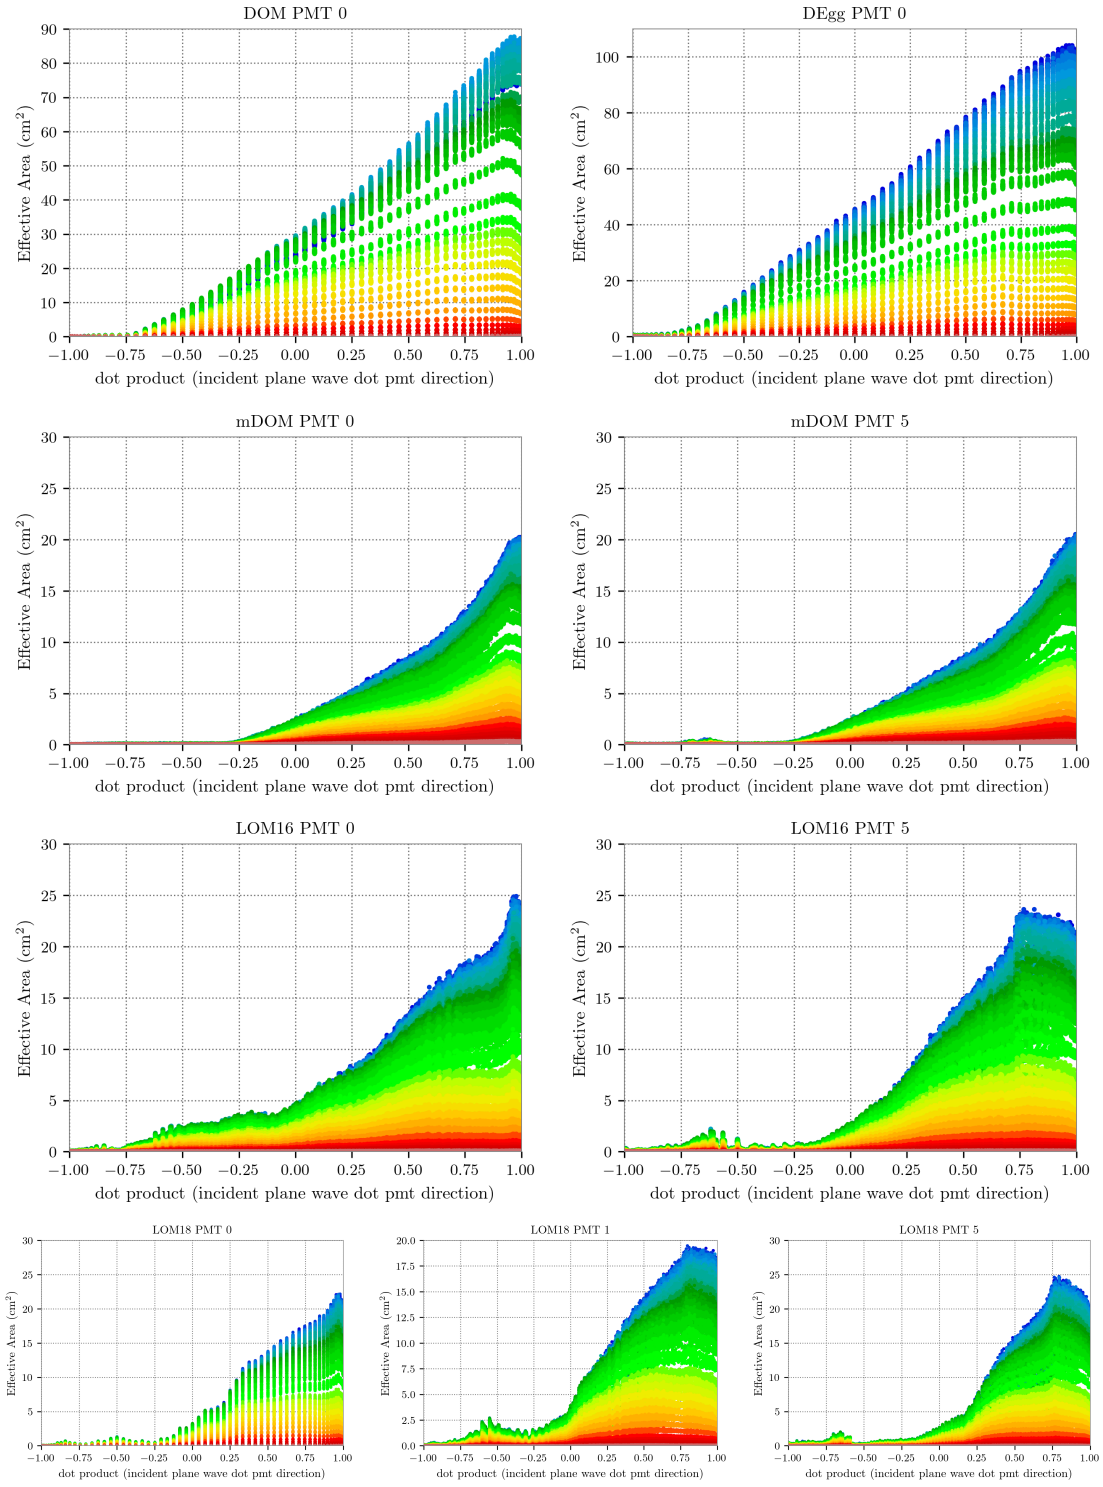


Figure B.4: Effective area as a function of the incidence angle between the PMT axis and the incident photon plane wave. Different wavelengths are color-coded (e.g., 380 nm in dark blue, 550 nm in green, 650 nm in red). The panels display the DOM and D-Egg (first row); the polar and equatorial PMTs of the mDOM (second row); the polar and equatorial PMTs of the LOM16 (third row); and the vertical "layer 3", polar and equatorial PMTs of the LOM18 (bottom row).

Finally, to quantify the physical dimensions defining the simulated geometries, Table B.1 lists the cross sections of the outer dimensions of each module's pressure vessel. Due to restrictions imposed by the smaller borehole diameter planned for IceCube-Gen2, the maximum allowable pressure vessel diameter is constrained to 12.5" $\approx$ 31.8 cm.

Table B.1: Size comparison of the optical modules. The listed values are measured from the CAD models and correspond to the projected cross sections of the outer glass pressure vessels, rounded to one decimal place. *Because the Gen2-OM design is not yet finalized, its projected cross section is estimated from the LOM height-to-diameter ratios.

Optical Module	Diameter (cm)	Height (cm)	Projected cross section (cm^2)
DOM	33.0	33.0	855.3
mDOM	35.6	40.5	1173.8
D-Egg	30.2	53.4	1386.2
LOM16	31.2	44.4	1163.1
LOM18	31.8	54.0	1467.1
Gen2-OM	31.8	48.6	$\approx$ 1300*

C PMT Manufacturer Data

Chapter 6 presented characterization results for the latest PMT versions based on three samples from each vendor. To provide broader context, this appendix shows the manufacturer quality assurance data for the complete production batches. A general overview of the measured performance parameters is shown in Figure C.1.

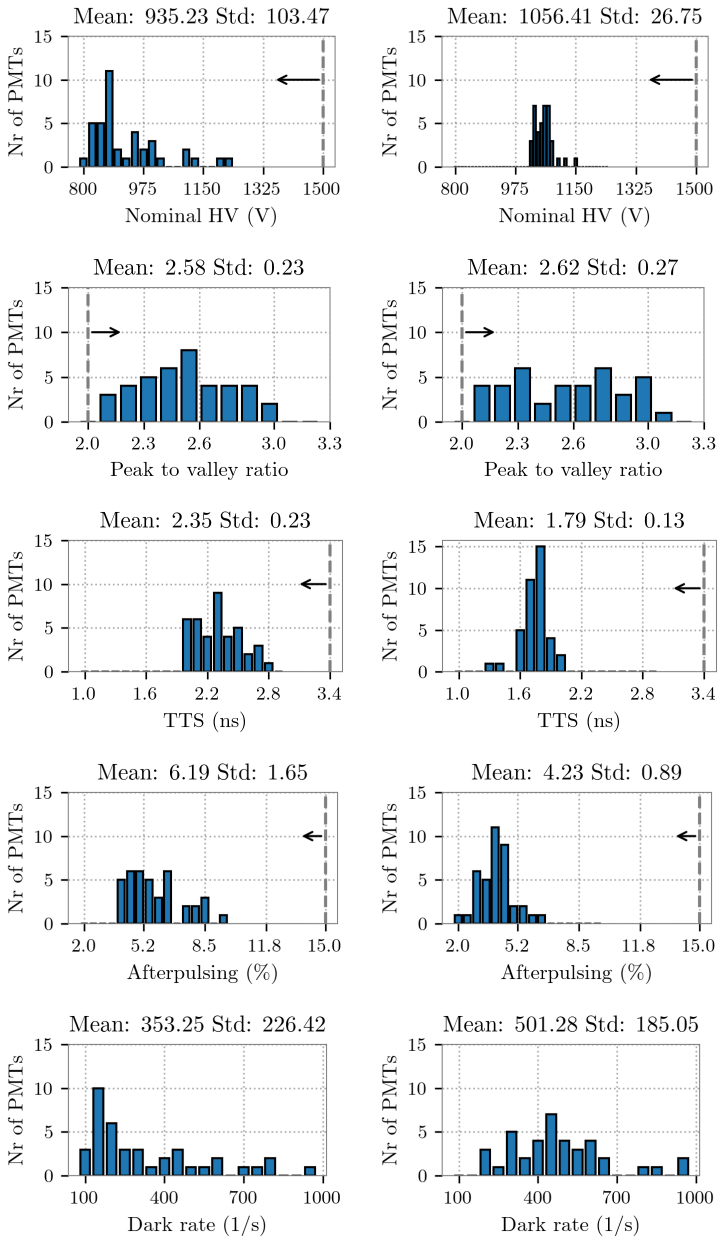


Figure C.1: Comparison of manufacturer data between NNVT **(left)** and Hamamatsu **(right)**. Acceptance requirements are indicated by gray vertical lines, while the corresponding batch mean and standard deviation are listed above each plot.

It must be emphasized that the measurements provided by NNVT and Hamamatsu cannot be compared directly with one another, nor strictly with the results of this study, owing to differences in their testing methodologies. A primary motivation for the independent characterization performed in this work was to establish a standardized and reliable baseline that evaluates both vendors under identical, experimentally relevant conditions. One notable example is the transit time spread (TTS). The TTS reported in the Hamamatsu batch data is significantly narrower than both the NNVT specification and the independent measurements obtained in this study. This difference could arise if Hamamatsu characterizes the TTS using either multiple-photoelectron illumination or a tightly collimated light source that illuminates only the center of the photocathode. Both methods effectively reduce the measured TTS compared to the full-area, single-photoelectron illumination conditions employed throughout this work.

As discussed in Section 6.3, the specific Hamamatsu PMT samples selected for detailed characterization in this work represented slight underperformers relative to their batch average, whereas the selected NNVT PMTs performed slightly above average in terms of quantum efficiency. Because Hamamatsu provided full wavelength-dependent QE spectra for every PMT, Figure C.2 compares the three selected study samples with the mean spectral response of the entire production batch.

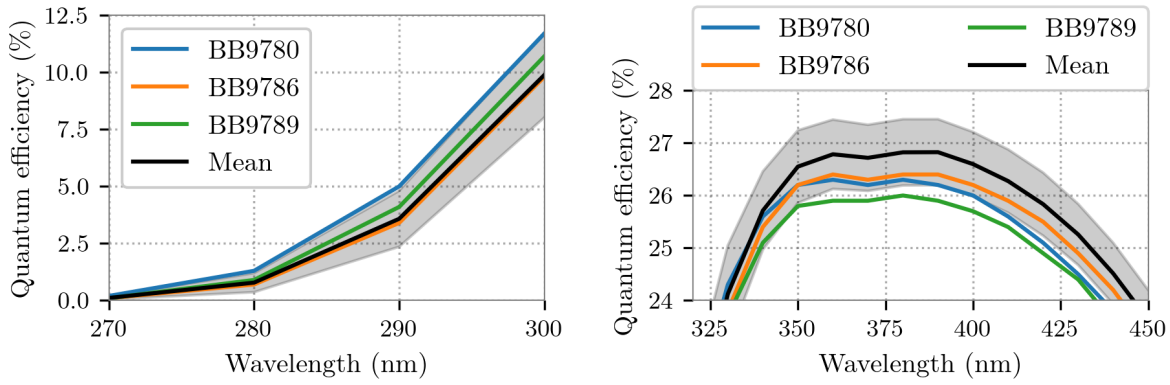


Figure C.2: Hamamatsu manufacturer data for spectral QE, highlighting the three tested study samples relative to the batch distribution.

NNVT provided single-wavelength QE measurements evaluated exclusively at 400 nm. The cumulative QE distributions at 400 nm for both vendors' production batches are shown in Figure C.3. Within their respective production distributions, the selected NNVT samples exhibit an absolute QE 1.27 % above the batch average, whereas the selected Hamamatsu PMTs are 0.63 % below their batch average at 400 nm.

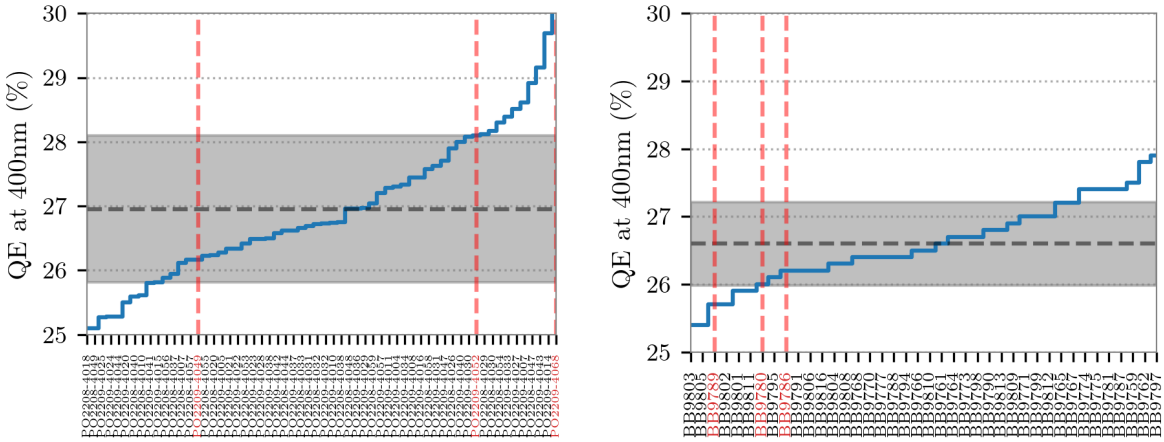


Figure C.3: Cumulative QE distributions at 400 nm from the manufacturer data for NNVT **(left)** and Hamamatsu **(right)**, highlighting the positions of the three tested samples.

Additionally, a mechanical inspection of the delivered Hamamatsu PMT batch was performed at Chiba University. The dimensions summarized in Figure C.4 confirm that all delivered PMTs satisfy the required physical envelope tolerances. A similar study conducted at WIPAC yielded consistent results for the NNVT PMTs, albeit based on pre-production units and a smaller sample size.

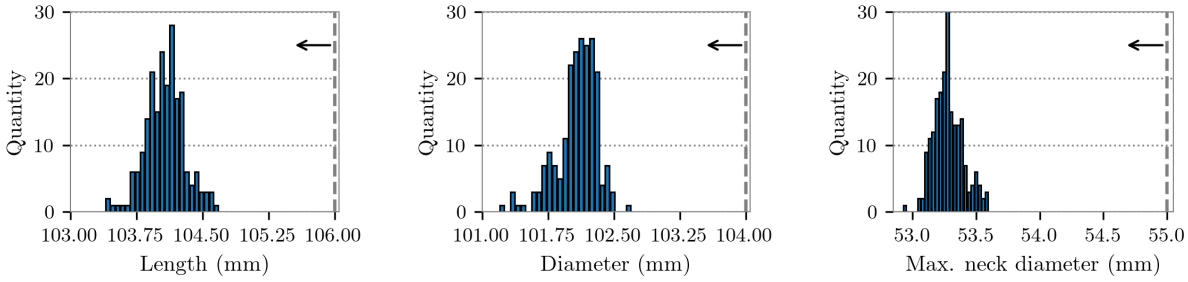


Figure C.4: Mechanical inspection results for the Hamamatsu PMT batch, confirming compliance with the specified dimensional tolerances (vertical lines).

D Supplementary Information on Module Background

Changes in Scintillation Yield

The initial determination of the Vitrovex α-yield was conducted in [91] and later extended using a fixed electron-to-alpha ratio of 1.7, as established in [183]. These earlier measurements investigated the temperature dependence of the scintillation yield by exposing small cylindrical Vitrovex glass samples to external radioactive sources and recording the resulting light yield with a standalone PMT. This experimental setup was subsequently modeled in Geant4 (Figure D.1) to account for the geometric acceptance and PMT detection efficiency. The intrinsic scintillation yield can then be determined by scaling the simulated light yield until the background-corrected PMT rates reproduce the measured values.

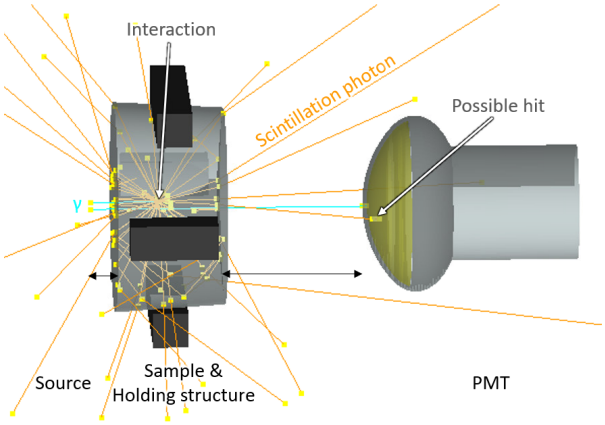


Figure D.1: Simulation setup used to determine the initial scintillation yield from small cylindrical glass samples, illustrated here for two incident γ-rays emitted by a radioactive source. Taken from [183 (p. 40)].

However, subsequent work in [73 (ch. 9)] introduced substantial revisions to the simulated PMT optical model. Combined with extensive Geant4 software updates over the intervening years, the effective detection efficiency of the simulation framework changed significantly. As a result, applying the original scintillation parameters within the updated framework produced full-module mDOM background rates that fell well below the previously validated baselines established in [73 (p. 185)]. To temporarily recover these baseline rates within the updated software (while maintaining the historical, higher isotope concentrations and narrower emission spectrum), a scaling factor of 1.35 had to be applied to the scintillation yields (Figure D.2). This recalibration was validated across multiple simulated temperatures and post-simulation artificial dead times (e.g., the 0 ns case shown).

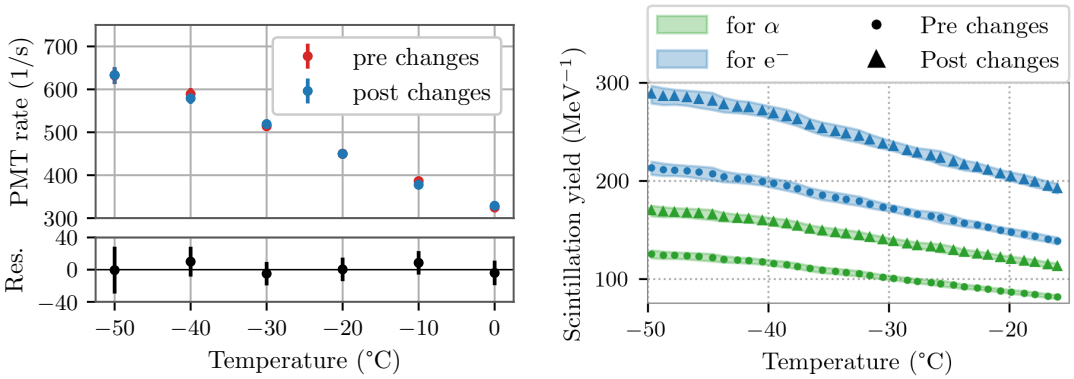


Figure D.2: (left) Average simulated mDOM PMT background rates in an ice environment before the PMT simulation changes reported in [73 (p. 185)] (red), reproduced after the software updates by applying the yield scaling factor (blue). The comparison is shown for a 0 ns artificial dead time with the intrinsic PMT dark rate included. **(right)** Comparison of the scintillation yields before and after the recalibration for each particle type. Error bands and markers indicate the systematic and statistical uncertainties associated with the yield determination.

Rather than re-simulating the historical small-cylinder setup to extract updated scintillation yields, this thesis adopts a direct calibration based on mDOM data. Fundamentally, solving the inverse problem for a small glass sample (matching simulated photons to measured PMT hits to determine the intrinsic yield) is mathematically equivalent to performing the same calibration using a fully assembled mDOM. In this case, the complete pressure vessel effectively acts as a large, self-irradiating glass sample monitored by 24 coupled PMTs.

By scaling the intrinsic yield directly against empirical mDOM target rates (derived via background subtraction as per Equation (9.5)), the simulation bypasses geometric scaling artifacts between the sample-testing environment and the final module configuration. With the integration of newly updated specific activities (Section 9.2) and a broadened emission spectrum (Section 9.3), this direct module-level parameterization of the scintillation yield represents an update to the background modeling methodology.

Spectra of Borosilicate Components

As discussed in Section 9.2.2, most characteristic FEPs fall below the decision threshold L_C, as reflected by the upper limits listed in Table 9.3. The corresponding spectra are shown in Figure D.3. The quartz sand could be characterized using all FEPs listed in Table 9.1, except for ^{40}K, whose signal could not be distinguished from the background. For the potassium nitrate sample, the prominent ^{40}K peak, exhibiting an activity approximately three orders of magnitude larger than all other FEPs across the investigated samples, increases the sample-induced background sufficiently that the spectral curves no longer overlap. This discrepancy is inconsequential, however, because a local background correction is applied prior to the FEP summation.

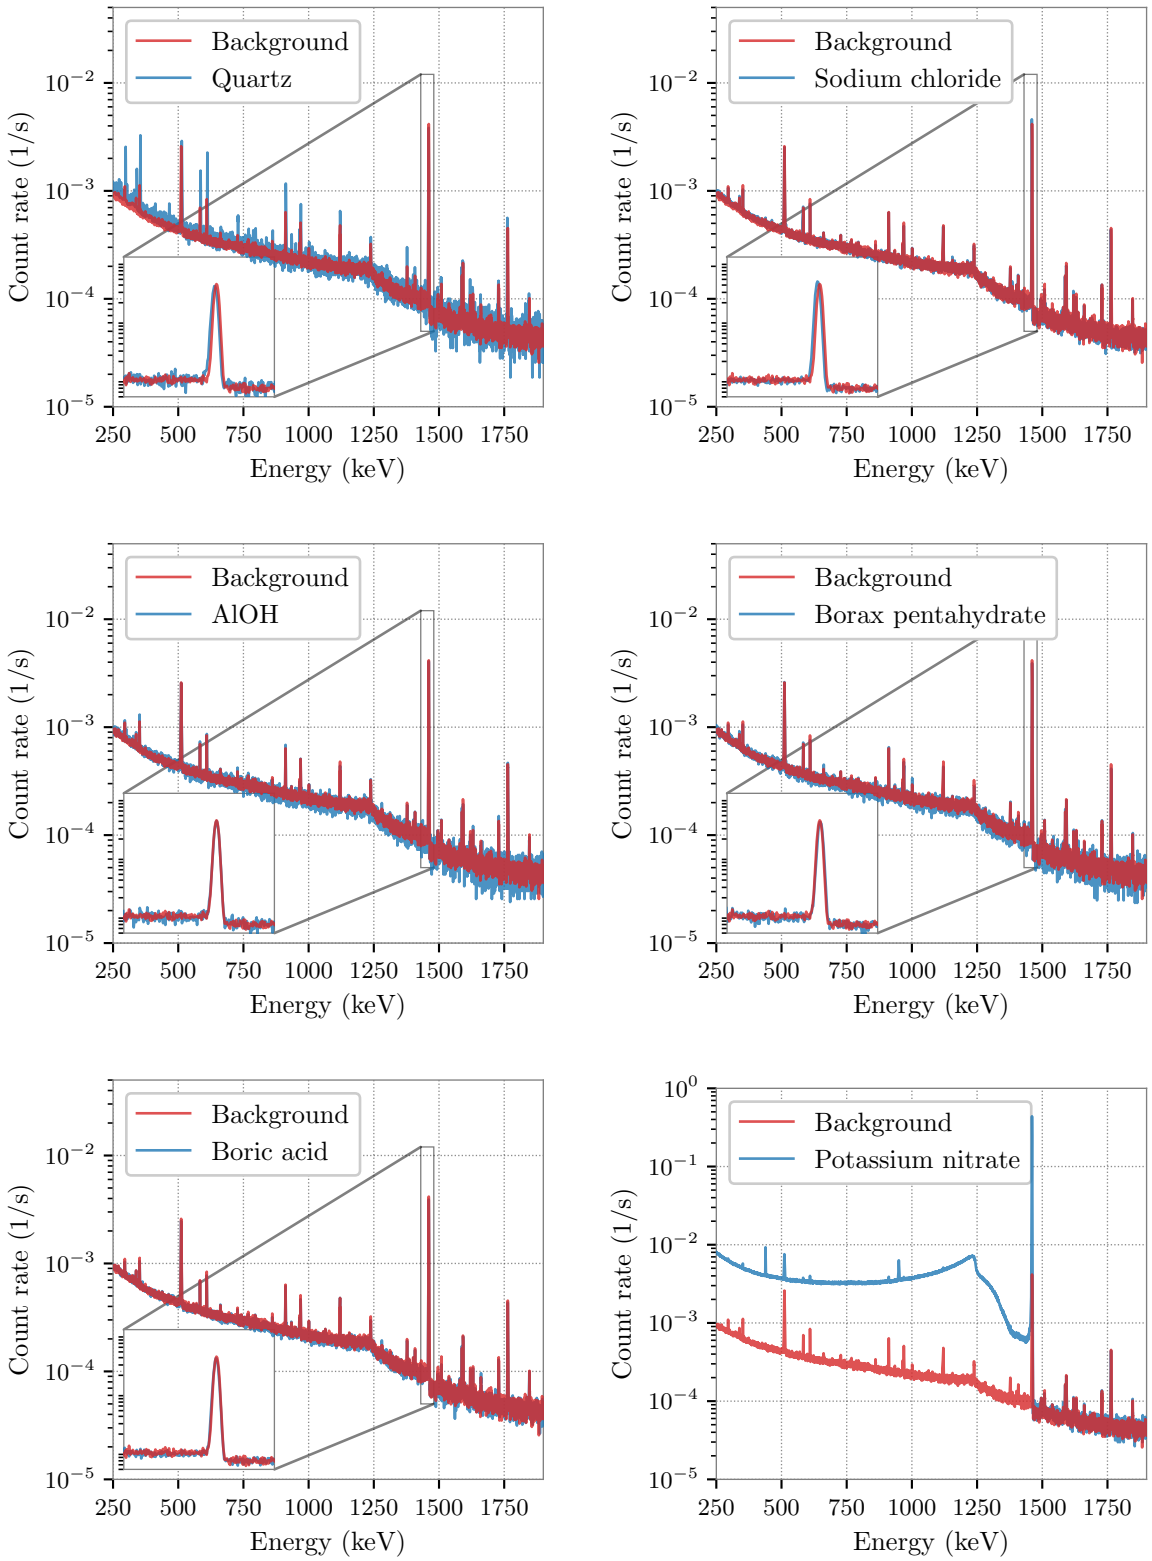


Figure D.3: Measured γ-ray spectra of the background acquired with the empty sample holder (red) and with the corresponding sample inserted (blue). The ^{40}K full-energy peak (FEP) is highlighted in each inset.

Acknowledgements

First and foremost, I would like to express my heartfelt gratitude to Alexander Kappes for his patience and invaluable guidance throughout this work. Having the opportunity to work with him in this exciting field since my early bachelor days has greatly impacted my personal development.
Similarly, I am grateful to my second advisor, Christian Weinheimer, and Jochen Heitger for evaluating this thesis.
I also wish to express my thanks to Michael Larson, who took me under his wing for the analysis part of this work. The beginning was challenging, yet he always made time to meet and patiently explain the details – time and again.

Furthermore, I wish to extend my thanks to each present and former member of our working group. It has been an absolute pleasure collaborating with you all over the years. The atmosphere was consistently welcoming and everyone was always ready to go the extra mile to offer assistance when needed – I regard you as personal friends. In particular, I want to give special recognition to Martin Unland, who, whether he likes it or not, solidified himself as somewhat of a mentor to me in the beginning – thank you for everything :)
Extending the circle, my time at the institute allowed me to meet many incredible people. From brainstorming with our engineers (Hans-Werner & Lukas) as well as colleagues in the mechanical (Georg) and electrical workshops (Bodo & Roland) to maintenance (Andreas), IT administration (Tobias) and other like-minded scientific colleagues from other working groups, I have learned way more on this journey than I initially anticipated. Whenever problems or challenging tasks arose, everyone was ready to lend a hand and come up with intricate solutions. I appreciate being able to call them my close colleagues, as the occasional burdensome task was usually handled as a joint venture.
As a member of the IceCube collaboration, I also had the privilege of meeting a diverse group of individuals from all over the world, which has broadened my worldview and left me with many cherished memories. A special shout-out goes to my counterparts in Chiba and Madison for their camaraderie (and for babysitting me in Japan) – you guys are awesome!
Closing out, special thanks to Martin, Javi, Aswathi, Yuya and Raffi who generously devoted their time to proofreading parts of this work, as well as Pia for designing the cover.

While I look forward to the future, I will miss many aspects of this chapter of my life as time moves on. The relationships I built have not only shaped my scientific journey but have also immensely enriched my personal development. I enjoyed every moment of this journey and would choose to do it all over again if given the opportunity.
Last but certainly not least, I am profoundly grateful to my wife, who has provided unwavering support throughout this demanding writing phase.

