

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

By

Jakob van Santen

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2014

Date of final oral examination: November 13, 2014

The dissertation is approved by the following members of the Final Oral Committee:

Albrecht Karle, Professor, Physics

Thomas K. Gaisser, Professor, Physics (University of Delaware)

Francis Halzen, Professor, Physics

Dan McCammon, Professor, Physics

Delia Tosi, Assistant Scientist, Wisconsin IceCube Particle Astrophysics Center

Justin Vandenbroucke, Assistant Professor, Physics

NEUTRINO INTERACTIONS IN ICECUBE ABOVE 1 TEV

Jakob van Santen

Under the supervision of Professor Albrecht Karle

At the University of Wisconsin-Madison

High-energy neutrinos are ideal cosmic messengers, produced whenever cosmic rays interact with matter or photons near their as-yet unknown acceleration sites, and carrying information about the conditions there to Earth without being deflected by magnetic fields or absorbed by intervening matter. At the same time, neutrinos produced in cosmic-ray air showers provide information about hadronic physics in kinematic regions that are difficult to probe with terrestrial accelerators.

This work presents an analysis of the energy spectrum and angular distribution of neutrinos that interacted inside the instrumented volume of the IceCube detector between May 2010 and May 2012. The analysis yielded new information about the energy spectrum of the recently discovered astrophysical neutrino flux as well as the maximum contribution to the neutrino flux from the prompt decays of charmed mesons produced in air showers.

New reconstruction, event selection, and simulation techniques were developed in the course of this work. These are described in Sections [3.4.3](#), [4.2](#), [5.2.2–5.2.4](#), and [5.4](#).

TABLE OF CONTENTS

	Page
LIST OF FIGURES	iv
LIST OF TABLES	viii
1 Introduction	1
2 The TeV neutrino landscape	7
2.1 Pion/kaon production and decay: conventional atmospheric neutrinos	7
2.1.1 Cascade equations	9
2.1.2 Measurements	14
2.2 Charmed meson production and decay: prompt atmospheric neutrinos	16
2.2.1 Limits from penetrating muon measurements	16
2.2.2 Limits from neutrino measurements	17
2.3 Astrophysical neutrinos	19
3 The IceCube Neutrino Observatory	21
3.1 Neutrino detection	21
3.1.1 Deep inelastic neutrino-nucleon scattering	21
3.1.2 Neutrino-electron scattering	23
3.1.3 Cherenkov radiation	23
3.1.4 Muon energy loss	26
3.1.5 Electron energy loss	27
3.1.6 Hadronic cascades	30
3.2 Detector	34
3.2.1 Hardware	34
3.2.2 Triggering and filtering	34
3.3 Optical properties of the South Pole ice	36
3.4 Reconstruction	37
3.4.1 First-guess algorithms	38
3.4.2 Muon track reconstruction: Pandel likelihoods	39
3.4.3 Cascade energy and angular reconstruction: Monopod	43
4 Event selection	47
4.1 Pre-selection	51
4.1.1 Online pre-selection	52
4.1.2 Offline pre-selection	53
4.2 Neutrino-level event selection	54
4.2.1 Clean-up cuts	54
4.2.2 Outer-layer veto	55

	Page
4.2.3 Inner track veto	59
4.2.4 Fiducial volume scaling	63
4.2.5 High-energy and upgoing-track event retention	65
4.3 Final sample	66
5 Analysis methods	68
5.1 Modeling the atmospheric neutrino flux	68
5.1.1 Neutrino event simulation	68
5.1.2 Conventional atmospheric neutrinos	70
5.1.3 Prompt atmospheric neutrinos	71
5.1.4 Cosmic ray composition correction	73
5.2 Vetoing atmospheric neutrinos	75
5.2.1 Modified cascade equations	77
5.2.2 Generalized atmospheric neutrino self-veto	79
5.2.3 Verification with full air shower simulation	87
5.2.4 Verification with full detector simulation	87
5.3 Modeling the astrophysical neutrino flux	91
5.4 Estimating the background from penetrating atmospheric muons	91
5.5 Observables	100
5.5.1 Deposited energy	100
5.5.2 Neutrino flavor discrimination	100
5.5.3 Zenith angle	102
5.5.4 Observable distributions for each component	106
5.6 Likelihood fitting	110
5.6.1 Forward folding	110
5.6.2 Unfolding	110
5.6.3 Energy scale correction	112
5.6.4 Confidence intervals	112
5.6.5 Sensitive energy range of the astrophysical spectrum fit	114
6 Results	117
6.1 Forward-folding likelihood fit	117
6.2 Descriptive statistics: unfolding	125
7 Conclusion and future directions	127
APPENDIX MuonGun	130
APPENDIX The 30 TeV down-going excess	144
Glossary	159

LIST OF FIGURES

Figure	Page
1.1 The all-particle cosmic ray spectrum	2
1.2 The cosmic ray spectrum above the knee	3
2.1 Atmospheric neutrino production	8
2.2 Approximate conventional neutrino and muon fluxes from pion and	13
2.3 Predictions of the prompt atmospheric neutrino flux	18
3.1 Feynman diagrams for deep inelastic neutrino-nucleon scattering.	22
3.2 Feynman diagram for neutrino-electron scattering	24
3.3 Neutrino-nucleon and antineutrino-electron scattering cross-sections	24
3.4 The Cherenkov effect	25
3.5 Refractive index of ice	26
3.6 Muon energy loss rates and secondary spectra	28
3.7 Muon ranges in ice	29
3.8 Cherenkov photon yield from electromagnetic cascades	31
3.9 Relative Cherenkov light yield of hadronic cascades	33
3.10 Absorption and effective scattering lengths of the South Pole ice	37
3.11 Detection geometry for cascades and tracks	40
3.12 Pandel time-residual PDFs	42
3.13 Amplitude and timing of detectable photons from cascades	44
4.1 Surface geometry of IceCube	48
4.2 Outline of the event selection	49
4.3 Geometry of the outer-layer veto	56

Figure	Page	
4.4	A demonstration of the outer-layer veto algorithm.	58
4.5	Illustration of the incoming-muon veto procedure	60
4.6	A demonstration of the inner downgoing track veto algorithm.	62
4.7	Muon background suppression as a function of energy	64
4.8	Fiducial volume scaling function	64
4.9	Identifying a neutrino-induced muon	65
4.10	Distribution of photon counts per event after each stage of the event selection	66
4.11	Neutrino effective area	67
5.1	Base conventional atmospheric neutrino flux parameterization	72
5.2	Base prompt atmospheric neutrino flux parameterization	74
5.3	Cosmic ray composition correction	75
5.4	Atmospheric neutrino self-veto	76
5.5	Minimum muon surface energy required to reach IceCube	80
5.6	Conventional atmospheric neutrino fluxes without accompanying muons	81
5.7	Fits to air shower lepton yields	84
5.8	Air shower lepton yields	85
5.9	Muon/neutrino energy ratio for various decays	86
5.10	Comparison of approximate self-veto passing rates with Monte Carlo calculation	88
5.11	Verification of the approximate veto probability calculation against the full detector simulation	90
5.12	Effective livetime of CORSIKA simulations	94
5.13	Muon multiplicity by selection level	95
5.14	Underground muon flux models	97
5.15	Effective livetime of CORSIKA and MuonGun simulations	99
5.16	Cascade energy resolution	101
5.17	Energy dependence of energy resolution	101
5.18	Zenith angle reconstruction error	104

Figure	Page
5.19 Zenith angle reconstruction error as a function of deposited energy	105
5.20 Zenith angle and deposited energy distributions for neutrinos and penetrating muons	107
5.21 Zenith angle and deposited energy distributions for starting track events	108
5.22 Zenith angle and deposited energy distributions for cascade events	109
5.23 Basis functions for energy spectrum unfolding	111
5.24 Energy scale correction	113
5.25 Sensitive energy ranges	116
6.1 Deposited energy spectra from the northern and southern skies	120
6.2 Zenith angle distribution	121
6.3 Verification of atmospheric neutrino veto with low-energy data	122
6.4 2D profile likelihood of the fit	123
6.5 Best-fit astrophysical flux as a function of prompt normalization	124
6.6 Unfolding the non-atmospheric excess	126
Appendix	
Figure	
A.1 Total single-muon flux	133
A.2 Total flux as a function of multiplicity	134
A.3 Radial distribution	135
A.4 Single-muon energy distribution	135
A.5 Multi-muon energy distribution	136
A.6 Distribution of zenith angles and intersection depths in CORSIKA simulation	139
A.7 Combined CORSIKA/MuonGun weighting	140
A.8 Comparison of the muon spectra generated by CORSIKA and MuonGun	141
A.9 Deficiencies in the muon bundle parameterization	143
B.1 Goodness of fit	145
B.2 Penetrating muon control sample	146

Figure	Page
B.3 Effect of variations in the self-veto probability calculation	148
B.4 Effect of systematic variations in ice properties	149
B.5 Vertex-string distance	150
B.6 Vertex-border distance	151
B.7 Time distribution of neutrino events near the excess	151
B.8 Effect of a source at the galactic center	152
B.9 Energy spectra fit with a galactic-center component	154
B.10 Zenith angle distributions fit with a galactic-center component	155
B.11 Arrival directions of neutrino events near the excess	156
B.12 Azimuth angle reconstruction error	157
B.13 Azimuth angle reconstruction error as a function of deposited energy	158

LIST OF TABLES

Table	Page
2.1 Meson decays to neutrinos	15
3.1 Electromagnetic fraction of hadronic showers	32
4.1 Number of events at each stage of the event selection	50
5.1 Parameters of $\cos \theta^*$ formula	71
5.2 Parameters of minimum muon surface energy formula	79
5.3 Parameters of the modified Elbert lepton yield formula	84
5.4 MuonGun simulations used in this analysis	98
5.5 Event type classification rates	103
6.1 Best fit parameters and number of events attributable to each component	117
Appendix	
Table	
A.1 Running time for CORSIKA and MuonGun	142
B.1 Best fit parameters for a galactic-center model	153

Chapter 1

Introduction

The Earth is constantly bombarded by high-energy atomic nuclei collectively known as the cosmic rays. Physicists have studied these particles since the discovery of ionizing radiation from space by Victor Hess in 1912 [1], and after just over 100 years of research, we know that the cosmic rays are primarily nuclei of all elements up to iron, but mostly free protons (79%) and helium nuclei (15%) [2]. Because they are charged, they can be deflected by magnetic fields, making their arrival directions at Earth isotropic to within 1 part in 10^4 [3]. From several GeV to several PeV, their energy spectrum follows a power law given by

$$\frac{dN}{dE} \sim 1.8 \times 10^4 \left(\frac{E}{1 \text{ GeV}} \right)^{-2.7} \frac{\text{nucleons}}{\text{GeV m}^2 \text{ sr s}} . \quad (1.1)$$

Above the “knee” near ~ 4 PeV, the spectrum softens from $\sim E^{-2.7}$ to $\sim E^{-3}$ before hardening again at the “ankle” near ~ 4 EeV [2]. Figure 1.1 shows the cosmic ray spectrum from the geomagnetic cutoff at a few GeV to the highest energies yet observed. Up to approximately 100 TeV the flux is large enough that it is practical to observe the particles directly with satellite- or balloon-borne detectors. Above 100 TeV, the fluxes become so small that cosmic rays can only be detected via the extensive air showers [4] of charged particles that an energetic nucleus initiates when it collides with another nucleus in the Earth’s atmosphere. Figure 1.2 shows a collection of all-particle energy spectra from various air-shower experiments. These are shown multiplied by $E^{2.6}$ to emphasize the knee and ankle.

What is still unknown is where the cosmic rays are accelerated to the high energies that we observe at Earth. While energy density arguments can be used to constrain the classes of sources that can be responsible for certain fractions of the flux (see e.g. [6] for a review), the scrambling of the directions of charged particles in magnetic fields makes it impractical to identify the sources from the arrival directions of cosmic rays at




Figure 1.1: The all-particle cosmic ray spectrum as a function of energy per particle from, reproduced from [5]. The spectrum is a nearly featureless power law over 10 orders of magnitude in energy.




Figure 1.2: The all-particle cosmic ray spectrum as a function of energy per nucleus, reproduced from [2]. With the exception of the satellite-borne Proton (Grigorov) and balloon-borne JACEE experiments, all spectra in this figure were measured indirectly through observations of extensive air showers by ground-based experiments. Such experiments can not directly observe the initiating particle, and so must infer its energy and type from the characteristics of the resulting shower.

most energies¹. In principle neutrino astronomy offers a way out: neutrinos will be produced whenever the cosmic rays interact with matter or radiation near their acceleration sites, and will travel in straight lines to the Earth without being absorbed by intervening gas or dust [6, 8–10]. However, despite years of searches with the large water and ice Cherenkov detectors Baikal NT-200, AMANDA, ANTARES, and now IceCube, built for neutrino astronomy, no sources of high-energy neutrinos have been identified yet [11]. Even though the sources are too weak to be detected individually yet, we can attempt to learn something about their properties by studying the quasi-diffuse astrophysical neutrino flux formed from the superposition of their individual neutrino outputs. Along the way, we can also study the flux of atmospheric neutrinos that form the background to searches for astrophysical neutrinos. These neutrinos are produced in air showers initiated by cosmic rays once they reach Earth², and the behavior of their energy spectrum and angular distribution can be used to constrain models of hadronic interactions at energies that are beyond the reach of fixed-target accelerator experiments in kinematical regions that are inaccessible to collider experiments.

This work presents such a study in the form of an analysis of neutrino interactions inside the instrumented volume of the IceCube detector from May 2010 to May 2012. Chapters 2 and 3 provide the background material required in order to understand the methods and results presented in Chapters 4–6.

Chapter 2 provides a review of the sources of neutrinos with energies greater than ~ 1 TeV. These are divided into two broad categories: atmospheric neutrinos, produced in cosmic-ray-induced air showers, and astrophysical neutrinos, produced in unknown, distant astronomical objects. Atmospheric neutrinos are further subdivided into a conventional component, whose energy spectrum is one power steeper than that of the cosmic rays at Earth, and is relatively well understood; and a prompt component whose energy spectrum follows that of the cosmic rays, but has yet to be conclusively observed. Above 100 TeV, the neutrino flux at Earth is dominated by a recently discovered [12, 13] astrophysical component whose energy spectrum is harder than that of the cosmic rays.

Chapter 3 provides an overview of neutrino detection in the IceCube Neutrino Observatory. Neutrinos interact rarely, but when they do, they transfer some of their energy to charged particles that move faster than the local speed of light in the ice that forms IceCube’s detection medium, inducing Cherenkov radiation. The range of these particles varies by type; electrons deposit most of their energy within a few meters, while muons can travel tens of kilometers. A small fraction of the induced Cherenkov photons are recorded by

¹ Nearby sources of protons may be identifiable via anisotropies in the directions of $\gtrsim 50$ EeV air showers. The Telescope Array experiment has recently presented evidence of a “hot spot” near the super-galactic plane in their > 57 EeV data [7]

²It would be more proper to call atmospheric neutrinos a “foreground” in the astronomical sense, since they are produced closer to the observer than astrophysical neutrinos. Here, however, “background” is simply meant as the opposite of “signal.”

IceCube’s photomultipliers, and the positions and times of photon detections are used to reconstruct the properties of the underlying event.

Chapter 4 presents the event selection developed for this work. It consists of three stages: an “online” pre-selection performed at the South Pole at the time the data were taken, an “offline” pre-selection performed later in coordination with other analyzers within the IceCube Collaboration, and a final selection specific to this work. While the pre-selections focused primarily on data reduction to satisfy resource constraints (data transmission bandwidth from the South Pole and storage space for associated Monte Carlo simulations, respectively), the final event selection was intended to remove as much of the background from penetrating atmospheric muons as possible in order to isolate a nearly pure sample of neutrino events. The final selection was based entirely on veto techniques, rejecting events where Cherenkov photons were detected at times and locations compatible with a penetrating muon entering the detector. Unlike previous veto implementations in IceCube, the rejection power of the veto used in this work is tunable, and was strengthened in inverse proportion to the energy deposited in the detector, making it possible to maintain sufficient background rejection down to a deposited energy of 1 TeV. An inverted variant of the veto technique was used to identify neutrino-induced muon events among the majority of neutrino-induced charged-particle shower events. 388 events passed the selection in 641 days of data-taking, of which approximately 92% were neutrinos.

Chapter 5 presents the analysis method applied to extract the contributions of conventional and prompt atmospheric neutrinos, penetrating atmospheric muons, and astrophysical neutrinos to the final event sample. Simulated neutrino and penetrating muon events were passed through the same selection that was applied to the data, and the surviving events weighted to a collection of assumed models to obtain predicted distributions of deposited energy, zenith angle, and presence or absence of a detectable out-going muon. The flux models were then adjusted to minimize the mismatch between the observed and predicted data distributions, as measured by a scoring function based on the Poisson distribution. While this is straightforward in principle, some new work was required to efficiently simulate the background from penetrating muons and properly model the fraction of down-going atmospheric neutrinos that are vetoed by penetrating muons produced in the same air shower.

Chapter 6 presents the results of the analysis method described in Chapter 5 to the event sample isolated in Chapter 4. The astrophysical neutrino flux was observable down to neutrino energies of 25 TeV, and the best fit flux was found to be

$$\Phi_\nu = 2.06_{-0.3}^{+0.4} \times 10^{-18} (E_\nu/10^5 \text{ GeV})^{-2.46 \pm 0.12} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1} \quad .$$

The data yielded no evidence for a neutrino flux from the prompt decays of charmed mesons in air showers, resulting in a new upper limit on the flux of prompt atmospheric neutrinos of 1.52 times the [Enberg, Reno, and Sarcevic \(ERS\)](#) prediction [14] at 90% confidence. The rate of neutrino events from the southern sky with deposited energies around 30 TeV exceeded the best-fit prediction by a noticeable margin, but could not be explained by known sources of systematic error in the neutrino acceptance of the detector. The excess is currently not statistically significant, however, and more data will be needed to investigate its origins.

Several of the methods developed in the course of this work were genuinely novel. A subset of these resulted in peer-reviewed publications, including the multidimensional B-spline histogram interpolation method used to smoothly parameterize photon transport in ice [15], the cascade vertex, direction, and energy reconstruction method based on it [16], the calculation of the fraction of atmospheric neutrinos that are vetoed by muons [17], and the result of the work itself [18]. Others, like the extension of the neutrino event simulation to treat full air showers as described in Section 5.2.4 and the targeted penetrating muon simulation scheme described in Appendix A remain unpublished, but will be quite useful tools in future iterations of this work.

Chapter 2

The TeV neutrino landscape

2.1 Pion/kaon production and decay: conventional atmospheric neutrinos

Hadronic interactions in extensive air showers produce copious numbers of pions. While π^0 decay promptly (on average, within 8×10^{-17} s) to a pair of photons, feeding the electromagnetic part of the shower, charged pions live $\sim 10^9$ times longer (2×10^{-8} s [2]). They either decay to $\mu\nu_\mu$ or, at high-energies, collide with a nucleus in the atmosphere to produce more, lower-energy pions, along with smaller numbers of other hadron species, as illustrated in Figure 2.1. This leads to a flux of atmospheric neutrinos whose energy spectrum follows that of the primary cosmic rays at low energies ($\lesssim 100$ GeV), but becomes one power steeper at high energies, where re-interaction becomes relatively more important due to relativistic time dilation. Kaons are also produced in air showers, albeit in ~ 10 times smaller numbers [20]. These too can decay to produce $\mu\nu_\mu$ (as well as a much smaller number of $e\nu_e$), and since they live roughly half as long as pions and transfer a larger fraction of their energy to a neutrino when they decay, they provide the dominant source of atmospheric neutrinos above 100 GeV despite their smaller number. The neutrinos produced in pion and kaon decays are collectively called “conventional” atmospheric neutrinos to distinguish them from the as-yet unobserved flux of “prompt” neutrinos from the decays of heavy, much shorter-lived mesons that are expected to provide the largest source of atmospheric neutrinos in the hundreds of TeV.

This section presents a brief outline of the theoretical motivation for the characteristic energy, angular, and flavor distribution of the conventional atmospheric neutrino flux before reviewing some relevant previous measurements.




Figure 2.1: An illustration of neutrino production in extensive air showers (reproduced from [19]).

2.1.1 Cascade equations

The development of an air shower in slant depth X in g/cm^2 can be described by a set of coupled differential equations for hadrons and leptons. Under some simplifying assumptions, these can be solved analytically to yield expressions for the flux of muon and muon neutrinos from the 2-body decays of charged pions and kaons. While these assumptions do not hold strictly in nature, the analytic solutions to the cascade equations provide some insight into the behavior of the energy spectra and zenith distributions of atmospheric neutrinos, and are useful for parameterizing the results of full Monte Carlo calculations as we will see in Section 5.1.2. This section provides a summary of the exhaustive treatment of the cascade equations given in Chapters 3–7 of [21].

In order to describe the spectrum of muons and neutrinos from meson decay, we first have to describe the spectrum of the parent mesons produced in the air shower. For hadron species i the differential number of particles N_i per unit energy is given by

$$\begin{aligned} \frac{dN_i(E, X)}{dX} = & - \left(\frac{1}{\lambda_i} + \frac{1}{d_i} \right) N_i(E, X) \\ & + \sum_j \int \frac{F_{ji}(E_i, E_j)}{E_i} \frac{N_j(E_j, X)}{\lambda_j} dE_j \quad . \end{aligned} \quad (2.1)$$

The first term involving the interaction length λ_i and the decay length d_i describes the loss of particles to collisions and decays. The second term described gains of particles from interactions of other hadron species. $F_{ji}(E_i, E_j)$ is the dimensionless total cross-section for a hadron of species j and energy E_j to produce an outgoing hadron of species i , defined as

$$F_{ji}(E_i, E_j) = E_i \frac{dn_i(E_i, E_j)}{dE_i} \quad , \quad (2.2)$$

where dn_i is the number of hadrons of type i produced in the energy bin $E_i \pm dE_i/2$.

This matrix of equations can be solved analytically for the fluxes of pions and kaons as a function of energy and slant depth under the following simplifying assumptions:

1. The primary cosmic ray spectrum can be described as featureless power-law flux of protons and neutrons over all energies.
2. Interaction lengths are independent of energy.

3. All differential cross-sections follow Feynman scaling, i.e. can be written in a universal form

$$F_{ji}(E_i/E_j) = F_{ji}(x_L) \quad .$$

4. The atmosphere is isothermal, i.e. its density decreases exponentially with altitude, and has no curvature.

While all of these are violated to some degree in nature, they allow for the derivation of illustrative functional forms for hadron and lepton fluxes in air showers. The first assumption allows all fluxes to be factorized into parts that depend only on energy and only on slant depth. The second and third allow the source term in Equation (2.1) to be rewritten in terms of $x_L \equiv E_i/E_j$:

$$\frac{dN_i(E, X)}{dX} = - \left(\frac{1}{\lambda_i} + \frac{1}{d_i} \right) N_i(E, X) + \sum_j \int_0^1 \frac{F_{ji}(E, E/x_L) N_j(E/x_L, X)}{\lambda_j} \frac{dx_L}{x_L^2} \quad . \quad (2.3)$$

The final assumption implies that the local density is a simple function of the slant depth X and zenith angle θ ,

$$\rho = \frac{X \cos \theta}{h_0} \quad , \quad (2.4)$$

where h_0 is the scale height of the atmosphere, approximately 6.4 km. Since the decay length in g/cm^2 is proportional to the local density, it too has a simple form in this approximation,

$$d_i = c\beta\tau_i \frac{E_i}{m_i} \rho \approx \frac{c\tau_i}{h_0 m_i c^2} E_i X \cos \theta = \frac{E_i}{\epsilon_i} X \cos \theta \quad , \quad (2.5)$$

where

$$\epsilon_i \equiv \frac{h_0 m_i c^2}{c\tau_i} \quad (2.6)$$

is called the ‘‘critical energy.’’ The ratio E/ϵ governs the relative importance of the decay term and interaction terms in Equation (2.3): if $E \gg \epsilon$ then decay can be neglected, and if $E \ll \epsilon$ then decay dominates. Table 2.1 gives critical energies for the meson species that contribute to the atmospheric neutrino flux.

In the high-energy limit, the flux of pions is

$$\Pi(E, X)_{E \gg \epsilon} = N(E, 0) \frac{Z_{N\pi}}{1 - Z_{NN}} \frac{\Lambda_\pi}{\Lambda_\pi - \Lambda_N} \left(e^{-X/\Lambda_\pi} - e^{-X/\Lambda_N} \right) \quad . \quad (2.7)$$

In the scaling limit, the pion spectrum is proportional to the nucleon spectrum at the top of the atmosphere $N(E, 0)$. The Z_{ij} are the spectrum-weighted moments of the inclusive cross-sections

$$Z_{ij} \equiv \int_0^1 (x_L)^{\gamma-1} F_{ij}(x_L) dx_L \quad , \quad (2.8)$$

where γ is the integral spectral index of the primary cosmic rays. The Λ_i are attenuation lengths defined as

$$\Lambda_i \equiv \frac{\lambda_i}{1 - Z_{ii}} \quad . \quad (2.9)$$

In the opposite limit where $E \ll \epsilon$, the pion flux acquires a zenith dependence through the decay length, and the spectrum becomes harder than the primary spectrum by one power of E as pions are lost to decay at a rate proportional to $1/E$:

$$\Pi(E, X)_{E \ll \epsilon} = N(E, 0) \frac{Z_{N\pi}}{\lambda_N} e^{-X/\Lambda_N} \frac{XE \cos \theta}{\epsilon_\pi} \quad . \quad (2.10)$$

The expression for the spectrum of kaons follows a similar form.

Given the meson fluxes we can turn to the muons and neutrinos produced in their decays. The spectrum of secondaries of species i is given by

$$\frac{dN_i(E, X)}{dX} = \sum_j Br(j \rightarrow i) \int_{E_{\min}}^{E_{\max}} \frac{dn_{ij}(E, E')}{dE} \frac{N_j(E', X)}{d_i} dE' \quad , \quad (2.11)$$

where $Br(j \rightarrow i)$ is the branching ratio of decays of particle j to final states that contain a particle i , $\frac{dn_{ij}(E, E')}{dE}$ is the spectrum of secondaries per decay. For 2-body decays $M \rightarrow \mu\nu$, in the relativistic limit, $\frac{dn_{ij}(E, E')}{dE}$ is simply a constant

$$\frac{dn}{dE_\nu} = \frac{dn}{dE_\mu} = \frac{1}{E_M(1 - r_M)} \quad , \quad (2.12)$$

where E_M is the lab-frame energy of the parent meson M and $r_M \equiv m_\mu^2/m_M^2$ is the squared ratio of the μ and parent meson masses. The lab-frame energies of the secondary muons and neutrinos are limited by

$$\begin{aligned} 0 \leq E_\nu &\leq (1 - r_M)E_M \quad \text{and} \\ r_M E_M &\leq E_\mu \leq E_M \quad . \end{aligned} \quad (2.13)$$

The relationship between the muon and neutrino energies will become important in the discussion of self-veto of atmospheric neutrinos by muons in Section 5.2. Inserting these limits and decay spectrum in Equation (2.11), the spectrum of muon neutrinos from charged pion and kaon decays is

$$\begin{aligned} \frac{dN_\nu(E, X)}{dX} = & \frac{\epsilon_\pi}{X \cos \theta (1 - r_\pi)} \int_{E_\nu/(1-r_\pi)}^\infty \frac{\Pi(E, X)}{E} \frac{dE}{E} \\ & + \frac{0.635 \epsilon_K}{X \cos \theta (1 - r_K)} \int_{E_\nu/(1-r_K)}^\infty \frac{K(E, X)}{E} \frac{dE}{E} \quad , \end{aligned} \quad (2.14)$$

where Π and K are the fluxes of charged pions and kaons, respectively, that solve Equation (2.3) (c.f. Equations 2.7 and 2.10), and we have used the branching ratios $Br(\pi^\pm \rightarrow \nu\mu) = 1$ and $Br(K^\pm \rightarrow \nu\mu) = 0.635$ [2].

The neutrino flux at depth X can be obtained by integrating Equation (2.14) from the top of the atmosphere. Taking the limit of large X and changing variables from E to $z \equiv E/E_\nu$

$$N_\nu = N_N(E_\nu) \left\{ \frac{Z_{N\pi}}{1 - r_\pi} \xi_\pi(E_\nu) I_\pi(E_\nu) + 0.635 \frac{Z_{NK}}{1 - r_K} \xi_K(E_\nu) I_K(E_\nu) \right\} \quad , \quad (2.15)$$

where $N_N(E_\nu)$ is the primary nucleon flux evaluated at E_ν , $\xi_i(E_\nu) \equiv \epsilon_i/(E_\nu \cos \theta)$, and

$$I_i(E_\nu) = \frac{\Lambda_i}{\lambda_N} \int_{1/(1-r_i)}^\infty \frac{dz}{z^{\gamma+2}} \left[\frac{1}{z + \xi_i(E_\nu)} - \frac{\Lambda_i/\Lambda_N - 1}{2z + \xi_i(E_\nu)} + \frac{(\Lambda_i/\Lambda_N - 1)^2}{3z + \xi_i(E_\nu)} - \dots \right] \quad . \quad (2.16)$$

Equation (2.16) can be evaluated in the low- and high-energy limits to yield

$$I_i(E_\nu) = \begin{cases} \frac{1}{\gamma+1} \frac{\Lambda_N}{\lambda_N} (1 - r_i)^{\gamma+1} \frac{1}{\xi_i(E_\nu)} & E_\nu \ll \epsilon_i \\ \frac{1}{\gamma+2} \frac{\Lambda_N}{\lambda_N} (1 - r_i)^{\gamma+2} \frac{\Lambda_i}{\Lambda_i - \Lambda_N} \ln \frac{\Lambda_i}{\Lambda_N} & E_\nu \gg \epsilon_i \end{cases} \quad . \quad (2.17)$$

We now have the essential features of the atmospheric ν_μ spectrum. In the limit of $E_\nu \ll \epsilon_i$ the $\xi_i(E_\nu)$ terms cancel, and the neutrino flux from each meson type is an image of the primary nucleon flux, while in the high-energy limit where $E_\nu \gg \epsilon_i$ it is steeper by one power of energy. For this reason, K^\pm provide the bulk of the atmospheric neutrino flux at TeV energies despite the fact that Z_{NK} is ~ 10 times smaller than $Z_{N\pi}$.

The low- and high-energy limits of Equation (2.15) can then be joined with an interpolation of the form

$$N_\nu(E_\nu) \approx \frac{N_N(E_\nu)}{1 - Z_{NN}} \sum_i \frac{Br(i \rightarrow \nu) A_{i\nu}}{1 + B_{i\nu} \cos \theta E_\nu / \epsilon_i} \quad , \quad (2.18)$$

where

$$A_{i\nu} \equiv Z_{Ni} \frac{(1-r_i)^\gamma}{\gamma+1} \quad \text{and} \quad (2.19)$$

$$B_{i\nu} \equiv \left(\frac{\gamma+1}{\gamma+2} \right) \left(\frac{1}{1-r_i} \right) \left(\frac{\Lambda_i - \Lambda_N}{\Lambda_i \ln(\Lambda_i/\Lambda_N)} \right) . \quad (2.20)$$

The flux of muons from the same processes can be obtained by integrating Equation (2.16) from 1 to $1/r_i$ instead of from $1/(1-r_i)$ to ∞ . The corresponding coefficients of the interpolant are

$$A_{i\mu} \equiv Z_{Ni} \frac{1-r_i^{\gamma+1}}{1-r_i} \frac{1}{\gamma+1} \quad \text{and} \quad (2.21)$$

$$B_{i\mu} \equiv \left(\frac{\gamma+1}{\gamma+2} \right) \left(\frac{1-(r_i)^{\gamma+2}}{1-(r_i)^{\gamma+2}} \right) \left(\frac{\Lambda_i - \Lambda_N}{\Lambda_i \ln(\Lambda_i/\Lambda_N)} \right) . \quad (2.22)$$

The difference in kinematics accounts for the differences in the muon-neutrino and muon spectra from the same decay processes shown in Figure 2.2. Even though the total number of meson decays contributing to each flux is the same, the muon flux at any given energy is higher, because it is derived from mesons at a lower energy (and thus higher flux) than the neutrino flux at the same energy. For the same reason, the kaon contribution is much less important in the muon flux than it is in the neutrino flux.

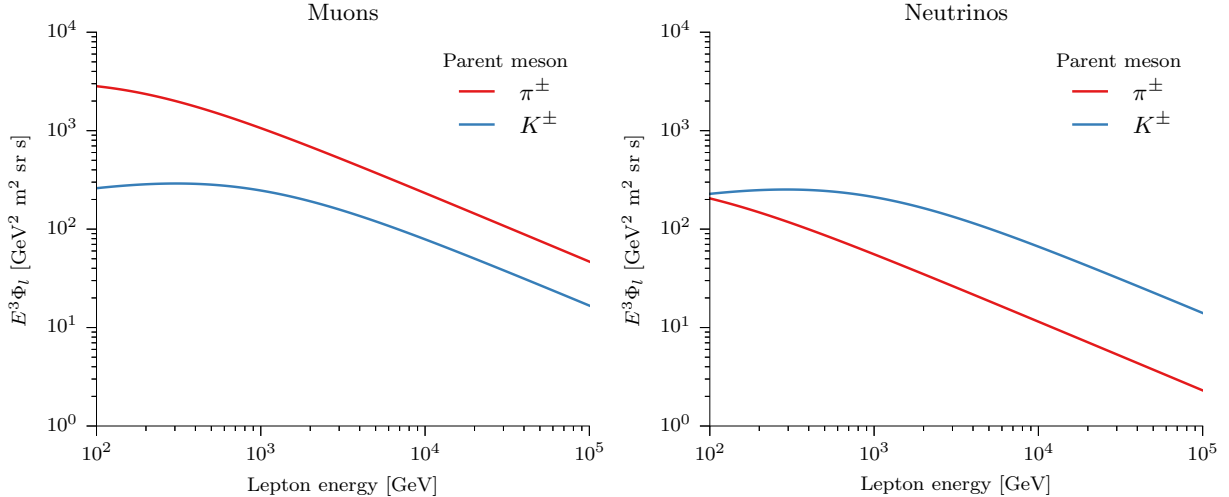


Figure 2.2: Approximate conventional atmospheric muon (left panel) and muon-neutrino (right panel) fluxes given by Equation (2.18). Even though the total number of meson decays contributing to each flux is the same, the muon flux at any given energy is higher, because it is derived from mesons at a lower energy (and thus higher flux) than the neutrino flux at the same energy. For the same reason, the kaon contribution is much less important in the muon flux than it is in the neutrino flux.

The flux of atmospheric ν_e is more difficult to treat. Because the weak interaction couples only to fermions with left-handed chirality and anti-fermions with right-handed chirality, the $\bar{\nu}_e$ in the decay $\pi^- \rightarrow e^- \bar{\nu}_e$ is always produced with right-handed helicity in the rest frame of the parent pion. Angular and linear momentum conservation require that the e^- also have right-handed helicity, but the weak interaction only couples to the left-handed chiral component of the electron wave function, which is small because of the electron's small mass, but non-zero. For the much more massive muon, this component is larger. The weak coupling suppresses the decay $\pi^- \rightarrow e^- \bar{\nu}_e$ by a factor of roughly

$$\frac{\Gamma(\pi^- \rightarrow e\nu_e)}{\Gamma(\pi^- \rightarrow \mu\nu_\mu)} \approx \left(\frac{m_e}{m_\mu}\right)^2 \left(\frac{m_\pi^2 - m_e^2}{m_\pi^2 - m_\mu^2}\right)^2 = 1.2 \times 10^{-4} \quad , \quad (2.23)$$

where m_e , m_μ , and m_π are the masses of the electron, muon, and pion respectively [22]. The suppression is even more extreme for 2-body kaon decay. As a result, the ν_e flux arises almost exclusively from 3-body decays of kaons such as $K^\pm \rightarrow \pi^0 e^\pm \nu_e$ and $K_L^0 \rightarrow \pi^\mp e^\pm \nu_e$ [2]. The chiral suppression is much weaker in 3-body decays, and the branching ratio to $e^\pm \nu_e$ is comparable to that for $\mu^\pm \nu_\mu$ as shown in Table 2.1.

While the decay spectrum $dn_{ij}(E, E')/dE$ in Equation (2.11) is a simple constant for 2-body decays, the spectrum for 3-body decays does not have a simple integrable form [23], and the constants $A_{i\nu}$ and $B_{i\nu}$ that appear in Equation (2.18) can't be evaluated analytically. Nonetheless $A_{i\nu}$ and $B_{i\nu}$ can still be determined by evaluating Equation (2.11) numerically or fitting them to the results of a full Monte Carlo calculation of the shower development as we will see in Section 5.1.2.

The decays of muons are a third source of both muon- and electron-neutrinos. Because of the extremely long mean lifetime of the muon ($2.2 \mu\text{s}$ [?]), however, muon decays provide less than less than 1% of the total atmospheric neutrino flux above 500 GeV [23], and will be ignored for the purposes of this work.

2.1.2 Measurements

Atmospheric neutrinos have a long observational history, starting with their nearly simultaneous discovery in 1965 with scintillation hodoscopes in deep mines at East Rand in South Africa [24] and Kolar Gold Fields in southern India [25], shortly after the experimental discovery of the ν_μ at the Brookhaven Alternating Gradient Synchrotron [26]. These first detectors were quite small and collected only a handful of nearly-horizontal muon events. High-statistics observations would have to wait until the 80s and 90s, when a generation of large-volume underground detectors originally built to search for proton decays (e.g. Fréjus

Table 2.1: Meson decays to final states with neutrinos [2]. The decay modes of the negatively charged mesons are the charge conjugates of the decay modes listed here. The decays of D_s^\pm are notable for their contribution to the quite small atmospheric ν_τ flux; the remaining D_s^\pm decay modes that involve ν_e in the final state are lumped together for clarity. Critical energies are taken from [21], assuming an isothermal atmosphere with a scale height of 6.4 km.

Particle	Lifetime [s]	$c\tau_0$ [cm]	Critical energy [GeV]	Final state	Branching fraction [%]
π^+	2.6×10^{-8}	780	115	$\mu^+ \nu_\mu$	99.9877
				$e^+ \nu_e$	0.012
K^+	1.2×10^{-8}	371	850	$\mu^+ \nu_\mu$	63.5
				$e^+ \nu_e$	0.0015
				$\pi^0 \mu^+ \nu_\mu$	3.4
				$\pi^0 e^+ \nu_e$	5.0
K_S^0	8.9×10^{-11}	2.68	1.2×10^5	$\pi^\pm e^\mp \nu_e$	0.07
K_L^0	5.1×10^{-8}	1154	205	$\pi^\pm \mu^\mp \nu_\mu$	27.0
				$\pi^\pm e^\mp \nu_e$	40.0
D^+	1.0×10^{-12}	0.028	4.3×10^7	$\bar{K}^0 \mu^+ \nu_\mu$	9.2
				$\bar{K}^0 e^+ \nu_e$	8.8
D^0	4.1×10^{-13}	0.013	9.2×10^7	$K^\pm \mu^\mp \nu_\mu$	3.5
				$K^\pm e^\mp \nu_e$	3.3
D_s^+	5×10^{-13}	0.015	7.5×10^7	$\tau^+ \nu_\tau$	5.5
				hadrons + $e^+ \nu_e$	9.8

[27], KamiokaNDE [28]) and magnetic monopoles (e.g. MACRO [29]) could be used to observe atmospheric neutrinos with a few GeV of energy. The zenith-angle-dependent deficit of neutrino-induced muons in Super-KamiokaNDE provided the first evidence that neutrinos oscillate, or change flavor, as they propagate [30]. While neutrino oscillation phenomena remain an active and interesting area of research today, the effects of atmospheric oscillations are negligible in the energy range of interest for this work ($\gtrsim 1$ TeV). A thorough treatment can be found in the excellent review of [31].

Because the spectrum falls off steeply with energy, the energy reach of these detectors was limited by their volume to a few tens of GeV in the case of Fréjus and MACRO, and 1 TeV in Super-K. Above 1 TeV, the only measurements come from large-volume Cherenkov detectors in naturally occurring water (e.g. ANTARES [32]) or ice (e.g. AMANDA [33, 34] and IceCube [35, 36]).

2.2 Charmed meson production and decay: prompt atmospheric neutrinos

The fluxes of atmospheric neutrinos and muons from meson decay are highly suppressed above the critical energy where interaction is more likely than decay for each meson family. This makes kaons the dominant source of atmospheric leptons above 1 TeV despite the fact that they are produced in far fewer numbers than pions. Eventually, however, the flux from kaon decay will be dominated by the flux from decays of shorter-lived mesons. The most important of these are the D mesons, which decay on picosecond timescales as shown in Table 2.1. Because of this extremely short decay time, the leptons from these decays are called “prompt” atmospheric leptons; their energy spectrum follows that of the primary cosmic rays up to energies in the 10s of PeV. Because their decay probability is nearly independent of the local density of the atmosphere, their zenith angle distribution is isotropic, like the primary cosmic rays.

This prompt flux, however, is expected to be quite small, and has yet to be conclusively observed in the energy spectra and angular distributions of either penetrating muons or neutrinos. Theoretical predictions of the flux depend strongly on assumptions about the $c\bar{c}$ production cross-section and fragmentation into hadrons, and vary by more than an order of magnitude, as illustrated in Figure 2.3.

2.2.1 Limits from penetrating muon measurements

The relevant measurements of penetrating muons come from indirect energy spectrum measurements in underground detectors. These measurements are indirect because the observed muons lose energy as they penetrate the rock above the detector, and so the inferred energy of each muon at the surface is a function

both of the energy and zenith angle measured at depth. The highest energies are typically only accessible by observing at large zenith angles, where the uncertainties in the density of the overburden and small errors in the zenith angle can lead to large errors in the inferred surface energy. The muon zenith angle distributions measured in the Baksan [37], Fréjus, [38], NUSEX [39], MACRO [40], and LVD [41, 42] experiments yielded no significant evidence for a prompt muon flux. A review of these measurements is given in [43].

2.2.2 Limits from neutrino measurements

It is in some ways easier to observe the prompt neutrino flux than it is to observe the prompt muon flux. Because neutrinos interact quite rarely, their energy spectrum can be measured directly without needing to correct for energy losses incurred while traveling to the detector. However, the small cross-section makes detections rare enough that gigaton-scale detectors are needed to collect them. Measurements of the energy spectrum of atmospheric ν_μ with AMANDA [33, 34], ANTARES [32], and partially-completed IceCube detector [36] have also yielded only upper limits on the prompt neutrino flux. In the case of [36], however, the upper limit excludes some of the more optimistic models such as [44]. These limits could in principle be improved by searching for ν_e instead of ν_μ to suppress the large background of atmospheric neutrinos from 2-body kaon decay. While it is not possible to strictly separate ν_e and ν_μ events in a sparsely-instrumented detector like IceCube, as we will see in Chapter 3, the signature of a large prompt component appears more prominently as a high-energy excess in the energy spectrum “cascade” events typical of all ν_e and a minority of ν_μ interactions than in the energy spectrum of positively-identifiable ν_μ events [45]. The complication, as already hinted at in Figure 2.3, is that if the prompt flux is not overly large, the excess flux of high-energy atmospheric neutrinos may be sub-dominant to a flux of high-energy extraterrestrial neutrinos produced in distant astrophysical objects. The expected properties of these “astrophysical” neutrinos are the subject of the next section.




Figure 2.3: A variety of theoretical predictions of the prompt atmospheric neutrino flux. The “Naumov” models are taken from [44], “Enberg” from [14], and “Martin” from [46]. Of these, only the RQPM model of [44] has been experimentally excluded [36]. The conventional fluxes of [47] are shown for reference, as well as an E^{-2} astrophysical flux with the normalization of [13]. The prompt flux consists of approximately equal parts ν_e and ν_μ , and exceeds the sub-dominant conventional ν_e flux at a much lower energy (and thus higher flux level) than the energy and flux level at which it exceeds the conventional ν_μ flux.

2.3 Astrophysical neutrinos

Like atmospheric neutrinos, high-energy astrophysical neutrinos are connected to the cosmic rays. To reach such high energies, the cosmic rays must be accelerated somewhere. Diffusive shock acceleration is the most promising mechanism for accelerating nuclei to the energies observed in the cosmic ray spectrum above the knee; this is expected to produce a flux of protons with a power-law spectrum at the source similar to $dN/dE \propto E^{-2}$ [6, 8–10]. If these nuclei encounter matter or radiation fields in the vicinity of their as-yet unknown acceleration sites, then their interactions produce high-energy charged pions, just as in nucleus-air collisions. Unlike pions produced in the relatively dense atmosphere of the Earth, they have a vanishingly small chance of interacting again, and nearly always decay in flight without losing energy. The muon produced in the pion decay will decay as well, producing 2 further neutrinos. In this way, each charged pion produces two ν_μ and one ν_e . Long-baseline oscillations transform this flavor ratio ($\nu_e : \nu_\mu : \nu_\tau$) of 1 : 2 : 0 at the source to approximately 1 : 1 : 1 at Earth [48, 49]. The spectrum of extraterrestrial neutrinos can easily deviate from this simple model, however. The spectral index in shock acceleration can vary strongly depending on the details of the acceleration environment, in some cases becoming as steep as 2.6 [50]. If the protons generating the neutrino flux interact in an environment with large enough magnetic fields that muons from pion decay tend to lose a large fraction of their energy to synchrotron radiation before decaying in flight, then the flavor ratio at earth can be closer to 1:1.8:1.8 than 1:1:1 [51].

These neutrinos serve as ideal astrophysical messengers, carrying information about conditions near their production sites to the Earth without being absorbed (like gamma rays) or being deflected by magnetic fields (like charged cosmic rays) [6, 10, 52, 53]. While the idea of extrasolar neutrino astronomy was proposed in a serious way as early as the late 1950s¹, it is unfortunately not quite a reality yet. Searches for steady point sources of neutrinos above the isotropic background of atmospheric neutrinos with IceCube [11] and ANTARES [55] as well as associations with high-energy transient events like gamma-ray bursts [56, 57] have yielded only upper limits on fluxes from individual sources of neutrinos.

Even if the sources of astrophysical neutrinos are too weak to be detected individually, their combined output creates a quasi-diffuse neutrino flux whose energy, flavor, and angular distribution is distinct from the characteristic distributions of atmospheric neutrinos. The flux level and energy spectrum of this diffuse flux can guide conjecture about the sources, and along with the existing limits on single point sources, set the sensitivity scale required to detect them individually. While hints of a characteristic high-energy excess

¹See [54] for an excellent history of the steps towards neutrino astronomy.

above the atmospheric neutrino spectrum appeared in early searches with the partially completed IceCube detector [36, 58], the first unmistakable sign came with the detection of two ~ 1 PeV neutrino-induced cascades during the first year of data-taking with the full IceCube detector [59]. A follow-up search [12] uncovered 15 further neutrino events of $\gtrsim 60$ TeV in two years of IceCube data, providing the first evidence of high-energy astrophysical neutrinos. A third year of data provided another 3 high-energy neutrino events (including one that deposited more than 2 PeV in the detector), further confirming the existence of a high-energy extraterrestrial neutrino flux [13].

As ground-breaking as it was, this follow-up search still left some questions unanswered due to its relatively high energy threshold. Since it was by design insensitive to details of the atmospheric neutrino spectrum, no conclusions could be drawn about the level of the prompt atmospheric neutrino flux, leading to some speculation in the community (most prominently in [60]) that prompt atmospheric neutrinos provided a large portion of the supposed astrophysical flux. For the same reason it could also not constrain the behavior of the observed astrophysical neutrino spectrum below 60 TeV. The work presented here is an attempt to rectify this by lowering the energy threshold of the search to 1 TeV, using novel techniques to reject nearly overwhelming backgrounds from penetrating muons and directly observe the transition from the relatively well-understood conventional atmospheric neutrino spectrum to the emergent astrophysical neutrino spectrum, and searching for evidence of a larger than expected flux of prompt atmospheric neutrinos in this transition region. Before describing the event selection developed as part of this work, we will take a brief detour to explain the technical details of the IceCube detector and the physics of neutrino interaction and detection that are necessary to understand the strategy behind the event selection and later, the data analysis presented in Chapter 5.

Chapter 3

The IceCube Neutrino Observatory

In order to understand the data analysis that will be presented in Chapter 5, we first have to understand the instrument that produced the data. This chapter presents a brief overview of the physics of high-energy neutrino interactions, the energy losses of secondary particles, the operating principle of the detector and basic data handling, and the reconstruction algorithms used to extract the properties of relativistic charged particles in the detector from the pattern of detected Cherenkov photons.

3.1 Neutrino detection

Since neutrinos interact only weakly, they are difficult to detect. Being uncharged, they do not emit light or exert forces¹ on electronic detectors. Accelerator experiments study the collisions of beams of particles with known momentum, producing both directly detectable particles and neutrinos. The charged particles, strongly-interacting particles, and photons will be registered in the detectors surrounding the interaction point, but the production of neutrinos must be inferred from missing transverse momentum. Observatory experiments like IceCube take a different approach, searching for the products of weak interactions of high-energy ($E > 1$ TeV) neutrinos with matter.


3.1.1 Deep inelastic neutrino-nucleon scattering

Deep-inelastic scattering (DIS) of neutrinos on nuclei in the target material can give rise to measurable signals. All neutrino flavors participate equally in the **charged-current (CC)** and **neutral-current (NC)** reactions, given by

¹With the exception of an extremely weak gravitational force that is neglected in the Standard Model [61].

$$\begin{aligned}\nu_l + N &\rightarrow l + X & (\text{CC}) \\ \nu_l + N &\rightarrow \nu_l + X & (\text{NC}) \quad ,\end{aligned}$$

where ν_l represents an incoming neutrino or antineutrino of a particular flavor (electron, muon, or tau), N the nucleon, l an outgoing charged antilepton or lepton of the appropriate flavor, and X the system of emerging hadrons. Figure 3.1 shows Feynman diagrams for these processes.



(a) Charged-current ν_τ scattering. The incoming neutrino is transformed into a charged lepton of the same flavor, and transfers some of its energy to the target nucleus. The diagrams for the other neutrino flavors are the same, with the neutrino and charged lepton exchanged for the appropriate flavor.

(b) Neutral-current scattering. The incoming neutrino remains a neutrino of the same flavor, but transfers some of its energy to the target nucleus.

Figure 3.1: Feynman diagrams for deep inelastic neutrino-nucleon scattering.

Due to the universality of the weak interaction, the cross-sections for these reactions only depend on the kinematics of the reaction and the momentum distribution of quarks within the nucleon. The cross-sections for an isoscalar target can be given in terms of the Bjorken scaling variables $x = Q^2/2M\nu$ and $y = \nu/E_\nu$ as [62–64]

$$\frac{d^2\sigma}{dx dy} = \frac{G_F^2 M E_\nu}{\pi} \begin{cases} 2 \left(\frac{M_W^2}{Q^2 + M_W^2} \right)^2 [xq(x, Q^2) + x\bar{q}(x, Q^2)(1-y)^2] & (\text{CC}) \\ \frac{1}{2} \left(\frac{M_Z^2}{Q^2 + M_Z^2} \right)^2 [xq^0(x, Q^2) + x\bar{q}^0(x, Q^2)(1-y)^2] & (\text{NC}) \end{cases} \quad , \quad (3.1)$$

where $Q^2 = -q^2$ is the 4-momentum transfer between the incident neutrino and outgoing lepton, $\nu = E_\nu - E_{\{l,\nu_l\}}$ is the energy loss in the target's rest frame, M is the mass of the nucleon², M_W and M_Z are the masses of the bosons that mediate the weak interaction, and $G_F = 1.16632 \times 10^{-5} \text{ GeV}^{-2}$ is the Fermi constant. The cross-sections are linear in energy up to a few TeV where the slope decreases because of the $1/Q^2$ term in the propagator. The quark distribution functions, q , \bar{q} , q^0 , and \bar{q}^0 depend on Q^2 and the parton momentum fraction x ; they must be extrapolated to low x by means of perturbative quantum chromodynamic calculations [65]. The total cross-sections, shown in Figure 3.3, can be obtained by integrating over all x and y ³.

3.1.2 Neutrino-electron scattering

For most energies of interest to IceCube, the charged-current scattering of electron antineutrinos on electrons in the target matter can be neglected. However, at a neutrino energy of 6.3 PeV, the center-of-mass energy (80.5 GeV) is large enough to produce a real W boson [62, 63]

$$\begin{aligned} \bar{\nu}_e + e &\rightarrow W^- \rightarrow \bar{\nu}_l + l \\ \bar{\nu}_e + e &\rightarrow W^- \rightarrow X \quad , \end{aligned}$$

as shown in Figure 3.2, which then decays into a leptonic system $\bar{\nu}_l + l$ or hadronic system X . At this energy, the total cross-section for the reaction is much larger than the cross-sections for neutrino-nucleon scattering.

3.1.3 Cherenkov radiation

Charged particles moving through an optical medium with a velocity $v = \beta c$ greater than the local speed of light c/n induce the medium to emit coherent radiation at the Cherenkov angle $\theta_C = \cos^{-1}(1/n\beta)$ as illustrated in Figure 3.4. The spectrum of these Cherenkov photons is given by [69]

$$\frac{dN}{dx d\lambda} = \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right) \quad , \quad (3.2)$$

²If we approximate the target as isoscalar, that is, an equal mix of protons and neutrons, then $M = \frac{M_p + M_n}{2}$

³All neutrino DIS cross-section calculations employ the formalism of [63]. CTEQ5 used the PDFs from [66]. CSS used ZEUS PDFs, while CSMS used HERAPDF 1.5.




Figure 3.2: Feynman diagram for resonant W^- production at a neutrino energy of 6.3 PeV.




Figure 3.3: Neutrino-nucleon and antineutrino-electron scattering cross-sections as a function of neutrino energy from [67] based on data from [64]. From bottom to top at low energy, the cross-sections are for $\bar{\nu}$ NC, ν NC, $\bar{\nu}$ CC, $\bar{\nu}$ total, ν CC, and ν total. The resonance peaked at 6.3 PeV is the antineutrino-electron resonance.




Figure 3.4: Illustration of the Cherenkov effect by Huygens' construction [68]. When a charged particle travels at $\beta = 0.5$, the spherical wave fronts are only slightly concentrated along the direction of travel. At $\beta = 1$ with $n = 1.33$, however, the wave fronts interfere constructively to produce a cone of light at an angle of 41° to the direction of the charged particle.

where λ is the wavelength, $\alpha \approx 1/137$ is the fine-structure constant. Integrated between 300 and 600 nm (the sensitivity window of the IceCube PMT) with the wavelength-dependent index of refraction shown in Figure 3.5 this yields 32565 photons per meter of relativistic charged-particle track. The total number of Cherenkov photons induced as a result of the energy deposition in the ice is a function of the total track length of relativistic charged particles in the event.




Figure 3.5: Refractive index of ice as a function of wavelength. The data points are from the compilation of [70], and the line is a B-spline interpolation of those points. The grey band indicates the range of sensitivity of the IceCube DOM.

3.1.4 Muon energy loss

TeV muons lose energy via four main processes: ionization, bremsstrahlung, e^+e^- pair production, and inelastic photonuclear interactions. The ionization loss rate increases only logarithmically with the muon energy [71], while the rate of losses to bremsstrahlung, pair-production, and photonuclear interactions is approximately proportional to energy. The average energy-loss rate is usually approximated as

$$-\left\langle \frac{dE_\mu}{dx} \right\rangle = a + bE_\mu \quad . \quad (3.3)$$

While the coefficients a and b themselves depend weakly on the muon energy, useful constant approximations for ice between 20 and 10^{11} GeV are $a = 0.246 \text{ GeV m}^{-1}$ and $b = 4.31 \times 10^{-3} \text{ m}^{-1}$ [72].

Figure 3.6 shows the energy loss rate of muons in water. As the energy of the muon increases, radiative processes begin to dominate the energy loss rate. Unlike ionization losses, these can be highly stochastic, causing the muon to lose a significant fraction of its energy in a single interaction. Nonetheless, the average fractional energy loss rate of a muon is quite small: a multi-TeV muon typically loses less than 1% of its energy transiting each meter of ice. This means that multi-TeV muons can penetrate large distances through ice, as shown in Figure 3.7.

3.1.5 Electron energy loss

Electrons lose energy in matter through the same processes as muons. Since they are much less massive than muons, however, the average fractional energy transfer is much larger, stopping electrons much more quickly. The critical energy E_c at which an electron's bremsstrahlung losses dominate over ionization is 79 MeV in water [74], much lower than the corresponding value for muons (1 TeV [74]). Above this energy, the mean energy loss rate is given approximately by


$$-\left\langle \frac{dE_{e^-}}{dx} \right\rangle = E_{e^-}/X_0 \quad , \quad (3.4)$$

where X_0 is the radiation length (0.39 m in ice [74]). On average, an electron loses all but $1/e$ of its initial energy each radiation length. The photons emitted in bremsstrahlung go on to produce electron-positron pairs in the electric fields of surrounding atoms [2], which themselves undergo bremsstrahlung, producing even more photons. This electromagnetic cascade continues until all of the initial energy has been distributed to electrons and positrons below the critical energy. As a consequence, the total length of charged-particle tracks in the cascade is proportional to the energy of the initial electron. The constant of proportionality derived from Monte Carlo simulation of electromagnetic cascades in ice is 5.21 GeV/m [75–78]. The longitudinal distribution of the cascade follows a gamma distribution

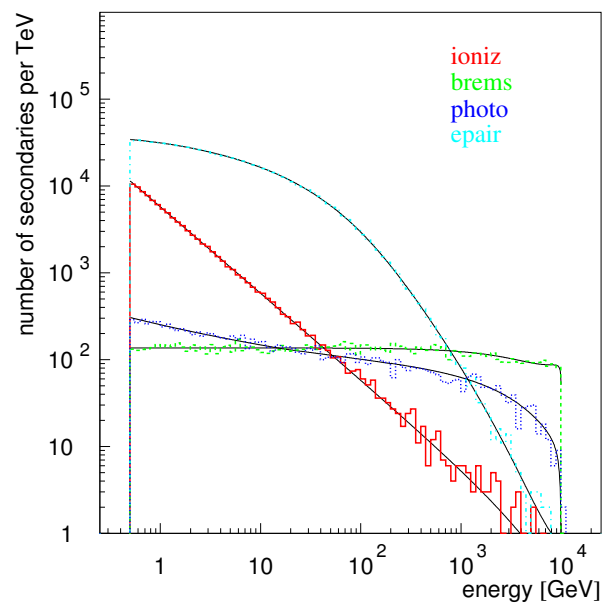
$$\frac{dE}{dt} = E_0 b \frac{bt^{a-1}e^{-bt}}{\Gamma(a)} \quad , \quad (3.5)$$

where $t \equiv x/X_0$ is the displacement in units of the electromagnetic radiation length,

$$a = 2.03 + 0.604 \log_{10}(E_0/\text{GeV}), \text{ and } b = 0.633$$




(a) Mean energy loss rates as a function of muon energy. Above a few TeV (the energy range of interest for IceCube), radiative processes are responsible for most of the energy loss.

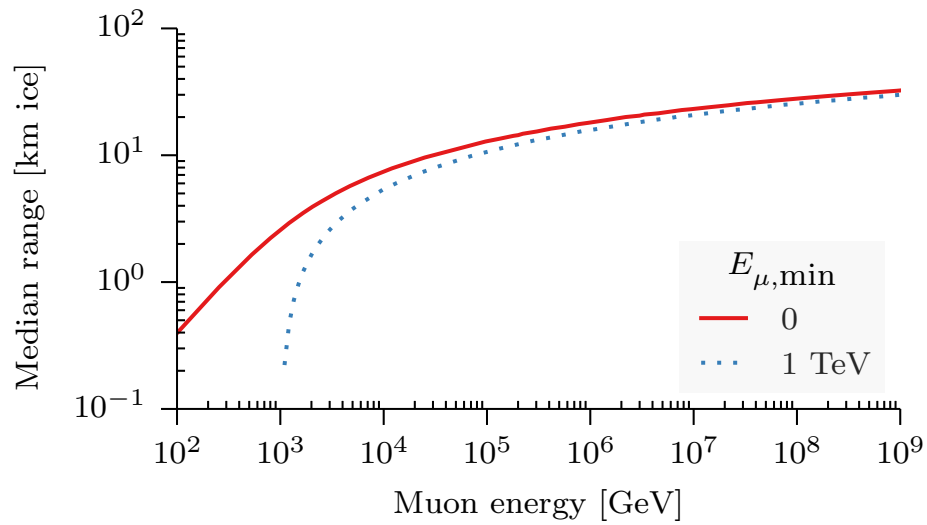


(b) Secondary spectra for 10 TeV muons in standard rock. Pair production accounts for most of the total energy loss, while bremsstrahlung is the dominant process in catastrophic losses, where the muon loses more than half its energy in a single interaction.

Figure 3.6: Muon energy loss rates and secondary spectra due to ionization (ioniz), bremsstrahlung (brems), pair production (epair), and photonuclear interactions (photo), taken from [72].



(a) Fraction of muons reaching a given distance for various initial muon energies. Solid lines show the fraction that merely survive, while the dotted lines show the fraction that survive with more than 1 TeV of kinetic energy.



(b) Median range of muons in ice as a function of initial energy. The solid line shows the median range, while the dotted line shows range where 50% of muons still have more than 1 TeV of kinetic energy.

Figure 3.7: Muon ranges in ice, calculated using PROPOSAL [73].

[75]. The longitudinal distribution and linear scaling of the total light output are illustrated in Figure 3.8. Since the charged-particle tracks that make up the cascade are nearly collinear with the initiating electron, the total Cherenkov photon emission of the cascade is sharply peaked at an angle of 41° from the initial electron direction.

3.1.6 Hadronic cascades

High-energy hadrons also lose energy rapidly in ice, though the resulting shower is slightly more elongated than a purely electromagnetic one because the nuclear interaction length in ice (0.91 m [74]) is longer than the electromagnetic radiation length. Because of the variety of particle types that may be produced in hadron-nucleus interactions, the development of hadronic showers is more complicated than that of electromagnetic showers. The total length of relativistic charged-particle tracks is no longer strictly proportional to the energy of the initiating hadron, because some energy can be lost to neutrons. Nonetheless, every hadronic shower has an intrinsic electromagnetic component due to prompt $\pi^0 \rightarrow \gamma\gamma$ decays that remove energy from the hadronic sector to feed a purely electromagnetic sub-cascade [80]. The average total Cherenkov light yield F of a hadronic cascade relative to an electromagnetic cascade with the same total energy can be treated with a purely phenomenological parameterization [76]

$$F = F_{\text{EM}} + (1 - F_{\text{EM}})f_0 \quad , \quad (3.6)$$

where F_{EM} is the fraction of the initial hadron energy that goes into the electromagnetic sector, and f_0 is the relative Cherenkov light yield of the purely hadronic part of the cascade. The increase of the electromagnetic fraction with energy can be described adequately by

$$F_{\text{EM}} = 1 - \left(\frac{E}{E_0}\right)^{-m} \quad , \quad (3.7)$$

where E is the total energy of the hadronic system and the pivot energy E_0 and power law index m are phenomenological parameters fitted to simulations. Since the visible fraction F depends on the stochastic production of π^0 , it is subject to large fluctuations. The distribution of F can be adequately described by a Gaussian distribution with a mean F and a standard deviation σ whose energy dependence is parameterized as

$$\sigma = F\sigma_0 \ln(E)^{-\gamma} \quad , \quad (3.8)$$




Figure 3.8: Cherenkov photon yield from electromagnetic cascades of various energies derived from GEANT simulation [79]. The upper panel shows the longitudinal profile, and the lower panel shows the total number of induced photons per unit of energy deposition. The constant of proportionality is approximately 169664 photons/GeV between 300 and 600 nm. Since each photon in this wavelength range carries less than 1 eV of energy on average, the energy loss to detected Cherenkov photons is less than 0.01%.

where E is again the total energy of the hadronic system and σ_0 and γ are model parameters. The above model parameters were fit to GEANT4 [79] simulations of hadronic jets from neutrino-nucleon scatterings generated with PYTHIA [81] in [77], yielding the parameters shown in Table 3.1 and plotted in Figure 3.9. These fits are valid for $E \gtrsim 500$ GeV.

E_0	m	f_0	σ_0	γ
0.399	0.130	0.467	0.379	01.160

Table 3.1: Parameters of Equations 3.6, 3.7, and 3.8 describing the Cherenkov light yield of hadronic cascades relative to purely electromagnetic cascades from [77].

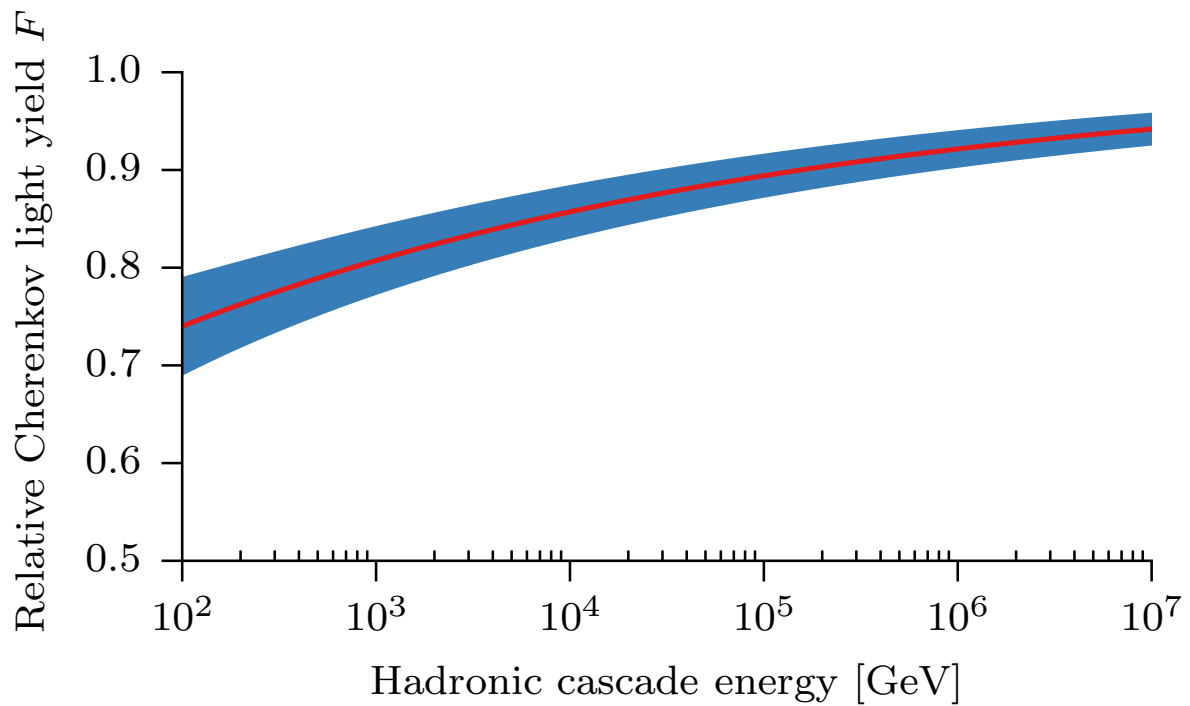


Figure 3.9: Relative Cherenkov light yield (visible fraction) of hadronic cascades relative to electromagnetic cascades with the same total energy as parameterized in [77]. The central line shows the mean scaling factor and the band its standard deviation.

3.2 Detector

3.2.1 Hardware

The IceCube detector [82, 83] consists of 5160 **Digital Optical Modules (DOMs)** buried in the ice at the South Pole, instrumenting a total volume of approximately 1 km^3 . The **DOMs** are attached to cables that provide power and communication with the data acquisition system on the surface of the glacier. Each of these “strings” hosts 60 **DOMs**; 78 of the strings are spaced 125 m apart on a hexagonal grid with **DOMs** placed every 17 m from 1450 to 2450 m below the surface, while the remaining 8 strings form the DeepCore in-fill array [83]. These in-fill strings are 30–60 m from the nearest string with 50 **DOMs** placed every 7 m between 2100 and 2450 m below the surface, where the ice is most transparent, and 10 **DOMs** placed every 10 m between 1750 and 1850 m below the surface. The data in this work were taken with the nearly-complete 79-string detector configuration from May 2010 to May 2011 and the first year of the complete 86-string detector from May 2011 to May 2012.

Each **DOM** consists of a 25 cm diameter **photomultiplier tube (PMT)** [84], power supply, and digitization electronics housed in a borosilicate glass pressure sphere. The **PMT** signal is digitized and stored for transmission to the surface whenever the **PMT** output current exceeds $1/4$ of the mean peak current of the pulse amplified from a single **photoelectron (PE)**. If a neighboring or next-to-neighboring **DOM** on the same string also triggers within $1 \mu\text{s}$ (local coincidence) the readout extends for $6.4 \mu\text{s}$, otherwise the readout only includes a 75 ns window around the peak current in the first 400 ns after the local trigger. The digitized waveforms are transmitted to the surface, where they are assembled into events by the data acquisition software pDAQ.

3.2.2 Triggering and filtering

Within the **data acquisition (DAQ)** software, a variety of trigger algorithms are applied to the incoming stream of digitized waveforms in order to recognize significant activity over the steady background of thermal noise and sub-threshold muons. The data that will be presented in Chapter 6 satisfied the simplest of these, the **simple multiplicity trigger with multiplicity 8 (SMT-8)**. The **simple multiplicity trigger (SMT)** slides a 5-microsecond-long window through the time-ordered stream of digitized waveforms. The trigger condition is satisfied whenever the window contains 8 or more local-coincidence hit records (**hard local coincidence (HLC)** hits). The **readout window** of the trigger consists of the time during which the condition is satisfied,

plus a 4-microsecond window before and a 6-microsecond window after. The **global trigger** window is the logical OR of all readout windows of all firing triggers. The hit records that fall within the global trigger window are assembled into an event that is sent to a high-level trigger system known as **Processing and Filtering (PnF)**.

The **PnF** [85, Sec. 3.4.4] system is responsible for deciding which triggered events are interesting enough to be sent over a satellite link to the IceCube Data Warehouse in Madison, WI. Since the bandwidth allocation is limited to 80 GB/day, only $\sim 10\%$ of triggering events can be transmitted⁴. The system consists of a cluster of 300 filtering clients and 1 control server. The control server accepts events from pDAQ, distributes them to filtering clients, and re-combines the processed events in time order before forwarding the events selected for transmission to the data transmission system. Each filter client is responsible for deciding whether an individual event should be transmitted.

This “**online**” filtering, so called because it happens on site and nearly in real time, is implemented in **IceTray**, a modular processing framework developed specifically for IceCube, but also used in **ANTARES** [86]. In this framework each global trigger is represented as a collection of key-value pairs called a “**frame**.” Each filtering client receives a frame containing the digitized waveforms collected during the **readout window** of the **global trigger**, and passes it through a chain of processing **modules** to extract the arrival times of individual photons and photon bunches at each DOM by deconvolving the characteristic single-**PE** pulse shape from the digitized waveforms [16] (see also [87, Section 3.4.1]). The resulting times and photon counts (**pulses**) are used to reconstruct the vertices, directions, and energies [16, 88] of the relativistic charged particles that induced the detected Cherenkov photons. The results of each processing stage are stored in the **frame** before it is passed to the next **module** in the chain, so that the last **modules** in the chain have access to the results of all previous **modules**. In the **PnF** system, the last **modules** in the chain are the **online filters** that implement the decision as to whether or not an event is interesting enough to transmit by satellite. Each detection channel has its own filter; the global filter decision is the logical OR of the individual filters. The event selection described in Chapter 4 used only events selected by the cascade filter.

⁴Starting in the 2012 data-taking season, most events were transmitted in a more compact format designed by the author of this work called SuperDST. Using this format, which encodes the leading-edge times and amplitudes of pulses extracted from the digitized waveforms, approximately 1/3 of all triggering events can be transmitted over the satellite in the same bandwidth allocation.

3.3 Optical properties of the South Pole ice

Once Cherenkov photons have been induced by the passage of relativistic charged particles through the ice, they must propagate through the glacier before they can be detected by IceCube’s photomultipliers. The vast majority of these will be absorbed before they can be detected, and the remainder will change direction in multiple small-angle scatters before they reach a photomultiplier. Therefore, an understanding of the optical properties of the medium is crucial for accurate directional and energy reconstruction.

IceCube is instrumented between vertical depths of 1450 and 2450 m below the surface of the South Pole ice sheet. The ice at these depths is formed from compressed snow that fell more than 25,000 years ago [89]. The air pockets and dust grains trapped in the snow remained as it was compressed into ice, and the strength of scattering and absorption of light at visible wavelengths is determined primarily by the concentration of these impurities [90]. At depths greater than ~ 1400 m the ice has been under pressure for long enough that the air bubbles have been compressed and integrated into the crystal structure of the ice. These air-ice clathrates have nearly the same index of refraction as the surrounding pure ice, and so do not cause visible-wavelength photons to scatter as they pass through [90]. Instead, the degree of optical scattering is determined primarily by the concentration of wind-deposited dust. The dust concentration is a tracer of global climate and is similar for ice layers deposited at the same time [89]. Since these ice layers are nearly, but not entirely horizontal [91], it is convenient to describe the scattering and absorption coefficients as a function of vertical depth.

Since no ice cores have been drilled to the relevant depths near the South Pole to date, the optical properties have been measured in situ using artificial pulsed light sources (LEDs) onboard the optical modules of AMANDA [90] and IceCube [91]. This measurement is complicated by the difficulty of modeling photon transport in the South Pole ice. In the limit where the scattering length is much longer than the absorption length, all photons travel in straight lines until they are absorbed. In the opposite limit where the absorption length is much longer than the scattering length, photon transport becomes pure diffusion. Both of these scenarios can be treated analytically. In the intermediate regime where the scattering and absorption lengths are similar, photon transport can only be described numerically with Monte Carlo simulations as described in [91]. In order to determine the parameters of the photon transport model, the light from on-board LED flashes were recorded with IceCube, and the same flashes simulated [92] with a particular model of photon transport. The simulations were repeated with different parameters to maximize agreement between the

experimental data and the simulations. Figure 3.10 shows the absorption and scattering coefficients as a function of depth determined in [91].

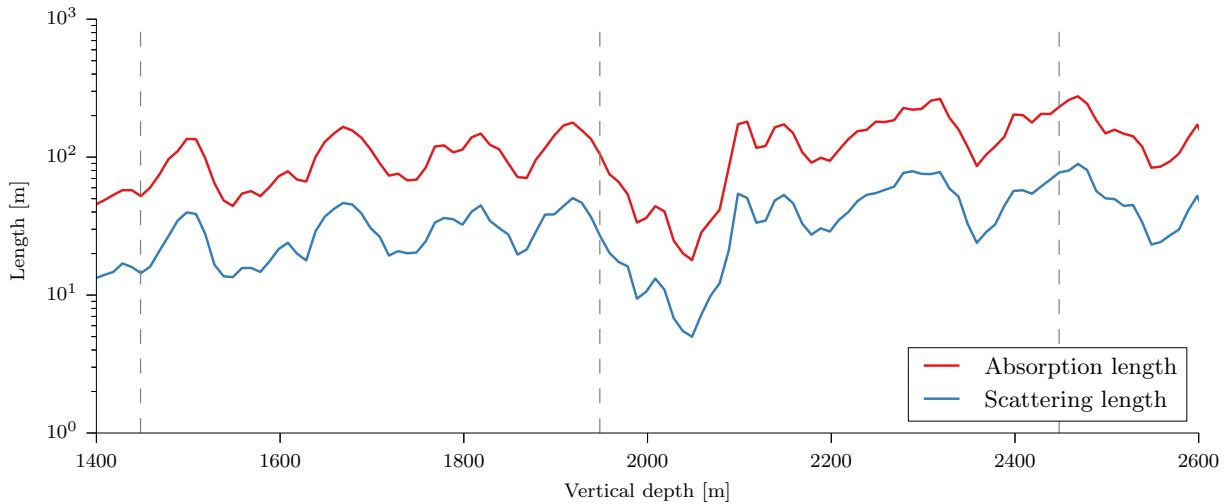


Figure 3.10: Absorption and effective scattering lengths as a function of vertical depth from [91]. The vertical dashed lines show the top, center, and bottom of the IceCube instrumented volume.

3.4 Reconstruction

Though the **pulses** recorded by the **DOMs** are the fundamental data in IceCube, they are not terribly interesting by themselves. When sifting through triggered events for signs of neutrino interactions and later, when analyzing the properties of those events, we are far more interested in the relativistic charged particles that induced the Cherenkov photons that were then detected in the **DOMs**. To get at these we have to reconstruct them from the observed pattern of **pulse** detection times and amplitudes. The reconstruction algorithms we use to calculate observables for each event are generally ordered in terms of computational complexity. The first group, called first-guess algorithms, are unaware of the physics of Cherenkov photon emission or transport in the South Pole ice, but do have closed-form solutions, and so are quite quick to calculate. The second group use parameterization of the approximate **pulse** time distributions to form a likelihood function that must be maximized numerically to find the best-fit parameters (vertex, time, and direction) of the underlying relativistic charged particles, starting from the first guess from the previous stage. The final stage of complexity is to replace the analytic approximations of the photon detection probabilities

with templates created with a full numerical simulation of photon propagation [93], described with a tensor-product B-spline surface [15]. This section describes these reconstruction algorithms to the level of detail required to understand the event selection that will be described in Chapter 4 and the observables that will be used in the analysis described in Chapter 5.

3.4.1 First-guess algorithms

3.4.1.1 Line Fit

The **line fit** is a first-guess reconstruction algorithm that treats the hit pattern as a plane wave moving through the detector. In this approximation the time of the first pulse detected in DOM i , t_i , is related to its position, $\vec{x}_{\text{DOM},i}$, by

$$\vec{x}_{\text{DOM},i} = \vec{x}_0 + t_i \cdot \vec{v} \quad ,$$

where \vec{x}_0 and \vec{v} are some position and velocity vector respectively. This is a linear problem that is solved by

$$\vec{x}_0 = \langle \vec{x}_{\text{DOM},i} \rangle - \vec{v}_{\text{LineFit}} \langle t_i \rangle$$

and

$$\vec{v}_{\text{LineFit}} = \frac{\langle t_i \cdot \vec{x}_{\text{DOM},i} \rangle - \langle \vec{x}_{\text{DOM},i} \rangle \langle t_i \rangle}{\langle t_i^2 \rangle - \langle t_i \rangle^2} \quad , \quad (3.9)$$

where $\langle x \rangle$ is the arithmetic mean over all pulses in the event [88]. In the context of cascade-like events we are only interested in the magnitude $v_{\text{LineFit}} \equiv |\vec{v}_{\text{LineFit}}|$ or **line fit speed**. For ideal, spherically symmetric cascade events v_{LineFit} will be close to 0, while for extended, minimum-ionizing tracks it will be close to c .

3.4.1.2 Tensor of Inertia

Tensor of inertia is a first-guess directional reconstruction algorithm that treats the hit pattern as a rigid body with “masses” given by the total photoelectron charge collected by each DOM. The pivot point of this imaginary body is the **center of gravity**

$$\vec{x}_{\text{CoG}} \equiv \frac{\sum_{\text{DOMs } i} q_i \cdot \vec{x}_i}{\sum_{\text{DOMs } i} q_i} \quad , \quad (3.10)$$

where q_i is the total photoelectron charge collected by each DOM. The elements of the **tensor of inertia** are given by

$$I^{kl} \equiv \frac{\sum_{\text{DOMs } i} q_i [\delta^{kl} (\vec{x}_i)^2 - \vec{x}_i^k \cdot \vec{x}_i^l]}{\sum_{\text{DOMs } i} q_i} . \quad (3.11)$$

This tensor has 3 eigenvalues corresponding to its 3 main axes. The smallest of these corresponds to the longest axis of the hit pattern. The **tensor-of-inertia eigenvalue ratio**

$$q_{\text{ToI}} \equiv \frac{\min_{k=0,1,2} e_k}{\sum_k e_k} , \quad (3.12)$$

where k runs over the three spatial dimensions, is a measure of the sphericity of the hit pattern. If the hit pattern is nearly spherical then $q_{\text{ToI}} \approx 1/3$, whereas if all hits fall on a single line it is 0.

3.4.2 Muon track reconstruction: Pandel likelihoods


The first-guess algorithms presented above do not account for the geometry of Cherenkov photon emission or optical scattering in the ice between emission and detection. The geometry is simple enough to take into account, as shown in Figure 3.11. For a point-like emitter like a cascade, the shortest-time path is a straight line from vertex to receiver at the local speed of light. Given the vertex time t_{vertex} , the earliest possible time that photons can be detected is

$$t_{\text{geo,point}} = t_{\text{vertex}} + \frac{n}{c} r , \quad (3.13)$$


where $r = |\vec{r}| = |\vec{x}_{\text{DOM}} - \vec{x}_{\text{vertex}}|$ is the distance between the vertex and the receiving DOM, n is the group index of refraction, and c is the speed of light in vacuum. For an extended emitter like a muon the position of the emission point changes with time. Given a reference vertex and time, the shortest-time path from vertex to receiver is along the muon track at the vacuum speed of light and then along the Cherenkov angle to the receiver. The earliest possible detection time in this case is

$$t_{\text{geo,track}} = t_{\text{vertex}} + \frac{\vec{r} \cdot \hat{n}_{\text{track}} + \rho \tan \theta_C}{c} , \quad (3.14)$$

where \vec{r} is the vector from the point on the muon track to the receiving DOM, \hat{n}_{track} is a unit vector along the direction of the track, θ_C is the Cherenkov angle, and $\rho = |\vec{r} - (\vec{r} \cdot \hat{n})\hat{n}|$ is impact parameter of the track with respect to the DOM. In the absence of scattering, all photons would be detected at exactly t_{geo} .



(a) In cascade events, all Cherenkov photons are induced near the vertex. The shortest time from source to receiver follows a straight path at the local speed of light c/n , shown as a solid blue line. A photon that takes this path with have a time residual of $\Delta t = 0$. Scattering, illustrated by the dotted line, lengthens the path the photon must take, introducing a time delay.



(b) In track events, Cherenkov photons are induced near the penetrating lepton as it travels at the vacuum speed of light c . Since this is larger than the local speed of light c/n , the shortest time from vertex to receiver is along the lepton path at c and then along the Cherenkov angle to the receiver at c/n . Photons induced at other points along the track or that take different paths from the same point always arrive later.

Figure 3.11: Detection geometry for cascades and tracks

Scattered photons, however, must cover more distance at the local speed of light, reaching the receivers after t_{geo} .

While there is no exact closed-form expression for the distribution of time delays in a heterogenous medium like the South Pole ice, there is a family of analytic approximations to the true time delay distribution that are collectively called “**Pandel functions**” [94]. These define a probability distribution in time

$$\frac{dP(\Delta t|r)}{dt} \quad ,$$

where r is the distance of closest approach of the emitter and receiver, and Δt is the **time residual**

$$\Delta t \equiv t - t_{\text{geo}} \quad , \quad (3.15)$$

where t is the detection time. Figure 3.12 shows these approximate **time residual** distributions evaluated for a variety of source-receiver distances [95, 96].

These **time residual** distributions depend directly only on the source-receiver distance, but acquire a dependence on the vertex, time, and direction of the assumed track through the definition of Δt . The most likely vertex, time, and direction of the track⁵ can be found by maximizing the likelihood function

$$L(x_{\text{vertex}}, t_{\text{vertex}}, \hat{n}_{\text{track}}) = \prod_{\text{pulses } i} \frac{dP(\Delta t_i | x_{\text{vertex}}, t_{\text{vertex}}, \hat{n}_{\text{track}})}{dt} \quad . \quad (3.16)$$

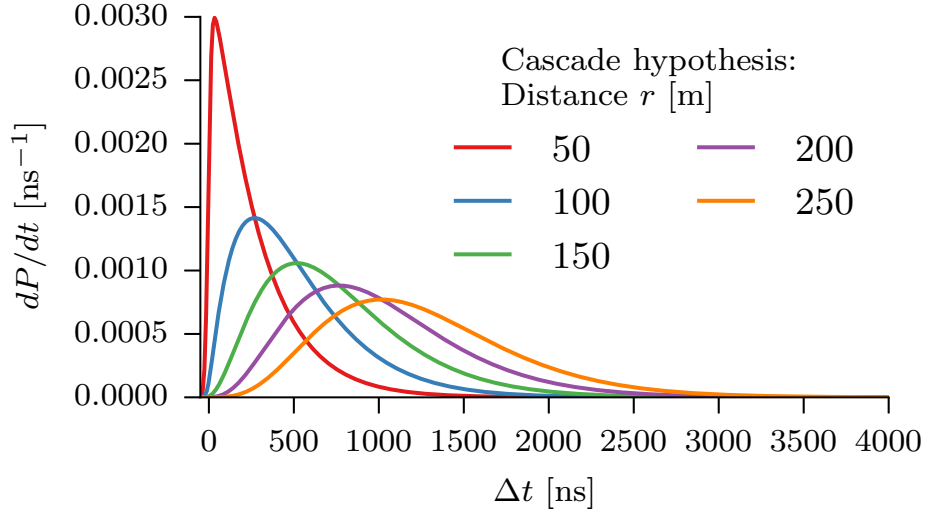
It is more convenient, however, to minimize the negative logarithm of the likelihood,

$$-\ln L = \sum_{\text{pulses } i} \ln \left(\frac{dP(\Delta t_i | x_{\text{vertex}}, t_{\text{vertex}}, \hat{n}_{\text{track}})}{dt} \right) \quad . \quad (3.17)$$


This formulation is correct as long as the sum runs over all detected pulses. Since this is true by construction when only one photon is detected by each DOM, this is called the “**single-photoelectron (SPE)**” likelihood function to distinguish it from an alternate formulation called “**multi-photoelectron (MPE)**” that uses the time residual distribution of the first of N photons instead [88]. Nonetheless, both are typically calculated using the time of the first recorded pulse on each DOM only.

The minimization of $-\ln L$ with respect to the track or cascade parameters yields a best-fit vertex, time, and in the case of tracks, a best-fit direction. The minimum value of Equation (3.17) is usually called simply

⁵The case for cascades is similar, except that the direction of the cascade does not affect the geometric time.



(a) Point cascade hypothesis (cf. Figure 3.11a).



(b) Infinite track hypothesis Point cascade hypothesis (cf. Figure 3.11b).

Figure 3.12: Pandel time-residual PDFs convolved with a 15-ns Gaussian to account for finite time resolution for cascades [95] and tracks [96]. The value of 15 ns was originally intended to account for the timing resolution of the **AMANDA** electronics, and is much smaller than the inherent time resolution of the **DOM** electronics (~ 4 ns [82]). The larger value is still used as a hedge against known mismatches between the real **time residual** distribution and the analytic approximation from the Pandel distribution.

the “log-likelihood.” This can be used as a quality parameter: small values indicate that the fit is generally good, while large values indicate that it is poor. Because the overall scale depends on the number of terms in the sum in Equation (3.17), however, it is better to define a **reduced log-likelihood** $-\ln L/n_{\text{dof}}$ where n_{dof} is the number of terms in the sum in Equation (3.17) minus the number of parameters allowed to vary in the minimization.

3.4.3 Cascade energy and angular reconstruction: Monopod


To reconstruct the energy and direction of cascades, we have to become more clever. As noted in Sections 3.1.4–3.1.6, electromagnetic energy losses in ice lead to showers of electrons, positrons, and photons. At energies above a few GeV, the total length of all charged-particle tracks in the shower, and thus the total number of induced Cherenkov photons, is proportional to the total energy in the shower. Hadronic showers behave similarly, albeit with a lower overall Cherenkov photon yield because some of the shower energy is lost to neutral particles and binding energies of hadrons. We use the light yield of a simulated 1 GeV electromagnetic cascade as the fundamental template for energy reconstruction, and infer the deposited energy by scaling this template up and down to match the total charge collected by each **DOM**.

Figure 3.13a shows the mean number of **PE** per GeV of electromagnetic-deposited energy that would be collected in **DOMs** at various distances and orientations with respect to a horizontal cascade at a depth of 1850 m ($z = 100$ m in the IceCube coordinate system), obtained from Photonics [93] simulation parameterized with a 5-dimensional tensor-product B-spline surface [15]. Using this template Λ , we can write the Poisson likelihood of observing k photons from an energy loss E as [16]

$$\begin{aligned}
 \mathcal{L} &= \frac{\lambda^k}{k!} \cdot e^{-\lambda} \\
 \lambda &\rightarrow E\Lambda \\
 &= \frac{(E\Lambda)^k}{k!} \cdot e^{-E\Lambda} \\
 \ln \mathcal{L} &= k \ln(E\Lambda) - E\Lambda - \ln(k!).
 \end{aligned} \tag{3.18}$$

Maximizing this with respect to energy, and adding the contributions from all **DOMs**:

$$\begin{aligned}
 0 = \frac{\partial \sum \ln \mathcal{L}}{\partial E} &= \sum_{\text{DOMs } j} (k_j \Lambda_j / E \Lambda_j - \Lambda_j) \\
 &= \sum k_j / E - \sum \Lambda_j \\
 \therefore E &= \sum k_j / \sum \Lambda_j.
 \end{aligned} \tag{3.19}$$



(a) Total observed light level as a function of radial distance from the source and observation angle with respect to the source direction. While scattering in the ice washes out the peak at the Cherenkov angle, the direction of the source remains visible as an asymmetry even at large distances.

(b) Normalized time distribution of detected photons at different distances for two observation angles. Photons detected at the Cherenkov angle have generally experienced the least scattering, and so are detected earlier and more closely bunched in time than those detected at other angles. In the approximation shown in Figure 3.12a, the time distribution does not depend on observation angle.

Figure 3.13: Distribution of detectable photons obtained from a Monte Carlo simulation [93] of a horizontal, 1 GeV electromagnetic cascade in the upper part of the IceCube detector. Both the number and time distribution of photons depend on the direction of the cascade, here oriented in the direction of observation angle 0. The distributions shown are made from spline tables (see [15]).

The generalization allowing additional contributions (e.g. **PMT** noise) is to replace the substitution $\lambda = E\Lambda$ in Equation 3.18 by $\lambda = E\Lambda + \rho$, where ρ is the expected number of noise photons. The likelihood (3.18) then becomes:

$$\ln \mathcal{L} = k \ln (E\Lambda + \rho) - (E\Lambda + \rho) - \ln (k!). \quad (3.20)$$

Maximizing with respect to E , as in Equation 3.19:

$$\begin{aligned} 0 &= \sum (k_j \Lambda_j / (E\Lambda_j + \rho_j) - \Lambda_j) \\ \sum \Lambda_j &= \sum k_j \Lambda_j / (E\Lambda_j + \rho_j). \end{aligned} \quad (3.21)$$

Unlike Equation 3.19, this does not have a closed form solution for E since Λ no longer cancels in the first term and E can therefore not be factored out. Solutions can, however, be easily obtained using gradient-descent numerical minimization algorithms.

In reality the number of **PE** ejected from the photocathode k is reconstructed from the height of amplified **PMT** pulses consisting of $\sim 10^7$ electrons for every initial **PE** [84], and so is not discretized in units of 1 **PE**. Since the $\sim 30\%$ fluctuations in amplitude of single-**PE PMT** pulses [84] are small compared to the Poisson fluctuations in the initial number of **PE**, it is sufficient to replace the factor $k!$ in Equation 3.20 with $\Gamma(k+1)$ [16, 97].

The likelihood function (3.20) depends on the direction and position of the cascade through the light yield templates Λ_i . We can exploit this to reconstruct the vertex and direction of the cascade by numerically maximizing Equation (3.20) with respect to the vertex and direction of the cascade, maximizing the likelihood over the energy dimension by solving Equation (3.21) at every step.

This is most effective when the time distribution of detected **PE** in each **DOM** is considered in addition to their total number. The time distribution of detectable photons depends on the relative position and orientation of source and receiver, as illustrated in Figure 3.13b. The majority of photons are emitted within a few degrees of the Cherenkov angle with respect to the cascade axis, and photons detected near this angle have the smallest average **time residual**. The **time residual** distributions of photons detected at other angles are wider, as photons have to follow scattered, longer paths to reach the **DOMs**. We reconstruct the vertex, direction, and time of the cascade by dividing the series of detected **PE** into time bins and replacing the total expected light yield Λ_i with the expectation in each time bin.

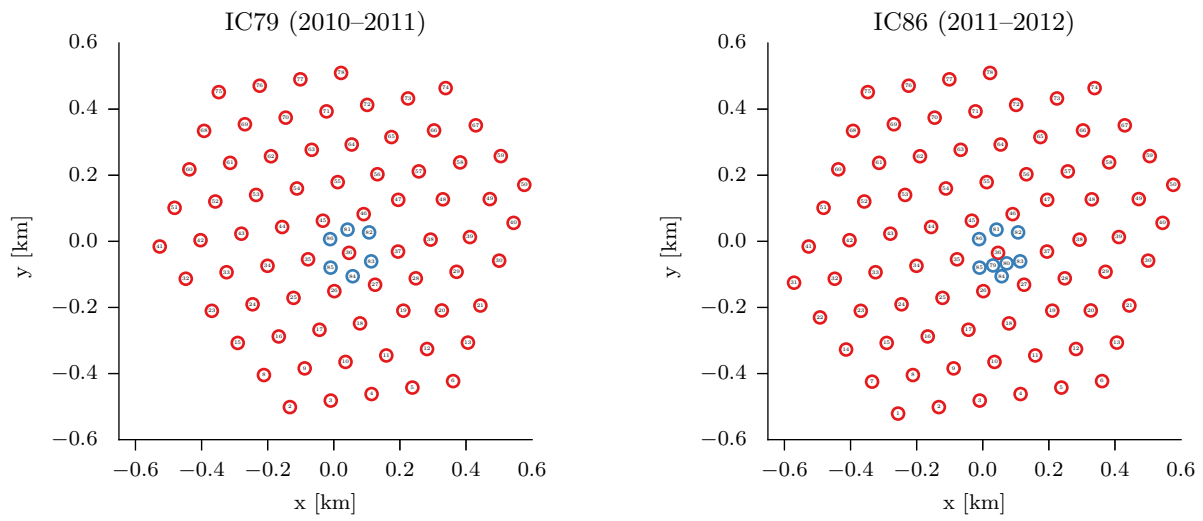
This combined vertex, time, direction, and energy reconstruction, called **Monopod**, is the single-source specialization of the more general Millipede reconstruction described in [87, Section 4.4], which itself is a generalization of the earlier Credo reconstruction [98]. The key advantage over the older Credo algorithm is computational efficiency. The optimization of the likelihood function over the vertex position, time, and direction of the cascade is aided by an analytic calculation of the vector gradient of Λ_j with respect to these 6 parameters. Because the derivatives of B-spline basis functions are linear combinations of B-splines of lower order, the value and gradient of the tensor-product B-spline surface that describes Λ_j can be evaluated in only $2\times$ the number of operations that it would take to calculate the value only, as opposed to the $6\text{--}12\times$ it would take to evaluate the gradient by finite differences [15]. When combined with the profiling out of the energy dimension by solving Equation (3.21), this reconstruction is ~ 8 times faster than the older Credo reconstruction. Its energy and angular resolution when applied to general event samples were presented in [16]. The resolutions for the event sample isolated in this work are presented in Section 5.5.1.

Chapter 4

Event selection

This chapter describes the selection procedure that was applied to obtain the event sample used in the analysis that will be described in Chapter 6. The selection consists of two major phases: a high-efficiency, 2-part pre-selection that reduces the data rate from ~ 3 kHz (~ 100 billion events per year) to ~ 0.3 Hz (~ 10 million events per year), and a harsh neutrino-level selection that reduces the sample further to 388 events in 2 years of data-taking, 95% of which are neutrino interactions within the fiducial volume of IceCube.

The data were taken with two slightly different detector configurations, called **IC79** and **IC86**. From May 2010 to May 2011 the detector ran with 79 **strings**, missing 5 **strings** on the grid-southwest side of the array and 2 **strings** in the **DeepCore** infill as shown in Figure 4.1. More importantly, though, the pre-selections applied to each year were different. Figure 4.2 shows an outline of the individual steps of the event selection, which are different for each data-taking year in the pre-selection phase but uniform for the final neutrino-level event selection. Table 4.1 gives the total number of events in the data sample and the expected contributions from various sources of neutrinos at every stage.



(a) Nearly-complete detector configuration used from May 2010 to May 2011.

(b) Complete detector configuration used from May 2011 to May 2012. One layer of string (5) was added to the grid-southwest side and 2 string to the DeepCore in-fill.

Figure 4.1: Surface geometry of IceCube **string** configurations used in this analysis.

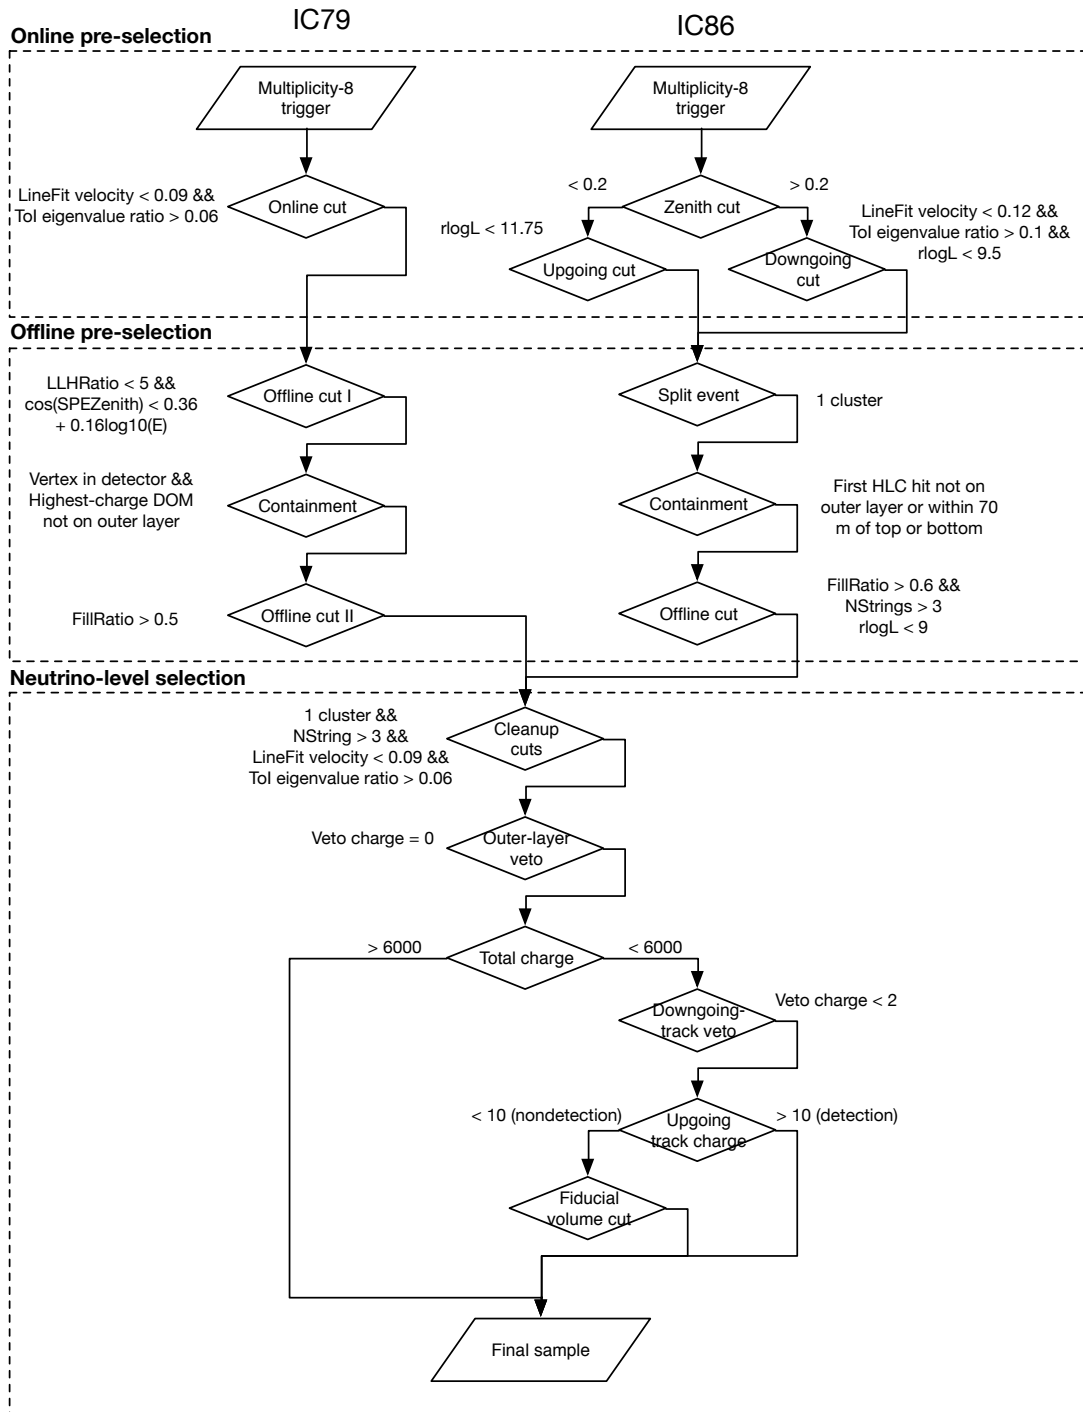


Figure 4.2: Outline of the event selection used to obtain the data sample from this analysis starting from the DOM multiplicity trigger.

Table 4.1: Number of events at each stage of the event selection for each of the data-taking periods used in this work. The data rates before the final selection are estimated from ~ 40 hours of data spread throughout each period. The conventional atmospheric neutrino rate is estimated from the model of [47] and the prompt atmospheric neutrino rate from the model of [14]. The E^{-2} astrophysical model is a 1 : 1 : 1 flux at the usual benchmark level ($10^{-18} \times (E/100 \text{ TeV})^{-2} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$, cf. Section 5.3), and the $E^{-2.46}$ astrophysical flux model is the best fit from this work (cf. Table 6.1). The final number of events represents the result of this event selection on the full 641-day data sample.

Selection	Data	Conv. atm.	Prompt atm.	E^{-2} astrophys.	$E^{-2.46}$ astrophys.
IC79 (312.5 days)					
Trigger	62503046140	-	-	-	-
Any filter	4773158264	17008	106	70	467
Cascade filter	754512648	6916	70	55	335
Cascade Level 3	20376876	4269	58	50	288
Containment	8923355	4145	57	48	282
Cleanup cut	3877563	3085	49	44	240
Outer veto	1575679	2525	39	35	192
Inner veto	405350	2109	34	28	168
Neutrino level	171	115	7	14	40
IC86 (328.9 days)					
Trigger	79116291072	-	-	-	-
Any filter	18912105539	20969	120	73	532
Cascade filter	752304582	14425	95	62	431
Cascade Level 3	8107396	6123	64	49	303
Containment	4144354	5652	54	43	269
Cleanup cut	1661411	3441	39	34	193
Outer veto	1023560	2962	32	27	160
Inner veto	296927	2737	30	25	150
Neutrino level	218	161	7	14	40

4.1 Pre-selection

In order to explain why the event selection is split between a pre-selection and a neutrino-level event selection, it is necessary to take a detour into the details of data handling and simulation in IceCube and to discuss how these somewhat tortured procedures arose from the constraints of limited data transmission bandwidth, computing time, and data storage capacity. The available satellite bandwidth limits the maximum possible passing rate of the initial “online” selection discussed in Section 3.2, and the only events available for later analysis are the ones that pass at least one of the online filter conditions. Analyses like the one that will be presented in Chapter 6 are not time-sensitive, and are typically designed “offline,” or after the data have already been taken, primarily because the computationally expensive simulations required to design an optimal event selection and then interpret the data have lagged several years behind data-taking. This is most severe for simulations of the penetrating atmospheric muon background: since an optimal event selection rejects nearly all penetrating muon events, many more events must be simulated to get a robust estimate of the passing rate than for neutrino events, where most events will survive the selection.

The space required to store these background simulations can be considerable: for example, a typical background simulation set (6939) requires 183 CPU-days of computing time and 140 GB of storage for each day of simulated detector livetime after the online filters and offline post-processing (discussed below) have been applied. Ideally one would like to have simulated a background sample at least as large as the one collected in the experiment, but storing the simulated data at such a low filter level would take an impractical 51 terabytes for only a single year of simulated livetime. Instead, high-statistics background simulations are only stored after further cuts have been applied to reduce the data rate by at least another factor of 10 from the passing rate of the online filters. In order to avoid wasting computing time, this further reduction is chosen to be as general to a detection channel as possible so that it can be shared between multiple analyzers at different institutions within the IceCube collaboration. Unfortunately, these high-statistics simulations only exist for the cascade channel, based on events that passed the online cascade filter. For this purely practical reason, the selection for this analysis is based on the online **cascade filter**. Unfortunately both stages of the pre-selection were chosen differently between the two data-taking years, necessitating more convoluted data handling than would have been strictly necessary had the needs of an analysis like this one been anticipated earlier. With the problem outlined, we can move on to discuss the details of each pre-selection.

4.1.1 Online pre-selection

The **IC79 online cascade filter**, shown on the left side of the upper panel of Figure 4.2, was quite simple, based only on the **line fit speed** v_{LineFit} and **tensor-of-inertia eigenvalue ratio** q_{ToI} ¹. The filter selects cascade-like events by requiring that the **line fit speed** be small (as happens when the hit pattern expands outwards from a single point instead of moving along a line) and that the **tensor-of-inertia eigenvalue ratio** be large (as happens when the hit pattern is spherical rather than elongated). The exact cut values were²

$$(v_{\text{LineFit}} \leq 0.09 \text{ [m ns}^{-1}\text{]}) \wedge (q_{\text{ToI}} \geq 0.06) \quad .$$

The passing rate of this filter averaged over the year was 27 Hz, or 1/85 of the 2315 Hz global trigger rate.

The **IC86 online cascade filter**, shown on the right side of the upper panel of Figure 4.2, was somewhat more complicated, based on more observables. In addition to the **line fit speed** and **tensor-of-inertia eigenvalue ratio**, the 2011 filter considered the zenith angle of the **SPE Pandel** likelihood fit with an infinite track hypothesis and the **reduced log-likelihood** from an **SPE Pandel** likelihood fit with a point-like, isotropically emitting cascade hypothesis³. The cut was parameterized on the cosine of the SPE track zenith angle: if the best-fit track hypothesis had $\cos \theta > 0.2$ ($\theta \lesssim 78^\circ$), like the vast majority of the penetrating muon background, then a cut similar to the IC79 filter was applied, augmented with an additional cut on the **reduced log-likelihood** $-\ln L_{\text{SPE,cascade}}/n_{\text{dof}}$. For $\cos \theta_{\text{SPE,track}} < 0.2$ the **line fit speed** and **tensor-of-inertia eigenvalue ratio** cuts were removed, and the **reduced log-likelihood** cut relaxed. This modified condition was intended to pass a larger fraction of low-energy, single-string events for which the **tensor of inertia** is poorly defined. However, this class of events proved to be too difficult to separate reliably from the remaining penetrating background in later selection stages, and was explicitly removed again in the next pre-selection stage. The exact cut values were

$$\left\{ \begin{array}{ll} (v_{\text{LineFit}} \leq 0.12) \wedge (q_{\text{ToI}} \geq 0.1) \wedge (-\ln L/n_{\text{dof}} < 9.5) & \cos \theta_{\text{SPE,track}} > 0.2 \\ (-\ln L/n_{\text{dof}} < 11.75) & \cos \theta_{\text{SPE,track}} < 0.2 \end{array} \right. \quad .$$

The passing rate of this filter averaged over the year was 29.2 Hz, or 1/95 of the 2784 Hz global trigger rate.

¹See Section 3.4.1.1 for a description of LineFit and TensorOfInertia.

² Logical operators are given in Boolean rather than C-style notation, so \wedge denotes logical AND rather than XOR, and \vee denotes logical OR

³See Section 3.4.2 for a description of the Pandel likelihood functions and related observables.

4.1.2 Offline pre-selection

The **offline** pre-selection⁴ for **IC79**, shown on the left side of the middle panel of Figure 4.2 consisted of two nested cuts. The first of these was based on the difference between the best-fit **log-likelihoods** of **Pandel** likelihood fits (see Section 3.4.2) with track and cascade hypotheses, with an additional 2-dimensional cut on the track fit zenith angle and reconstructed cascade energy. The energy-dependent zenith angle cut removes the large population of penetrating muon background events that are reconstructed as straight-downgoing with a deposited energy less than 1 TeV. The exact cut values were

$$(\ln L_{\text{track}} - \ln L_{\text{cascade}} < 5) \wedge (\cos \theta_{\text{track}} < 0.36 + 0.16 \log_{10}(E_{\text{cascade}}/\text{GeV})) \quad .$$

After this first cut, the sample is split into a contained and an un-contained sample. The un-contained sample consists of all events where the cascade vertex is reconstructed outside the polygon formed by the outer **strings** (see Figure 4.1), or where the **DOM** with the most collected charge in the event was on one of the outer **strings**. This part of the sample was not used for this analysis. The remaining contained sample undergoes another cut on an observable called the **fill ratio** [99], which is the fraction of DOMs inside a sphere centered on the reconstructed cascade vertex that trigger. The radius of the sphere is a function of the average distance of hit DOMs to the vertex. For true cascade events the **fill ratio** is close to 1, while for track-like events it is smaller. In this second cut the selection requires that

$$R_{\text{fill}} > 0.5$$

The year-averaged passing rate for this offline pre-selection is 0.35 Hz.

The offline pre-selection for **IC86**, shown on the right side of the middle panel of Figure 4.2, was designed independently. It starts by attempting to detect the presence of muons from multiple independent air showers in the single detector readout⁵. Events in this class look neither like single muon tracks or cascades, and tend to confuse reconstruction algorithms that expect one or the other event type. Multiple muons are detected by applying an algorithm called **TopologicalSplitter** [100] that sorts the hit pattern into clusters that could be causally connected. If hits from neighboring clusters are separated by a sufficiently small spacetime interval, the clusters are merged. The number of such clusters left when the algorithm terminates is related, but not

⁴In IceCube jargon this detection-channel-specific pre-selection is called “Level 3” in analogy to Level 1, the stream of events available to the online filters, and Level 2, the stream of events available for offline analysis.

⁵In IceCube jargon, these are called “coincident events.”

necessarily exactly equal to, the number of independent air showers contributing to the detector readout. Nonetheless the cluster-merging tolerances can be chosen to be sufficiently small that having only one cluster can be taken as evidence that there is no second muon. For the purposes of this analysis any events with multiple clusters are discarded. Next, the events are divided into a contained and an uncontained sample, though in contrast to the IC79 pre-selection the distinction is made solely based on the location of the first HLC hit in the event. If it is on the outer layer of strings (see Figure 4.1) or in the top or bottom 70 m of instrumented volume, then the event is considered uncontained and discarded for the purposes of this analysis. Finally, the single, contained events are subjected to cuts on the fill ratio, reduced log-likelihood of the Pandel likelihood cascade reconstruction, and the minimum number of non-DeepCore strings with HLC hits. The condition is

$$(-\ln L_{\text{cascade}}/n_{\text{dof}} < 9) \wedge (R_{\text{fill}} > 0.6) \wedge n_{\text{strings}} > 3 \quad .$$

The year-averaged passing rate for this offline pre-selection is 0.30 Hz.

4.2 Neutrino-level event selection

The next stage of the event selection is to reduce the pre-selected sample of nearly 20 million events to a sample containing mostly neutrinos. This proceeds in 4 stages: a set of clean-up cuts to homogenize the samples from the two different pre-selections, an outer-layer veto to remove penetrating muons that trigger the outer layer of the detector, an inner down-going track veto to remove lower-energy muons that miss the outer layer, and a charge-dependent fiducial volume cut that effectively makes the efficiency of the inner veto proportional to the background-to-signal ratio as a function of deposited charge.

The description of each step of the selection begins with a summary of its purpose, set in italic type inside a box as follows:

<i>Purpose of the selection step</i>

4.2.1 Clean-up cuts

<i>Homogenize IC79 and IC86 samples</i>

A careful reader will notice that the first stage of the neutrino-level event selection shown in Figure 4.2 repeats selections made earlier in the selection chains for both data samples. This is intentional and intended to select the most reliable subset from both selections. The single-cluster selection and minimum **HLC string** count remove potentially un-reconstructable events from the sample. This is important, as the inner down-going track veto requires a reliably reconstructed cascade vertex in order to function. By the time the inner veto was developed, however, the **IC79** selection had already been frozen for production of high-statistics penetrating muon background simulation, and so the cut had to be added after the fact. The remaining cuts were necessary because at the time when the neutrino-level event selection was developed, no high-statistics penetrating muon background simulation existed for the **IC86** sample, and the background for **IC86** had to be estimated by analogy to the **IC79** sample, which had been thoroughly simulated. The **line fit speed** and **tensor-of-inertia eigenvalue ratio** cuts were added to select only the subset of the **IC86** pre-selection that would have passed the **IC79 online filter**.

4.2.2 Outer-layer veto


Reject events that start on the edge of the detector

The next stage of the selection uses the outer-most layer of **DOMs** shown in Figure 4.3 as an active muon veto in a way that is very similar, but not exactly identical, to the veto used in [12, 13]. The basic principle is the same, though: we define an event start time, and demand that no **DOMs** in the veto region trigger before that time. The start time is defined to be the time when a minimum number of **PE** Q_{start} have been collected in the entire detector. This threshold is proportional to the total collected charge Q_{total} between fixed limits:

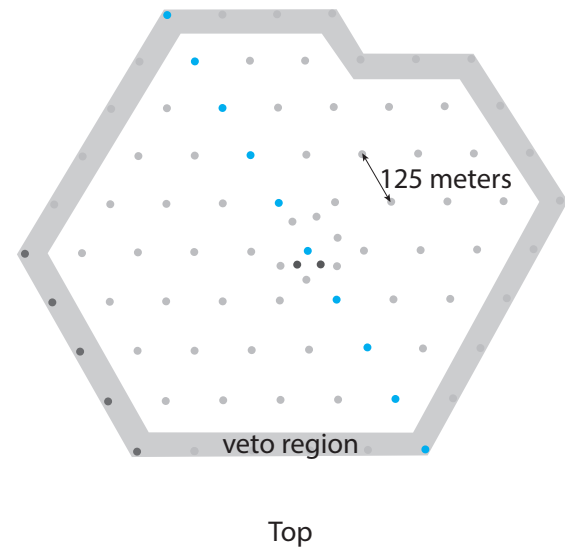
$$Q_{\text{start}} \equiv \begin{cases} 3 & Q_{\text{total}} < 72 \\ Q_{\text{total}}/24 & 72 \leq Q_{\text{total}} < 6000 \\ 250 & Q_{\text{total}} \geq 6000 \end{cases} \quad . \quad (4.1)$$

The constant of proportionality was chosen for consistency with [12, 13].


Figure 4.4 shows a demonstration of the veto algorithm. All **HLC pulses** are arranged in time order, and 3-microsecond-wide time window slides through the time-ordered stream of **pulses** until it contains a total charge of Q_{start} . If any of the pulses currently in the window came from **DOMs** in the veto region,



(a) Cross-section view of the detector geometry, showing the outer-layer veto region in gray. (Drawing by C. Kopper)



(b) Top view of the detector geometry. The strings in the cross-section at left are shown in blue. The strings shown in darker gray dots on the grid-southwest side were part of the IC86 run only; for IC79 the veto region is one layer further in. (Drawing by C. Kopper)



(c) Optical attenuation length at 400 nm as a function of vertical depth. The dust-layer veto region, shown as a gray band, extends from the depth of maximum absorption to just below the next lower absorption minimum. If the veto region were restricted to the outer strings, then muons could pass the veto much more easily if they entered the detector at the depth near the absorption maximum.

Figure 4.3: Geometry of the outer-layer veto


then the event is vetoed⁶. The penetrating muon event shown in Figure 4.4a starts in the veto region before continuing on into the fiducial region, whereas the contained-vertex neutrino event shown in Figure 4.4b starts in the fiducial region. Some photons do diffuse out into the veto region, but only after the veto has been disarmed by photons induced at the neutrino interaction vertex.

The veto region, shown in Figure 4.3 consists of all DOMs on outer strings, plus


1. the top 90 m of the detector, measured from the first DOM on the deepest string,
2. the bottommost active DOM on each string, and
3. all DOMs between vertical depths of 2050 m (the depth of maximum optical absorption) and 2170 m (just below the next absorption minimum).

These choices bear some explanation. The top layer is thicker than the side layers because the penetrating muon rate is highest from the zenith, while the bottom layer is only one DOM thick because the only up-going muons are neutrino-induced. There is an additional potential background from large stochastic losses of highly inclined penetrating muons underneath the detector. A single layer is sufficient to protect against this background, however, since in order for the event to appear to start in the fiducial volume the photons would have to be less likely to trigger DOMs closer to their point of origin than further away. The veto region in the middle of the detector is necessary because of the particular optical properties of the South Pole glacier. The attenuation length of the ice at the wavelengths where the DOMs are sensitive, shown in Figure 4.3c, is on average on the order of half a string spacing or more, making it generally difficult for TeV muons to pass between two strings undetected. The exception is Dust Peak D [90] centered at 2050 m depth, known in IceCube jargon simply as “the dust layer.” At this dust peak the attenuation length drops to 20 m, raising the effective threshold of the DOMs at that depth significantly. Inclined muons can sneak between two strings in the outer layer and appear to “start” at the depth of around 2100 meters when the absorption length increases and a previously-undetected energy deposition rate suddenly becomes detectable. The dust layer veto region covers the depths where such events would appear. Again, it only needs to cover the region below the dust peak, because all penetrating muons are down-going.

⁶This is more stringent than the condition used in [12, 13], where an event was vetoed if it contained at least 3 PE or pulses from 3 DOMs, whichever was larger.



(a) Photoelectrons collected per 25 ns in all DOMs with local coincidence in the veto region (red) and fiducial region (blue) as a function of time in a penetrating muon background event with 16644 PE total. The 3-microsecond-wide veto window ends once 250 PE have been collected anywhere in the detector. The event is vetoed if the veto window contains any hits from the veto region (red). Hits in the fiducial volume (blue) do not cause the event to be vetoed. In this example, the veto window extends from 7043 to 10043 ns and contains hits from the veto region (red), so the event is rejected.



(b) A neutrino interaction in the fiducial region with 14660 PE total. The veto window (in this example, 6898–9898 ns) contains only hits from the fiducial region (blue) and no hits from the veto region (red), so the event is accepted. Photons from the event do not have to be completely contained in the fiducial volume: in this example, photons do diffuse back out to the veto layer (see inset plot), but do not fall in the veto time window.

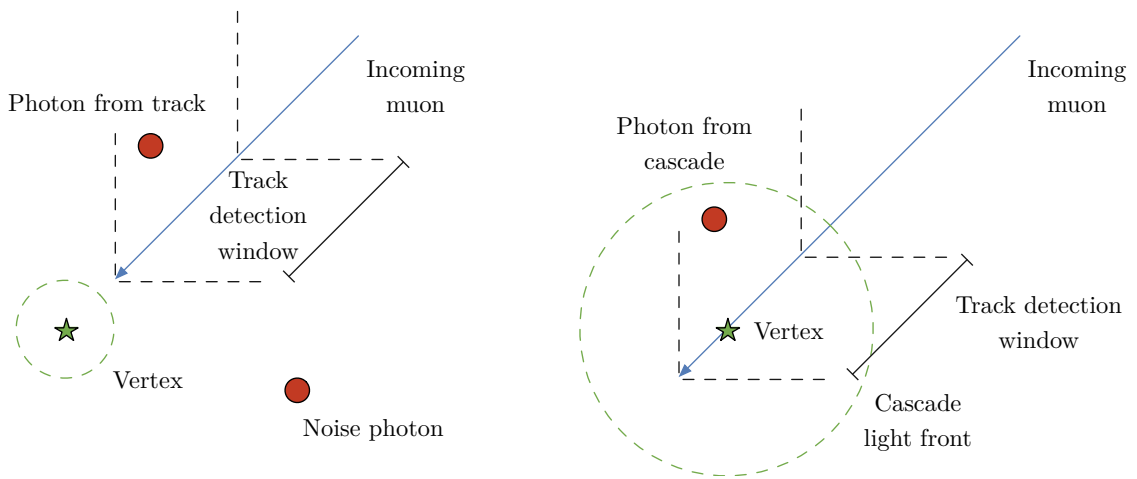
Figure 4.4: A demonstration of the outer-layer veto algorithm.

4.2.3 Inner track veto

Search for signs of dim penetrating muons before the alleged neutrino vertex

For very bright events with thousands of detected photons, the outer-layer veto is sufficient to suppress the muon background below the level of atmospheric neutrinos (Figure 4.10, center panel). As the energy threshold is lowered, the number of background muons increases rapidly, while their average energy loss rate decreases, a combination that overwhelms the ability of the single layer to reject incoming muons. In order to extend the selection to lower energies, a second kind of veto is required. The first modification is to remove the requirements that veto photons be detected on the outermost layer of DOMs and in local coincidence. This allows isolated photon detections anywhere in the instrumented volume to veto an incoming track, which lowers the energy threshold but also loses signal events to spurious vetoes caused by noise. In order to mitigate the signal loss, a second modification is required: pulses are only considered for veto if they are detected at a time and position consistent with an incoming track but inconsistent with the cascade vertex (as reconstructed with Monopod), as shown in Figure 4.5a.

Figure 4.6 shows a demonstration of this inner downgoing-track veto. The quantities on the x axis are time residuals Δt , or the difference between the photon detection time and the earliest possible arrival time for Cherenkov photons as defined in Equations (3.13) and (3.14) and illustrated in Figure 3.11. In the absence of scattering and noise, Δt would be identically zero for all detected photons given the correct track or cascade hypothesis. Scattered photons, however, must cover more distance at the local speed of light, making the time residual positive. Negative time residuals, on the other hand, can only arise when the event hypothesis is wrong. For example, if an event that looks like a contained neutrino interaction, with a muon track emanating from a large cascade, is in fact caused by a large stochastic energy loss of a penetrating muon, then the Cherenkov photons the muon induces before the large stochastic loss may be detected well before the times expected for photons induced at the cascade vertex. This is the situation depicted in the upper panel of Figure 4.6a. This event passed the outer layer veto, but a number of photons were detected with negative time residuals with respect to the reconstructed cascade vertex (inside the gray band). These photon detections are a combination of Cherenkov photons from the muon whose stochastic loss dominates the event as well as random noise. Simply requiring that there be no negative time residuals would reject this event, but also a large fraction of true neutrino events, because while the Cherenkov photons from the muon would not exist, the random noise would. In order to mitigate this signal loss, photons are only considered



(a) Penetrating muon before its largest energy loss. The dashed grey lines mark the positions at which photons induced by a muon would be detected with minimal and maximal delay. The photon that falls inside this window is counted towards the veto total, while the random noise photon that falls outside the window is not.

(b) Penetrating muon after its largest energy loss. The dashed circle marks the positions where photons propagating from the vertex at the speed of light in ice would be detected with minimal delay. Here the photon is not counted towards the veto since it is detected at a time compatible with propagation from the reconstructed vertex.

Figure 4.5: An illustration of the incoming-muon veto procedure. Each panel shows a snapshot in time with the current position of the muon marked by the blue arrowhead and the position of the reconstructed vertex marked by a green star. (a) shows a penetrating muon before its largest energy loss with a photon detection that counts towards the veto, while (b) shows the same configuration after the largest energy loss with an ambiguous photon detection that does not count towards the veto.

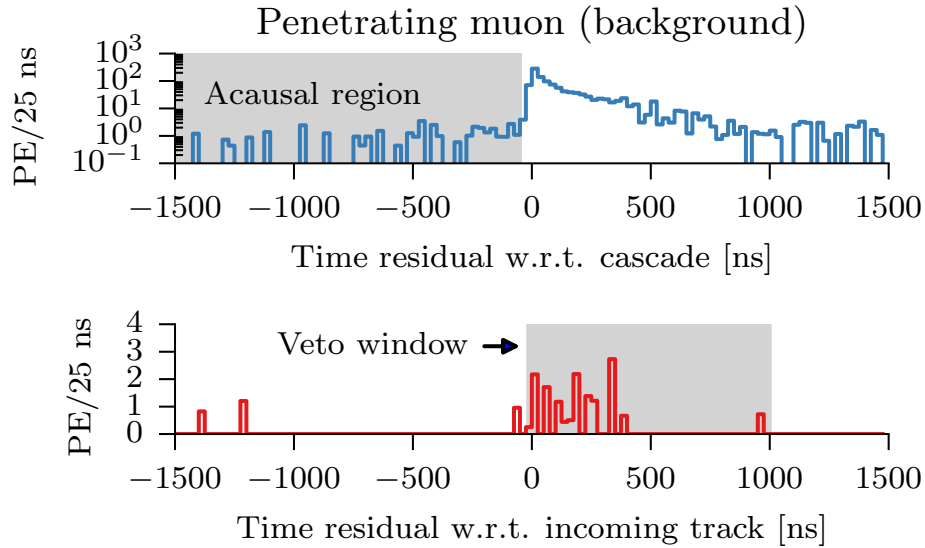
for veto if their **time residuals** are compatible with an incoming muon track. The lower panel of Figure 4.6a shows the photon detections from within the gray band from the upper panel, but the **time residuals** are calculated with respect to the incoming muon track rather than the reconstructed cascade. 18 of these have $-15 < \Delta t_{\text{track}} < 1000$ ns, and so are compatible with the incoming track. The remaining noise photons are uniformly distributed in time and less likely to fall in the track veto window. Figure 4.6b shows a neutrino event with a similar total charge. In this case there are also photons with negative **time residuals**, but none of them fall in the veto window of an incoming track.

This veto construction of course requires a track hypothesis that is reasonably close to the true incoming track, as the veto window of a terribly misreconstructed track could potentially miss all of the photons detected from the true track. This presents a complication for a selection like this one, where the most resilient background comes from single muons with a single disproportionately large stochastic energy loss. Such events can be completely dominated by the photons induced by the largest energy loss, causing track reconstruction algorithms like those presented in Section 3.4.2 that assume uniform light emission to fail to find the correct direction. The position and time of the vertex can, however, be reconstructed reliably regardless of the presence of a detectable muon track. In order to ensure that veto photons will be found, the search is repeated for each of 104 different down-going track hypotheses (chosen from the upper hemisphere of a HEALpix [101] grid) that pass through the reconstructed vertex. The track hypothesis with the largest number of associated veto photons is considered the best. A photon is associated with an incoming track if it is detected

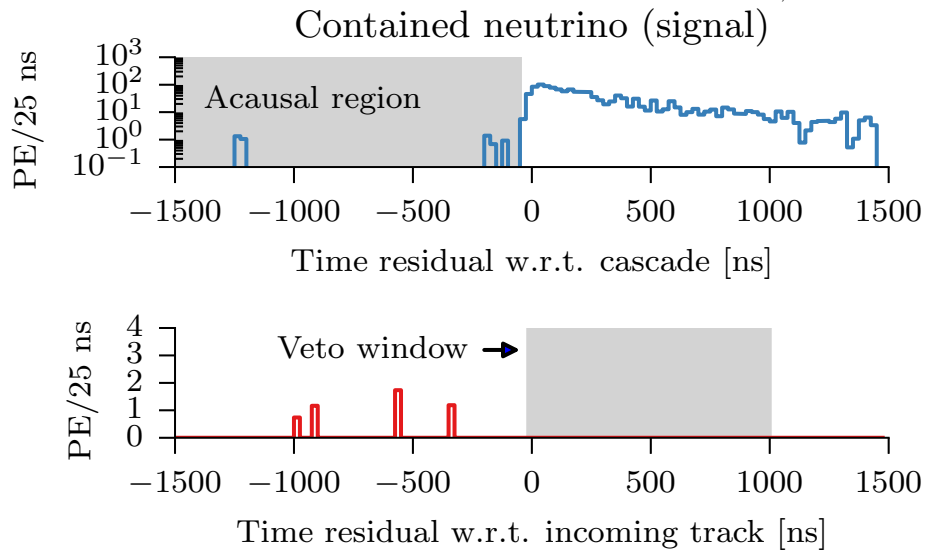
- at least 50 ns before the earliest possible time for a photon induced at the previously reconstructed vertex,
- between 15 ns before and 1000 ns after the earliest possible time for a photon induced by the hypothetical muon,
- and no more than 100 m from the hypothetical muon trajectory.

The event is rejected if the best track has more than two associated photons. This threshold is chosen to avoid spurious vetoes of neutrino events when the grid search finds a noise photon⁷ that can be connected with a muon track passing through the neutrino interaction vertex.

⁷Because each DOM triggers in isolation at a rate of approximately 585 Hz (much higher than the 13 Hz local coincidence trigger rate) [102], it is quite often possible to connect one photon to the vertex.



(a) The upper panel shows photoelectrons collected per 25 ns in all DOMs as a function of time residual with respect to the reconstructed cascade vertex for a penetrating muon event with 1234 PE that passed the outer-layer veto. The photoelectrons with negative time residuals (inside the gray band) can't be causally connected to the cascade vertex. The time residuals of these PE are shown again in the lower panel, this time with respect to the best-fit incoming muon track. Since 18 PE fall within the 1-microsecond-wide veto window for the track, the event is rejected.



(b) A neutrino interaction in the fiducial region with 1183 PE. Though PE are collected in the acausal region, none of them are associated with an incoming muon track, so the event is retained.

Figure 4.6: A demonstration of the inner downgoing track veto algorithm.

4.2.4 Fiducial volume scaling

Scale veto thickness up with decreasing brightness to make the muon rejection efficiency proportional to the muon rate

Since the effectiveness of the track-based veto is proportional to the probability of detecting at least 2 photons from an incoming muon before the reconstructed vertex, it increases in proportion to the number of detectable photons the muon induces and the number of PMTs it passes on its way to the vertex as shown in Figure 4.7. This relationship can be exploited to maintain sufficient penetrating muon rejection at low energies by requiring a minimum distance between the reconstructed vertex and the edges of the instrumented volume that increases as the number of collected photons decreases. In order to be accepted, the **homogenized total charge**⁸ Q_{total} and vertex position must satisfy

$$\log_{10}(Q_{\text{total}} - 100) > 3.41 - \frac{(d_{\text{side}})^{1.74}}{17266} \quad (4.2)$$

$$\log_{10}(Q_{\text{total}} - 100) > 3.40 - \frac{(d_{\text{top}} - 100)^{1.88}}{23710} \quad , \quad (4.3)$$

where d_{side} is the distance from the nearest face of the instrumented volume in meters and $d_{\text{side}} = 500 - z$ is the distance from the top of the instrumented volume. Figure 4.8 shows the fiducial volume corresponding to (4.3) at 4 different charge thresholds. For the dimmest events the fiducial volume is reduced to the DeepCore subarray with the remainder of the detector used as a veto; as the total charge increases the veto reduces to the outermost layer of PMTs as in [12].

⁸The **homogenized total charge** is the sum of charges of all **HLC pulses** detected on non-**DeepCore DOMs** that do not contribute more than 50% of the total. This quantity has a smaller variance with respect to deposited energy than a naïve sum over all pulse charges.




Figure 4.7: Fraction of pre-selected penetrating muon background events (Section 4.1.1) that pass the veto conditions (Sections 4.2.2 and 4.2.3), derived from MC simulation. The outer-layer veto reduces the background of the highest-energy muons by 10^4 , but degrades rapidly at lower energies. The incoming-track veto scales in a similar way with respect to energy, but is more sensitive because it considers isolated photon detections. In contrast to the outer-layer veto, its efficiency also improves with increasing distance d from the detector border of the reconstructed vertex.

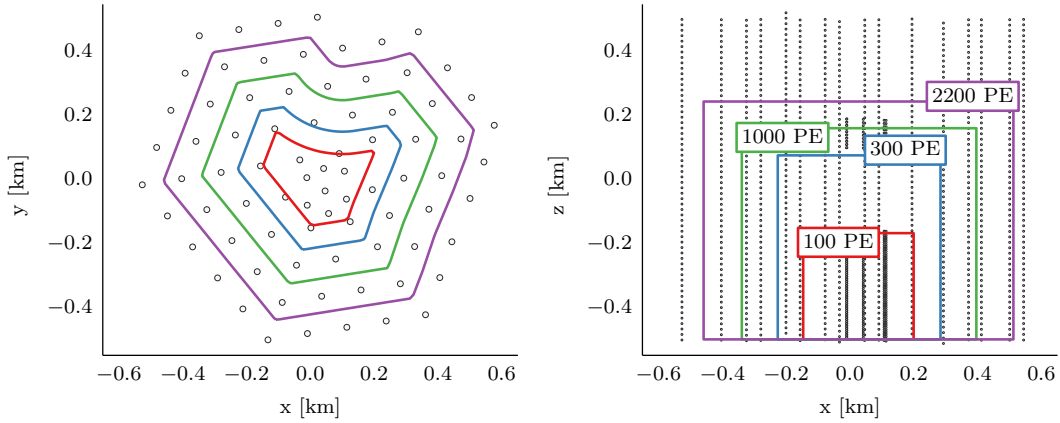


Figure 4.8: Fiducial volume scaling function evaluated at four different photon counts. Left: Top view, showing the positions of the IceCube strings and the boundaries of the fiducial volume for events with a given total photon count. Right: Side view, showing the modules along strings.

4.2.5 High-energy and upgoing-track event retention

Relax fiducial volume cut for clear neutrino candidates

There are two kinds of events that can be classified as neutrino events without recourse to the charge-dependent fiducial volume cut. The first are events that pass the outer layer veto with a homogenized total charge of more than 6000 PE. As shown in Figure 4.2, these are kept explicitly for consistency with the selection of [12]. The second are events with upward-going muon tracks, which can only be induced by neutrino interactions. These are identified by inverting the incoming-track veto as shown in Figure 4.9 and selecting events with more than 10 PE associated with an outgoing, upward track. The grid search for the best track hypothesis is done over 104 candidate directions as with the incoming-track search, and uses the **Monopod** vertex as an anchor point, but the directions are chosen from the lower hemisphere of a HEALpix [101] grid rather than the upper hemisphere. The search uses HLC pulses only and a higher acceptance threshold than the rejection threshold for incoming muons in order to avoid false positives⁹.




Figure 4.9: Identifying a neutrino-induced muon by inverting the incoming-track veto (cf. Figure 4.5). Photons induced at the cascade vertex spread outwards at the speed of light in ice, while the muon moves at the speed of light in vacuum. Eventually the muon out-runs the light front from the cascade, and photons collected in the track detection window can be used to positively identify an out-going muon in the event.

⁹ The most insidious of these are sub-threshold muons that pass very close to or even through a DOM, producing an isolated SLC pulse with up to 100 PE. These quite rare events are not accounted for in simulation, but can be avoided by using HLC pulses only.

4.3 Final sample

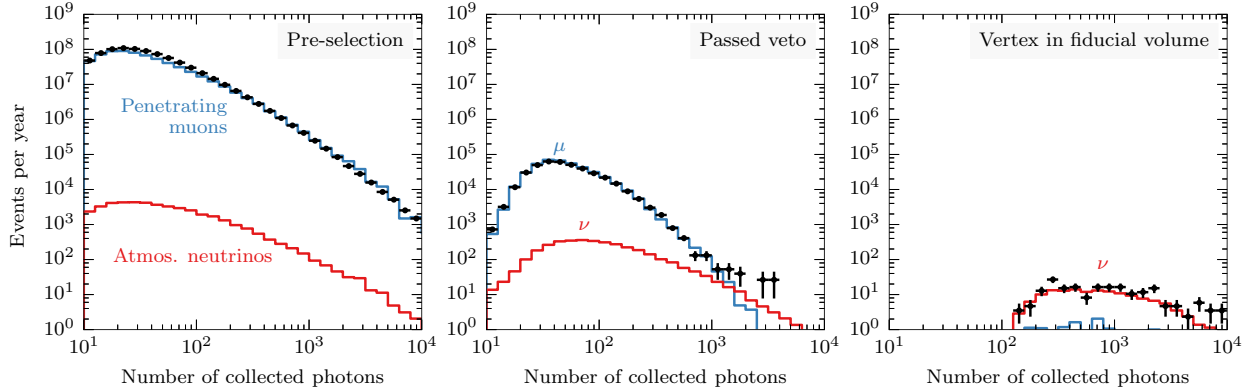
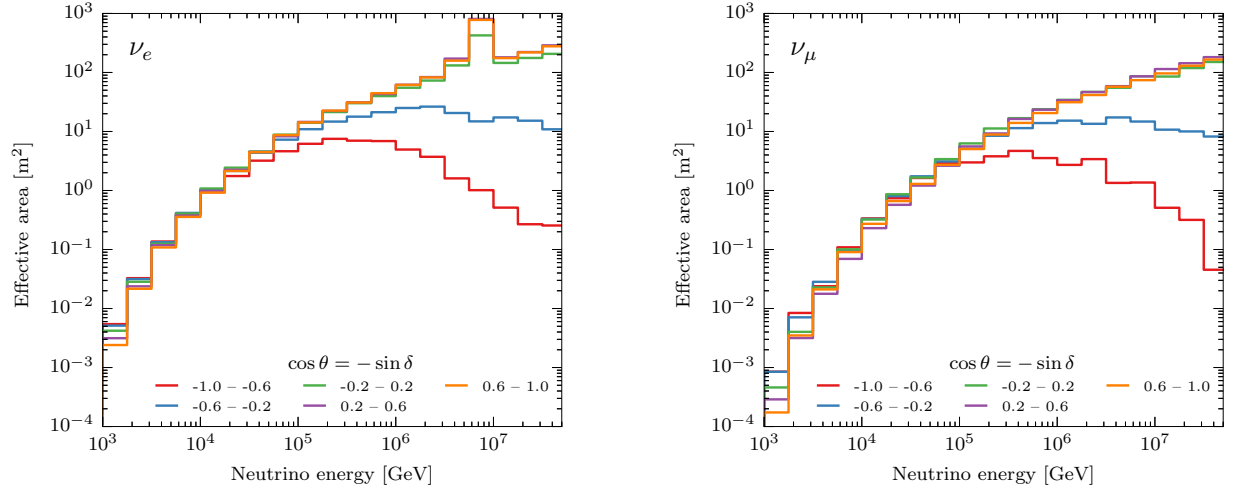



Figure 4.10: Distribution of photon counts per event after each stage of the event selection. The total number of collected photons is on average proportional to the total deposited energy; for example, 10^3 photons correspond to roughly 10 TeV deposited energy. The stepped lines show the prediction from Monte Carlo simulation of penetrating atmospheric muons (blue) atmospheric neutrinos (red), while the points show experimental data. Left: Pre-selected events transmitted from the South Pole (Section 4.1.1). Center: Removed events with veto hits (Sections 4.2.2 and 4.2.3). Right: Fiducial volume scaled with photon count (Section 4.2.4).

The selection presented in the previous section creates the neutrino-dominated sample shown in the right panel of Figure 4.10. 283 cascade and 105 track events passed the final selection criteria in 641 days of data-taking. Of those 388 events, 106 deposited more than 10 TeV at the cascade vertex, and 9 deposited more than 100 TeV. At high energies the selection overlaps nearly completely with the selection of [12]: 7 of the 9 events depositing more than 100 TeV were also in the previous selection. The effective area of the selection is shown in Figure 4.11.



(a) Electron neutrinos. The peak just below 10 PeV is the Glashow resonance (see Section 3.1.2), and the reduction at large zenith angles is due to neutrino absorption in the Earth.

(b) Muon neutrinos. The fluctuations visible at high energies for up-going neutrinos are statistical artifacts.



(c) Tau neutrinos. The effective area at high energies for up-going neutrinos is much higher than for other neutrino flavors, because τ lepton produced in CC interactions in the Earth “re-generate” ν_τ [48].

Figure 4.11: Neutrino effective area of this selection as a function of neutrino energy and zenith angle (or equivalently, declination) for the 3 neutrino flavors. The effective areas averaged between neutrinos and antineutrinos.

Chapter 5

Analysis methods

The events that pass the selection defined in the previous chapter arise from four sources: conventional atmospheric neutrinos, prompt atmospheric neutrinos, penetrating atmospheric muons, and high-energy extraterrestrial neutrinos. We wish to disentangle the contributions of each of these components to the observed sample of events. We can do this by exploiting the fact that each of these components produces a distinct distribution in the observables chosen for this analysis: reconstructed deposited energy, reconstructed zenith angle, and the presence or absence of a detectable outgoing muon track. We then extract the contribution of each component by fitting the distribution of experimental observables to a combination of template functions, one for each component. This chapter presents the experimental observables, the techniques used to estimate their distributions for each component from Monte Carlo simulations, and the fitting technique used to extract the interesting physics parameters from the data.

5.1 Modeling the atmospheric neutrino flux

5.1.1 Neutrino event simulation

Neutrino events are simulated in IceCube by a specialized software package called NeutrinoGenerator [103, 104]. Some knowledge of its inner workings is required to understand how counts of simulated events are converted into predicted experimental event rates. This subsection will give a brief overview of the neutrino event simulation and how event weights are calculated.

The simulation of each neutrino event consists of three phases: injection, propagation, and final interaction. In the injection phase, a primary neutrino is chosen. The flavor (ν_e , ν_μ , or ν_τ) is chosen ahead of time for the entire simulation run, and ν or $\bar{\nu}$ chosen with 50% probability for each injected neutrino. The energy is drawn from a power law, typically $dN/dE \propto E^{-1}$ or E^{-2} , and the direction in the IceCube coordinate

system is chosen uniformly in ϕ and $\cos\theta$. The neutrino is then placed on the surface of the Earth, aimed along the chosen direction at a target surface surrounding the IceCube instrumented volume that defines the final interaction volume.

In the propagation phase the neutrino is moved from the surface of the Earth, through the various intervening matter to the final interaction volume. In each step ΔX , the neutrino may interact with a probability given by

$$p_{\text{interaction}} = 1 - e^{-\Delta x \sigma n} \quad , \quad (5.1)$$

where σ is the total interaction cross-section and n is the number density of targets (nucleons for neutrino-nucleon scattering, electrons for neutrino-electron scattering). Since the total neutrino-nucleon cross-section is on the order of 10^{-35} cm^2 at 1 TeV (rising to $\sim 10^{-33} \text{ cm}^2$ at 1 PeV) and nucleon densities in standard rock are on the order of 10^{30} cm^{-3} , interactions are quite rare. When they do occur, the neutrino is lost, and a new event is started. The lone exception are **CC** interactions of ν_τ , where the τ lepton can decay before losing much energy to produce another ν_τ , a phenomenon known as ν_τ regeneration [48].

The final interaction volume is a cylinder whose axis is parallel to the neutrino direction and centered on the origin of the IceCube coordinate system. The position of the final interaction in this volume, Δx , is drawn from a uniform distribution $dp_{\text{interaction,generated}}/dx = 1/L$ over the length of the cylinder, L . The weight associated with this forced interaction is

$$w_{\text{interaction}} \equiv \frac{dp_{\text{interaction,nature}}/dx}{dp_{\text{interaction,generated}}/dx} = Ln\sigma e^{-\Delta x \sigma n} \quad , \quad (5.2)$$

where again σ is the total interaction cross section. This is simply proportional to the differential interaction probability per unit length (which, to first order, is proportional to the cross-section), with a constant of proportionality given by the total length of the interaction volume. Once the interaction position has been chosen, a process and final state are chosen proportional to their partial cross-sections. In order to efficiently simulate neutrino-induced muon events that can reach the instrumented volume from far away, L is chosen such that the most energetic muon that could be produced in a **CC** interaction would only reach the edge of the minimum final interaction volume in 0.01% of cases, using the maximum range parameterization found in [72].

These final states are then passed through the remainder of the simulation chain and subjected to the same event selection as the experimental data. In order to calculate the expected experimental data rate, one

adds up the weights for all remaining simulated neutrino events. These are simply the ratio of the number of events expected in nature to the number that were simulated, or

$$w_i = w_{\text{interaction}} \frac{\frac{dN(E_i, \theta_i, T_i)_{\text{nature}}}{dE dA d\Omega dt}}{\frac{dN(E_i, \theta_i, T_i)_{\text{generated}}}{dE dA d\Omega}} \quad , \quad (5.3)$$

where $\frac{dN(E_i, \theta_i, T_i)_{\text{nature}}}{dE dA d\Omega dt}$ is the flux, differential in energy and solid angle, of neutrinos of type T_i and energy E_i at zenith angle θ_i . All units cancel except for a s^{-1} from the flux model, leaving a rate in s^{-1} . This allows us to convert from a neutrino flux model to distributions of observables in the experimental data.

We model each of the three neutrino sources we consider (conventional atmospheric neutrinos, prompt atmospheric neutrinos, and astrophysical neutrinos) using the same set of simulated events and simply exchanging the flux model in the numerator of Equation (5.3)¹. For upward-going neutrinos this is entirely appropriate, as neutrinos from different sources are indistinguishable on an event-by-event basis. Down-going atmospheric neutrinos, on the other hand, are often accompanied by atmospheric muons from the same air shower, whereas astrophysical neutrinos are never accompanied by muons. Since the event selection removes events with penetrating muons, however, we can account for the difference by reducing the effective atmospheric neutrino flux in the downgoing region. This correction will be presented in Section 5.2.

5.1.2 Conventional atmospheric neutrinos

We use the 3-dimensional Monte Carlo calculation of [47, 105] to model the conventional atmospheric neutrino flux. The primary advantage of this calculation over other, mostly older calculations like [23, 106–112] is that it was explicitly tuned to reproduce the atmospheric muon spectra measured precisely by the BESS spectrometer at 30, 2770, and 37000 m above sea level [113, 114] and the L3 detector at LEP at 430 m above sea level [115]. As a practical matter, it is also the benchmark model used in previous IceCube analyses [12, 13, 36, 58, 116].

Because it is a Monte Carlo calculation, however, it is defined on a discrete grid of neutrino energies and zenith angles. We interpolate between these points and extrapolate beyond the upper energy limit of the explicit calculation at 10 TeV using a parameterization of the form (2.18), with one modification. The curvature of the atmosphere becomes important for zenith angles larger than 60° , spoiling the simple relation between slant depth and local density given in Equation (2.4) that leads to the $1/\cos\theta$ term in the

¹The flux models for conventional and prompt atmospheric neutrinos are presented in the next two subsections, and the model for the astrophysical neutrino flux is presented in Section 5.3.

expression for lepton fluxes given in Equation (2.18). The zenith dependence of lepton fluxes in a curved atmosphere can be described by replacing $\cos \theta$ with the ad-hoc form [117]

$$\cos \theta^* \equiv \sqrt{\frac{\cos^2 \theta + p_1^2 + p_2 \cos^{p_3} \theta + p_4 \cos^{p_5} \theta}{1 + p_1^2 + p_2 + p_4}}, \quad (5.4)$$

using the parameters given in Table 5.1.

Table 5.1: Parameters of Equation (5.4) from [117].

p_1	p_2	p_3	p_4	p_5
0.102573	-0.068287	0.958633	0.0407253	0.817285

The points in Figure 5.1 show the MC points of [47], while the solid lines show the interpolation with the form of (2.18) and the effective local density correction from Equation (5.4). The parameters of the interpolation are taken from [118], and evaluated with [119].

5.1.3 Prompt atmospheric neutrinos

The modeling of the flux of atmospheric neutrinos from the prompt decays of charmed mesons is more uncertain. Since this flux has not yet been conclusively observed in either muons or neutrinos, there are no measurements that can be used to tune models directly. In addition, there are significant theoretical uncertainties in the production cross-sections for charm quark pairs in the forward region [120]. Accordingly, predictions for the flux [14, 44, 120–124] have varied by more than an order of magnitude in normalization, though some like the RPQM model of [44] have already been directly excluded by IceCube neutrino data [36].

We choose [14] as a representative model in keeping with previous analyses of IceCube neutrino data [12, 13, 36, 58, 116]. The prediction consists of three parts, a perturbative QCD calculation of the $p + p \rightarrow c\bar{c}$ production cross-section in the dipole framework with a phenomenological treatment of gluon saturation, a parameterization of $c\bar{c}$ fragmentation into final-state hadrons, and the numerical convolution of these with a primary cosmic ray flux using cascade equations similar to those presented in Section 2.1.1. Since the critical energy for charmed mesons is quite high², the local density of the atmosphere is irrelevant, and the predicted flux is isotropic. Figure 5.2 shows the flux predicted in [14] along with the parameterization thereof used in this analysis. The thin grey lines show the quoted theoretical uncertainty in the charm

²43 PeV for D^\pm mesons [21]





(a) ν_μ flux as a function of energy(b) ν_μ flux as a function of zenith angle(c) ν_e flux as a function of energy(d) ν_e flux as a function of zenith angle

Figure 5.1: Parameterizations of the conventional atmospheric neutrino flux calculation of [47] used as the basis for flux modeling in this analysis. In each panel the points are from [47] and the solid lines are a parameterization of the form (2.18) with the local density correction (5.4).

production cross-section, obtained by varying the input parton distribution functions, factorization scale, and assumed charm quark mass. The predicted flux of $\nu_e + \bar{\nu}_e$ and $\nu_\mu + \bar{\nu}_\mu$ are assumed to be identical due to the approximately equal branching ratios to ν_e and ν_μ in the dominant process, 3-body D meson decays. The ~ 20 times smaller flux of $\nu_\tau + \bar{\nu}_\tau$ from the decay $D_s \rightarrow \tau\nu_\tau$ (5% branching ratio [2]) is ignored.

5.1.4 Cosmic ray composition correction

Predictions of the atmospheric neutrino flux depend critically on knowledge of the primary cosmic ray flux, particularly in their extrapolations to high energies where direct measurements of the muon energy spectrum can't be used to check the calculation. The published calculations that we intend to use as templates for the atmospheric neutrino flux [14, 47], however, both assumed different primary cosmic ray fluxes, neither of which agrees particularly well with recent measurements of the cosmic ray spectrum above the knee [125]. In order to obtain templates of the atmospheric neutrino flux components based on a consistent primary cosmic ray flux model, we apply the method first presented in [36] to calculate an approximate correction as a function of neutrino energy.

There is more than one way to obtain the muon and neutrino spectra from the primary cosmic ray spectrum. Instead of representing the lepton flux due to a power law primary spectrum as in Equation (2.18), we can represent the spectrum of secondary leptons from a delta function of energy, and later convolve this with an arbitrary primary flux to obtain a lepton flux. A purely phenomenological parameterization of this “lepton yield” due to [126] is

$$N_l(> E_l, A, E, \theta) = K_l \frac{A}{E_l \cos \theta} x^{-p_1} (1-x)^{p_2} \quad , \quad (5.5)$$

where A is the atomic number of the primary, $x \equiv E_l/E_p A$ is the fraction of the primary energy per nucleon carried by the lepton, and K_l , p_1 , and p_2 are fit parameters given in Table 5.3 (p. 84). For obvious reasons we will refer to Equation (5.5) as the Elbert formula. We can convolve this yield with the primary cosmic ray flux to obtain a neutrino flux,

$$\frac{dN_l}{dE_l dA d\Omega dt} = \sum_i \int dE_p \frac{dN_{p,i}}{dE_p dA d\Omega dt} \frac{dN_l(E_l, A_i, E, \theta)}{dE_l} \quad , \quad (5.6)$$

where the sum i runs over all the chemical elements in the cosmic ray flux. It should be noted that the form (5.5) assumes that the meson re-interaction dominates over decay, so the slope of the neutrino spectrum does




Figure 5.2: Parameterization of the prompt atmospheric neutrino flux calculation of [14] used as the basis for flux modeling in this analysis. The points show the results of the baseline calculation, and the solid red line shows a 5th degree polynomial used to describe the energy dependence of the flux in log-log space. The thin grey lines parameterize the minimum and maximum flux levels obtained in [14] when theoretical inputs are varied within reasonable limits. The flux is treated as equal for ν_e and ν_μ and isotropic. The 10 times smaller contribution of ν_τ from D_s decay is neglected.

not turn over at low energies like the cascade-equation solution (2.18). Nonetheless, this integral captures the relationship between primary cosmic ray energies and the resulting neutrino energies. To estimate how much the neutrino flux would change if the full Monte Carlo flux calculation were redone with a different cosmic ray flux assumption, we evaluate Equation (5.6) with the target flux model, and then again with the flux model used in the published calculation, and take the ratio. Figure 5.3 shows the re-weighting functions used to correct [47], which assumed the knee-less flux parameterization given in [105], and [14], which assumed the parameterization of [127], where the knee is too high in energy. Recent advancements in the efficient and precise calculation of atmospheric cascades [128] should obviate the need for these sorts of corrections in future analyses.




Figure 5.3: Atmospheric neutrino flux correction used to account for the cosmic ray knee as presented in [36]. The red line is the correction from the cosmic ray flux parameterization of [105] to [125], used to adjust [47]; the blue line is the correction from [127] to [125], used to adjust [14].

5.2 Vetoing atmospheric neutrinos

Atmospheric neutrinos are produced in cosmic-ray air showers, and those same air showers also produce muons. For muon neutrinos from the northern sky³ that have been collected in past neutrino analyses [36, 116] this is of little consequence, as the bulk of the Earth completely attenuates the accompanying

³ The “north” and “south” denote halves of the celestial sphere as observed from the Geographic South Pole. The majority of atmospheric neutrinos observed from the northern sky are produced in air showers that reach ground level at points in the southern terrestrial hemisphere.




Figure 5.4: An illustration (not to scale) of how downgoing atmospheric neutrinos that interact inside the fiducial volume of IceCube can be vetoed by muons produced in the same air shower. The left panel shows an air shower where a $K^+ \rightarrow \mu^+ \nu_\mu$ decay produces a neutrino that undergoes a CC interaction inside the fiducial volume of IceCube, shown as a white rectangle. The μ^+ from the same decay triggers the outer-layer veto, causing the event to be classified as penetrating muon background. The left panel shows the analogue process for ν_e in which the partner lepton in the $K^+ \rightarrow \pi^0 e^+ \nu_e$ decay can't penetrate the > 1.5 km of ice overburden to trigger the veto. In this case, the veto can only come from muons produced in other decays in the same shower.

muon flux. Event selections like the one presented in Chapter 4, however, are sensitive to neutrinos from the southern sky as well, but explicitly reject penetrating muons using veto techniques. This means that the event selection also removes atmospheric neutrinos that are accompanied to depth by muons produced in the same air shower as illustrated in Figure 5.4. Since astrophysical neutrinos always arrive without accompanying muons, the ability to detect muon-less neutrinos from the southern sky makes it possible to classify some neutrinos as astrophysical with a much higher degree of certainty in the southern sky than in the northern sky, where they are mixed in with a truly irreducible background of atmospheric neutrinos. This simple but powerful fact was first pointed out in [129], but widely ignored until it was exploited in the analysis that provided the first robust evidence for the high-energy astrophysical neutrinos [12].

This section will outline the modifications to the published atmospheric neutrino flux models [47] and [14] necessary to convert the prediction of the total neutrino flux to a prediction of the flux that can reach IceCube with no detectable accompanying muons. First we review the analytic calculation of [129] for the fraction of neutrinos accompanied by the muon produced in the same pion or kaon decay. We then apply a different method to account for muons produced elsewhere in the same air shower to extend the self-veto prediction to ν_e and neutrinos from charmed meson decay as in [17]. Finally, we show that the modified model is a good approximation to direct Monte Carlo simulation of the entire air shower.

5.2.1 Modified cascade equations

The most obvious candidates for self-veto are muon neutrinos, where a partner muon is always produced in the same parent meson decay. In 2-body decays, the energies of the muon and neutrino are directly correlated. If we return to Equation (2.16), we can rewrite the lower limit of the integral for meson species i as

$$\begin{aligned} z_i &\equiv \frac{E_M}{E_\nu} = \frac{E_\mu + E_\nu}{E_\nu} \\ &= 1 + \frac{E_\mu}{E_\nu} \quad , \end{aligned} \tag{5.7}$$

where E_M is the energy of the parent meson, E_ν is the energy at which the neutrino flux is being evaluated, and E_μ is the energy of the associated muon. This means that increasing values z correspond to increasingly large muon energies. To evaluate the flux of neutrinos where the partner muon is under some energy threshold $E_{\mu,\min}$, we divide the integral into two parts: the first, which we directly evaluate, covers the part of the

phase space where the muon is guaranteed to be above the energy threshold, while the remainder covers the part of the phase space where the muon is below threshold. It is straightforward to obtain this remainder by subtraction from the integral over the full phase space. We replace the lower limit of the integral for meson species i with

$$z_{i,\min} \equiv \max \left\{ \frac{1}{1-r_i}, 1 + \frac{E_{\mu,\min}}{E_\nu} \right\} , \quad (5.8)$$

where again $r_i \equiv m_\mu^2/m_i^2$, removing the part of the decay phase space where $E_\mu \leq E_{\mu,\min}$. This leads to a new set of coefficients for Equation (2.18) (cf. Equation (2.20)) [129],

$$A_{i\nu}(E_{\mu,\min}) \equiv \frac{1}{(z_{i,\min})^{\gamma+1}} \left(\frac{Z_{Ni}}{1-r_i} \right) \left(\frac{1}{\gamma+1} \right) \quad \text{and} \quad (5.9)$$

$$B_{i\nu}(E_{\mu,\min}) \equiv z_{i,\min} \left(\frac{\gamma+1}{\gamma+2} \right) \left(\frac{\Lambda_i - \Lambda_N}{\Lambda_i \ln(\Lambda_i/\Lambda_N)} \right) . \quad (5.10)$$

These coefficients are now no longer constants, but instead functions of both $E_{\mu,\min}$ and E_ν through their dependence on $z_{i,\min}$. The unaccompanied flux is then just

$$\frac{d\Phi_{\nu,\text{unvetoed}}}{dE_\nu} = \frac{d\Phi_\nu}{dE_\nu}(E_{\mu,\min} = 0) - \frac{d\Phi_\nu}{dE_\nu}(E_{\mu,\min}) , \quad (5.11)$$

where the first term is evaluated using the coefficients of (2.20) and the second is evaluated using the coefficients of (5.10).

It now remains to determine an appropriate value for the minimum muon energy at production $E_{\mu,\min}$. In [129] this was taken to be the minimum surface energy required for a muon to penetrate an overburden X average,

$$E_{\mu,\min}(X) = 0.73 \text{ TeV} \times [\exp(X/2.8 \text{ kmwe}) - 1] ,$$

taken from the parameterization of the mean muon range found in [72]. This is a necessary but not sufficient condition for an event being vetoed, as the detection process has an energy threshold. To account for this threshold, we require that the muon not only penetrate the given overburden, but also survive with an energy greater than the threshold 50% of the time. This was determined with Monte Carlo simulation by propagating a monoenergetic beam of muons through a slab of ice with **PRopagator with Optimal Precision and Optimized Speed for All Leptons (PROPOSAL)** and adjusting the initial energy until the median final energy matched the desired value. The points in Figure 5.5a show these surface energies as a function of

overburden for various choices of the median final energy. The solid lines are a parameterization given by

$$\log_{10} E_{\mu,\min}(X, E_{\mu,\text{med}}) = a + \frac{b}{10^4} X + \frac{c}{10^{10}} X^2 \quad , \quad (5.12)$$

X is the overburden in meters, $E_{\mu,\text{med}}$ is the required median energy at depth, and a , b , and c are polynomials in $\log_{10}(E_{\mu,\text{med}})$ of the form

$$p_0 + p_1 \log_{10}(E_{\mu,\text{med}}) + p_2 \log_{10}(E_{\mu,\text{med}})^2 \quad . \quad (5.13)$$

The parameters of Equation (5.13) are given in Table 5.2. Figure 5.5b shows the slant depth of the ice as a function of zenith angle for the range of vertical depths relevant for IceCube.

Table 5.2: Parameters of Equation (5.13).

Parameter	p_0	p_1	p_2
a	2.793	-0.476	0.187
b	2.069	-0.201	0.023
c	-2.689	3.882	0


Figure 5.6a shows Equation (5.11) evaluated as a function of neutrino energy with a surface energy threshold corresponding to 1 TeV median energy at depth according to Equation (5.13). In order to apply this as a correction to [47], we take the ratio of the muon-less flux shown in solid lines to the total flux shown in dotted lines. This passing fraction is shown in Figure 5.6b. A kink appears in each of these curves when the minimum possible energy of the partner muon from pion decay exceeds the detection threshold, suppressing the pion component completely and leaving only neutrinos from kaon decay. This occurs at (cf. Equation (2.13))

$$E_\nu = \left(\frac{1}{r_\pi} - 1 \right) E_{\mu,\text{med}} = 0.76 \times E_{\mu,\text{med}} \quad . \quad (5.14)$$


The same mechanism suppresses the flux from 2-body kaon decay completely above $21.1 \times E_{\mu,\text{med}}$.

5.2.2 Generalized atmospheric neutrino self-veto

Up to now we have only treated self-veto by the direct partner muon from the same 2-body meson decay. This is unsatisfactory for two related reasons. First, it does not treat ν_e , where the partner lepton is an electron that can't penetrate the ice overburden rather than a muon. In such cases the vetoing muon can only come from a different decay vertex in the same shower. While ν_e are sub-dominant in the conventional




(a) Surface energy required for a muon to survive with various median energies as a function of slant depth. The solid lines are a parameterization of the data points, which were obtained from simulations with PROPOSAL [73].




(b) Slant depth as a function of zenith angle for vertical depths relevant for IceCube.

Figure 5.5: Quantities required to calculate the minimum surface energy required for a muon to reach IceCube from various zenith angles.



(a) Conventional neutrino fluxes evaluated without regard to accompanying muon energy (dotted lines) and with a maximum accompanying muon energy at 1950 m depth of 1 TeV (solid lines).



(b) Passing fraction, taken as the ratio of the solid to dotted lines in (a). A kink appears in each curve when the pion component becomes completely suppressed, leaving only the kaon component.

Figure 5.6: Conventional neutrino fluxes evaluated according to Equation 5.11. The constants in (2.20) and (5.10) were taken from [21]. The surface energies corresponding to a median energy at 1950 m depth of 1 TeV were 2.4, 5.2, and 57.2 TeV at 0, 60, and 78 degrees zenith, respectively.

atmospheric neutrino flux, they provide fully half of the prompt atmospheric neutrino flux. Furthermore, the treatment above applies only to 2-body decay, and it is not immediately obvious that it should describe the fraction of unaccompanied ν_μ from the 3-body D meson decays that provide the bulk of the prompt ν_μ flux. This section addresses both of these weaknesses, first by showing how to calculate the number of muons from other parts of the same air shower and then by showing that the analytic treatment of the previous section also applies approximately to neutrinos from 3-body decays.

First we turn to the problem of calculating the probability that a neutrino will be accompanied by a muon from a different part of the same air shower, following the treatment in [17]. While the energies of the neutrino and muon are no longer directly correlated, they are still related through the total energy of the air shower that produced both. We can approach this problem with a technique very similar to the one used to account for the cosmic ray knee in Section 5.1.4, using the lepton yield (cf. Equation (5.5)) to connect neutrino and muon energies to the energies of primaries that produced them.

One can write a response function that [17]

gives the distribution of primary energy of nuclei of mass A that produce leptons of a given energy E_ℓ as

$$R_\ell(A, E, E_\ell, \theta) = \phi_N(A, E) \times \frac{dN_\ell(> E_\ell, A, E, \theta)}{dE_\ell} . \quad (5.15)$$

Then the flux of leptons is

$$\phi_\ell(E_\ell, \theta) = \Sigma_A \int dE R_\ell(A, E, E_\ell, \theta) . \quad (5.16)$$

To estimate the passing rate of neutrinos we evaluate

$$P_\nu(E_\nu, \theta) = \frac{\Sigma_A \int dE R_\nu P(N_\mu = 0)}{\Sigma_A \int dE R_\nu} , \quad (5.17)$$

where $P(N_\mu = 0|A, E, E_{\mu, \text{med}}, \theta)$ is the probability that no muons from a shower initiated by a cosmic ray of the given mass, energy, and zenith angle penetrate to the depth of the detector without dropping below the detection threshold $E_{\mu, \text{med}}$. This can be approximated as the Poisson probability

$$P(N_\mu = 0|A, E, E_{\mu, \text{med}}, \theta) = e^{-N_\mu(A, E, E_{\mu, \text{med}}, \theta)} , \quad (5.18)$$

where $E_{\mu,\min}(\theta)$ is the surface energy required to reach the detector with $E_{\mu,\text{med}}$ 50% of the time, parameterized according to Equation (5.13), and N_μ is the cumulative muon yield evaluated at that energy. The central idea of the estimate is to weight the probability of zero muons according to the weights that give rise to the flux of neutrinos of a given E_ν .

It remains to find a parameterization of the yield $dN_l(> E_l, A, E, \theta)/dE_l$ for the processes we have to consider: μ and ν_μ production in 2-body pion and kaon decay, ν_e from 3-body kaon decay, and ν_μ and ν_e production in 3-body charmed meson decay. This was done by simulating proton-initiated showers with **COsmic Ray Simulations for KAscade (CORSIKA)** at fixed energies for various zenith angles and tabulating the energies of penetrating secondary leptons in terms of $x \equiv E_l/E_p$, separated by the flavor of the lepton and the flavor of the heaviest quark in the parent hadron. In order to obtain a sufficient number of high-energy leptons, **CORSIKA** was modified to shorten the mean lifetimes of kaons and charged pions such that the probability of decay in flight was at least 50%. A weight was calculated for each meson to account for the bias towards decay over interaction⁴. Figure 5.7 shows the secondary lepton energy spectra obtained from this simulation, while Figure 5.8b shows an example of the zenith dependence of the yield. The conventional yields can be well described by a slightly modified form of the original Elbert parameterization given by

$$N_l(> E_l, A, E, \theta) = K_l \frac{A}{E_l \cos \theta^*} x^{-p_1} (1 - x^{p_3})^{p_2} \quad , \quad (5.19)$$

where again $x \equiv E_l/E_p$ and $\cos \theta^*$ is a local density correction (cf. Equation (5.4)). The best fit parameters for μ , ν_μ , and ν_e are given in Table 5.3. The decay probability of the parent meson is proportional to $1/E_l \cos^* \theta$. The same form can be made to describe leptons from the decays of charmed mesons like the D^\pm that decay promptly before they can re-interact by removing the decay-probability factor:

$$N_l(> E_l, A, E, \theta) = K_l A x^{-p_1} (1 - x^{p_3})^{p_2} \quad . \quad (5.20)$$

Figure 5.8a shows these parameterizations evaluated at a single primary energy and zenith angle.

⁴This weighting scheme is feasible for a calculation like this one, where only the properties of individual leptons are of interest. In calculations that involve the ensemble of particles in each shower, like the full Monte Carlo calculation of the unaccompanied neutrino flux that we will see shortly, this is no longer true. If we forced all mesons to decay, then we would generate only very atypical showers with vanishingly small weights, and never encounter the average case we are interested in. If, on the other hand, we anoint only one meson to decay, we would need to be able to evaluate the probability of choosing that specific meson from the population of all the mesons that could have been produced in the shower in order to calculate the correct weight. This is an altogether more difficult problem.

Table 5.3: Parameters of the modified Elbert formula for different lepton flavors and production processes.

Parameterization	K	p_1	p_2	p_3	Equation
Elbert μ	14.5	0.757	5.25	—	(5.5)
Conventional μ	49.5	0.626	4.94	0.580	(5.19)
Conventional ν_μ	79.9	0.463	4.37	0.316	(5.19)
Conventional ν_e	0.548	0.669	8.05	0.722	(5.19)
Charm ν_μ and ν_e	0.000780	0.604	7.34	0.767	(5.20)






(a) Conventional μ from SIBYLL(b) Conventional ν_μ from SIBYLL(c) Conventional ν_e from SIBYLL(d) Prompt ν_μ from DPMJET

Figure 5.7: Energy spectra for various lepton flavors and production processes from **CORSIKA** simulations of vertical, proton-initiated showers with hadronic interactions from SIBYLL 2.1 [130] and DPMJET 2.55 [131]. The solid lines are the parameterizations of Equation (5.19) for conventional leptons or Equation (5.20) for leptons from charmed particles, evaluated using the parameters in Table 5.3.



(a) Cumulative lepton yields for vertical, 1 PeV proton showers.


(b) Zenith angle dependence of the cumulative conventional ν_e yield evaluated at 10% of the primary energy for different primary energies.

Figure 5.8: Lepton yields from the modified Elbert formula. Each curve in (a) shows Equation (5.19) evaluated with one of the parameter sets from Table 5.3. The points in (b) show yields from CORSIKA simulation, the dotted lines a $1/\cos\theta$ dependence, and the solid lines a $1/\cos\theta^*$ dependence (cf. Equation 5.4). All conventional lepton yields have the same zenith dependence.

We can now describe the probability that atmospheric neutrinos from any process will be accompanied by muons from other parts of the same air shower as well as the probability that a ν_μ will be accompanied by a muon from the same 2-body pion or kaon decay. What is missing for a complete description of conventional and prompt atmospheric neutrinos is the probability that a ν_μ produced in a 3-body charmed meson decay will be accompanied by a muon from the same vertex. In these decays, e.g. $D^+ \rightarrow \bar{K}^0 \mu^+ \nu_\mu$, the lepton pair shares energy with a kaon. The most obvious consequence of this is that there is no longer a 1-to-1 correspondence between the lab-frame energies μ and ν energies. The relevant property of the decay is however the distribution of E_μ/E_ν , particularly how often $E_\mu \ll E_\nu$, making the neutrino unvetoable. This is shown in Figure 5.9 for 2-body π and K decays as well as 3-body D decays. Even though the distribution of E_μ/E_ν extends to $-\infty$ in the 3-body D decay, the fraction of the distribution where $E_\mu < E_\nu$ is smaller than for 2-body K decays. Given this fact it is reasonable to expect that the simple analytic form of the unvetoable fraction derived for 2-body π and K decays derived in the previous section will provide a reasonable description of the unvetoable fraction of ν_μ from charmed meson decays. This is in fact the case, as we will see shortly see verified with a full Monte Carlo calculation of the passing fraction.




Figure 5.9: Cumulative distribution of muon lab-frame energies from various ultrarelativistic meson decays in units of the energy of the neutrino from the same decay. The fraction of unvetoable neutrinos is governed by the fraction of the cumulative distribution at small values of E_μ/E_ν . The lab-frame energies of the D daughter particles were sampled according to the phase space density of 3-body decays according to the Raubold-Lynch algorithm [132]. For 2-body decays the sampling is trivial, and was done directly.

5.2.3 Verification with full air shower simulation

It is important to verify that the approximations made in the previous section accurately reflect the distributions of muons and neutrinos in realistic air showers. We do this by simulating air showers with **CORSIKA**, propagating the muons through an overburden of ice corresponding to their zenith angle (see Figure 5.5b), and tabulating, for neutrino in the shower, the neutrino energy, flavor, parent meson type, and the number of muons with more than 1 TeV at depth. The primary cosmic rays were then weighted to the spectrum of [125], and the veto passing fraction calculated for each zenith angle, neutrino type, energy, and parent meson type as the ratio of the flux accompanied by zero muons to the total flux. We then compare this to the approximation derived in the previous sections. Since ν_μ may be vetoed either by a muon from the same vertex or from the rest of the shower, we approximate the passing rate as

$$P_{\text{total}} \approx P_{\text{correlated}} \times P_{\text{uncorrelated}} \quad , \quad (5.21)$$

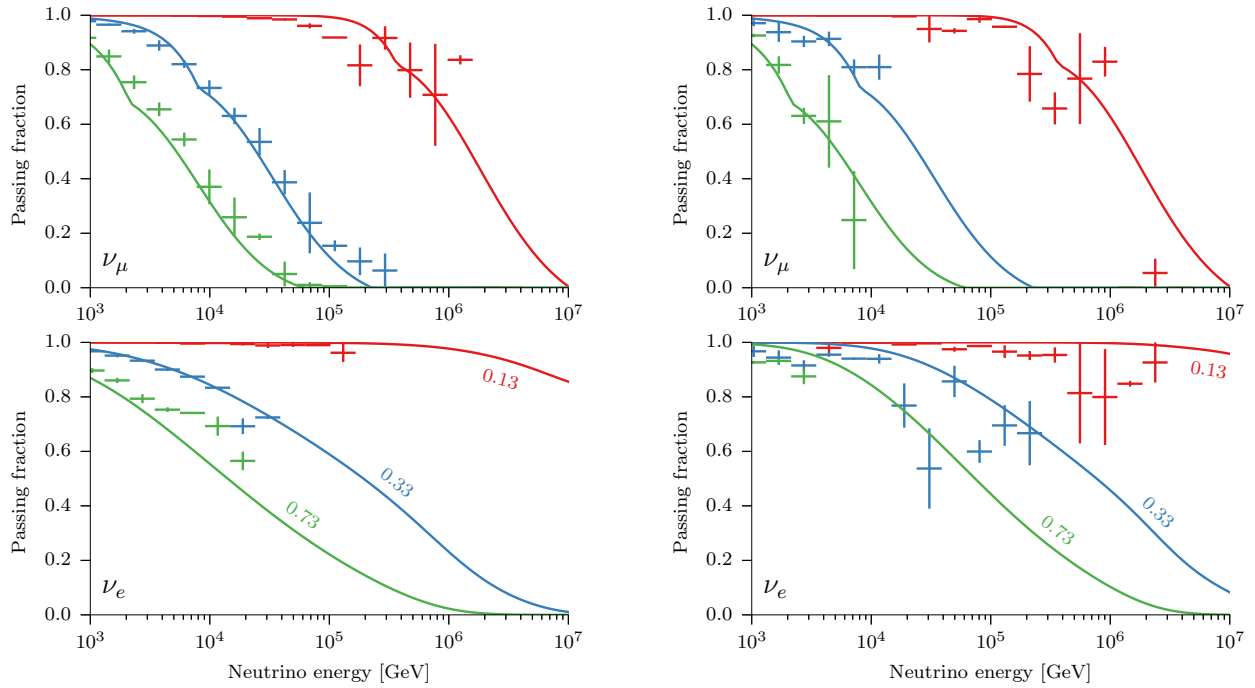
where the first factor is the passing rate from Equation (5.11) [129] and the second is from Equation (5.17) [17]. While this approximation accounts for the correlated muon more than once, it nonetheless describes the full Monte Carlo calculation quite well. For ν_e there is no partner muon, and the passing rate is described well by Equation (5.17) alone.

Figure 5.10 shows a comparison of the passing rates calculated from the full Monte Carlo simulation with the approximations developed in the previous sections.

5.2.4 Verification with full detector simulation

The one remaining assumption that must be verified before the approximations developed here can be used for real data analysis is that the effect of the veto cuts presented in Chapter 4 can be approximated by a simple muon energy threshold above which all atmospheric neutrinos are removed from the sample. This was demonstrated by applying the full detector simulation and event selection to air showers simulated with **CORSIKA** where one neutrino was forced to interact within the fiducial volume of IceCube. As usual, the only subtlety in this process is calculating correct weights. In this simulation, the expression for the weight is

$$w_i = w_{\text{interaction}} w_{\text{selection}} \frac{\frac{dN(E_i, A_i)_{\text{nature}}}{dEdAd\Omega dt}}{\frac{dN(E_i, A_i)_{\text{generated}}}{dEdAd\Omega}} \quad , \quad (5.22)$$



(a) Conventional neutrinos from hadronic interactions simulated with SIBYLL. Top panel: ν_μ with solid lines showing the passing rate from (5.21). Bottom panel: ν_e with solid lines showing the passing rate from (5.17).

(b) Prompt neutrinos from hadronic interactions simulated with DPMJET. Top panel: ν_μ with solid lines showing the passing rate from (5.21). While this calculation only applies strictly to 2-body decays of pions and kaons, it provides an adequate description of the muon/neutrino correlation in 3-body decays of D mesons as well. Bottom panel: ν_e with solid lines showing the passing rate from (5.17).

Figure 5.10: Comparison of approximate self-veto passing rates (solid lines) with Monte Carlo (crosses) for atmospheric neutrinos at 3 values of $\cos \theta$.

where $w_{\text{interaction}}$ is the same forced-interaction weight defined in Equation 5.2, but the flux weights are now a function of the primary cosmic ray flux model rather than the neutrino flux model, and an extra term $w_{\text{selection}}$ has appeared to compensate for the fact that only one neutrino has been selected out of all neutrinos in the bundle. If we selected a neutrino at random, then the selection weight would simply be the number of available neutrinos N_ν , accounting for the fact that one neutrino is representative of the larger population from which it was drawn. Selecting neutrinos this way would be quite inefficient, though, since most of the neutrinos are found at low energies as can be seen from Figure 5.7. Instead, we select with a probability proportional to the energy of the neutrino,

$$p_{\text{biased}} = \frac{E_i}{\sum_j E_j} \quad . \quad (5.23)$$

In this case the selection weight is

$$w_{\text{selection}} = N_\nu \frac{p_{\text{natural}}}{p_{\text{biased}}} = \frac{1}{p_{\text{biased}}} \quad . \quad (5.24)$$

We then weight the single-neutrino simulation by (5.3), multiplied by an additional factor (5.21) to approximately account for the effect of the veto cuts on the unsimulated accompanying muons.

Figure 5.11 shows the zenith distributions of conventional ν_μ and ν_e above 10 TeV for the single-neutrino simulation and the full CORSIKA atmospheric neutrino simulation. For this event selection, the full simulation can be approximated by weighting the single-neutrino simulation by an additional factor (5.21) using a muon energy threshold of 100 GeV. It should not be a surprise that the energy threshold for vetoing atmospheric neutrinos is a factor of 10 below the energy threshold of the overall event selection. The veto was designed to reject millions of penetrating atmospheric muons for every neutrino it accepts. With this large background, the veto must be able to reject muons that have experienced quite rare under-fluctuations in their energy deposition rate. Neutrino events, however are different: since the events pass the veto because of the rare interaction of the neutrino itself, the starting population is small. Only the average properties of the accompanying muons are relevant, and the average 100 GeV muon can be rejected at least 50% of the time.

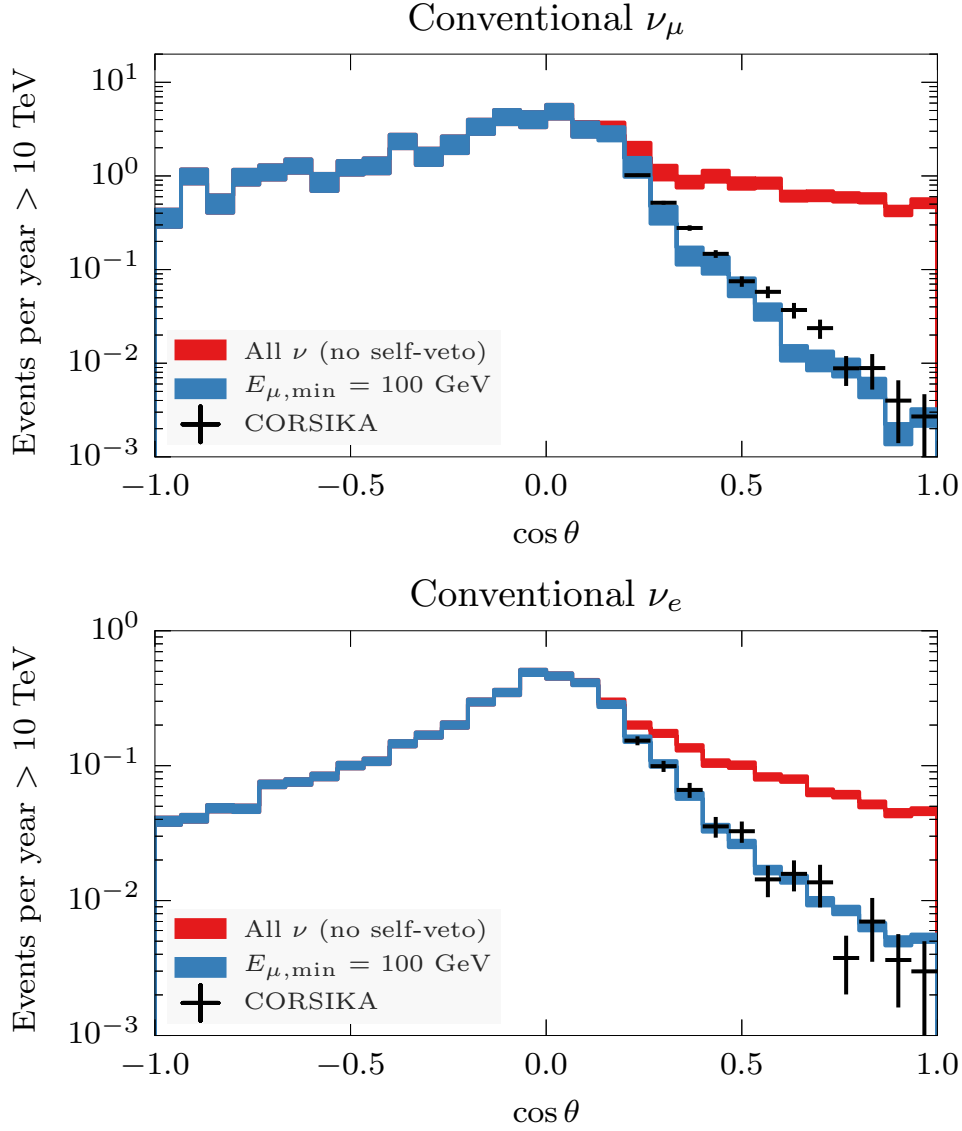


Figure 5.11: Verification of the approximate veto probability calculation against the full detector simulation. The colored bands show the true zenith angle distribution for all conventional atmospheric neutrinos with energies greater than 10 TeV predicted from single-neutrino simulation. The red band is weighted to the flux of [47], while the blue band is weighted to [47] modified to account for self-veto with $E_{\mu,\text{med}} = 100$ GeV. The width of the bands gives the statistical error on the contents of each bin. The black crosses are a full **CORSIKA** air-shower simulation with SIBYLL 2.1 hadronic interactions where one neutrino from each shower was forced to interact in the fiducial volume of IceCube. The showers were weighted to the primary flux model of [125]. The effect of the veto cuts on the down-going atmospheric neutrino flux can be approximated reasonably by $E_{\mu,\text{med}} = 100$ GeV.

5.3 Modeling the astrophysical neutrino flux

Overlaid on the atmospheric components is a flux of high-energy neutrinos of astrophysical origin. Their energy distribution is harder than those of any of the other sources of neutrinos, and we found that they are the dominant source of events with more than 100 TeV deposited energy [12, 13]. The Earth absorbs a significant fraction of upward-going neutrinos above 100 TeV [63], so the highest-energy of these are concentrated around the horizon and in the southern sky. Since the sources of these neutrinos are unknown, the shape of their energy and angular distribution cannot be predicted exactly, and given the limited number of neutrino events that can be detected, only very simple models can be tested. Neutrinos associated with the extragalactic sources of the highest-energy cosmic rays are assumed to be isotropically distributed and follow a power law energy distribution of approximately E^{-2} [8] and arrive at the Earth as equal parts ν_e , ν_μ , and ν_τ due to oscillations [49]. We parameterize the diffuse astrophysical neutrino flux as

$$\Phi_{\text{astro}} = \Phi_0 \left(\frac{E}{E_0} \right)^{-\gamma}, \quad (5.25)$$

where Φ_0 is the $\nu + \bar{\nu}$ flux for each flavor at $E_0 = 10^5$ GeV in units of $\text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$ and γ is the spectral index. More generally, γ can be allowed to vary to account for the spectra expected from specific classes of sources, for example TeV photon emission from active galactic nuclei ($2.2 \lesssim \gamma \lesssim 2.6$) [6] or interactions of cosmic rays with dense gas clouds while magnetically confined in starburst galaxies ($2.0 \lesssim \gamma \lesssim 2.25$) [133].

5.4 Estimating the background from penetrating atmospheric muons

One of the more difficult parts of an analysis is obtaining a robust estimate of the background from penetrating atmospheric muons. In rare cases penetrating muons can produce event signatures that are indistinguishable from neutrino events. Even though these events are rare they can contribute significant numbers of events to the final sample because penetrating muons outnumber interacting neutrinos 100,000:1, and any realizable neutrino selection will have some contamination of penetrating atmospheric muons. For a selection like this one, the contamination must be estimated from simulation.

In IceCube this is done by simulating cosmic-ray air showers with **CORSIKA** down to the surface of the glacier (2831 meters above sea level), then propagating the muons through the firn and ice to the detector with **Muon Monte Carlo (MMC)** and running the detector simulation. These simulated penetrating muon

events are then run through the same selection process as the experimental data. It is often convenient to draw the energies of the primary particles that initiate the air showers from a power law spectrum, and then weight the simulated spectrum to a realistic spectrum after the fact. The weight for each event i is simply the ratio of the number of nucleons of the given species A and energy E expected per unit energy, area, solid angle, and time to the number generated in the simulation per unit energy, area, and solid angle:

$$w_i = \frac{\frac{dN(E_i, A_i)_{\text{nature}}}{dE dA d\Omega dt}}{\frac{dN(E_i, A_i)_{\text{generated}}}{dE dA d\Omega}} \quad . \quad (5.26)$$

All the units cancel except for an extra factor of $1/t$, so this is a fractional contribution to a rate. To get the predicted background event rate under a particular selection, one simply sums the weights of all events that survive the selection.

The immediate problem that arises is that when no events survive the selection, the estimated rate is identically zero. This could happen because the rate is actually zero, or because not enough events have been simulated that one would pass the selection. Here it is useful to consider the concept of an effective livetime. There are various ways to define this, but the simplest definition is that the effective livetime is the experimental observation time it would take to accumulate the number of events that were simulated. The effective livetime as a function of primary energy and species is then just the inverse of the weighting function:

$$T_{\text{eff}}(E, A) = \frac{1}{w(E, A)} \quad . \quad (5.27)$$

The estimated passing rate is only robust if the effective livetime is of the same order or greater than the observation time, or the period of the expected signal, whichever is smaller. Since in this scheme it is a function of both primary energy and species rather than experimental observables, it is necessary to simulate sufficient livetime everywhere in the range of primaries that could contribute to the range of observables that would pass the event selection.

Figure 5.12 shows the effective livetime of the **CORSIKA** simulations used to develop the event selection as a function of energy per nucleon. The energy per nucleon is a useful quantity because it provides a strict upper bound on the energy of the of the most energetic secondary particle that can be produced in the shower. For example, a shower initiated by a 10 TeV proton or a 560 TeV iron nucleus can contain no muons or neutrinos with energies greater than 10 TeV, which in turn means that no more than 10 TeV can be

deposited in the detector⁵. The largest contribution to the effective livetime comes from the dedicated high-statistics simulations for the cascade detection channel that had the offline pre-selection applied before being written to disk, but these were only able to reach 1 year of effective livetime by simulating only primaries with more than 30 TeV/nucleon. This is problematic for this analysis, since the background statistics for events depositing between 1 and 30 TeV are vastly under-simulated. To further complicate things, these dedicated simulations had not yet been produced for IC86 at the time the selection was developed, and as of the time of writing are still not complete.

The simulations shown in Figure 5.12 already took 400 CPU-years to produce, so simply simulating 10 times more showers was not a viable option. The particular qualities of the event selection presented in Chapter 4, however, make a full air shower simulation unnecessary. For the purposes of background estimation, the function of the CORSIKA air shower simulation is to convert a distribution of cosmic ray primaries into a distribution of muon bundles at the depth of the detector. With potentially thousands of muons per bundle, the problem has too many dimensions to handle analytically, and so is treated by Monte Carlo. These large muon bundles, however, are not a relevant background for a veto-based selection. Figure 5.13 shows the cumulative distributions of muon bundle multiplicity at the border of the instrumented volume⁶ after each of the cuts presented in Section 4.2. Each of the veto cuts preferentially selects single-muon events, while rejecting high-multiplicity bundles with increasing efficiency. The reason for this is quite simple. Penetrating muons can appear as neutrino-like events if their energy loss rate (and thus Cherenkov photon rate) fluctuates under the detection threshold while passing near DOMs that could participate in the veto before fluctuating upwards to produce a large stochastic loss that mimics the cascade at the vertex of a neutrino interaction. In order for a bundle to exhibit this behavior, all muons in the bundle would have to experience a correlated under-fluctuation in their energy loss rates, an occurrence that is less and less likely the larger the number of muons in the bundle becomes.

The overwhelming dominance of single muons is convenient, as it reduces the dimensionality of the problem enough that it becomes possible to de-couple the air shower simulation from the detector simulation when estimating the background without losing any information. Inspired by the previous work of [134], we parameterize the distribution of single muons and bundles at depth, and use this parameterization to

⁵In the general case, it is the total energy of all muons that determines how much can be deposited in the detector. However, as we will see shortly, for a veto-based selection like this one the background is dominated by events where only a single muon survives to the detector.

⁶Approximated as an upright cylinder with a full height of 1000 m and a radius of 500 m, centered on the origin of the IceCube coordinate system.




Figure 5.12: Effective livetime of centrally-produced CORSIKA penetrating muon background simulations for each year of data-taking as a function of energy per nucleon. The sudden uptick at 30 TeV/nucleon in the IC79 is the dedicated high-statistics simulation that was produced with the offline pre-selection applied. For IC86 no such dedicated simulation existed at the time the analysis was performed, and the livetime at low energy/nucleon was only a few days.




Figure 5.13: Number of muons that reach the detector border (derived from simulation) after various cuts. After each cut the background becomes more and more dominated by single muons over bundles.

inject single muons at the depth of the detector that can later be weighted to account for the frequency with which they would occur assuming some primary composition and hadronic interaction model. This parameterization and its associated sampling and weighting methods are presented in detail in Appendix A. Figure 5.14 shows the single-muon flux assumed for the background estimate at 2 km vertical depth (the center of IceCube) at a zenith angle of 60 degrees. The largest contribution to the muon flux comes from the decays of relatively long-lived π and K mesons. Because their mean lifetimes in the lab frame increase with energy, their decays become increasingly rare at high energies, steepening the resulting muon spectrum [23]. The spectrum of muons from the decays of short-lived, charmed mesons, however, steepens at a much higher energy, so these could potentially provide the dominant source of single muons at sufficiently high energies. In order to obtain a conservative upper bound on the background for penetrating single muons, we take the single-muon flux as the sum of the conventional π/K and prompt charmed-meson fluxes from CORSIKA simulations with two different hadronic models. The conventional component is treated with SIBYLL 2.1 [130], which only treats u , d , and s quarks⁷, but also predicts the largest overall muon flux [136]. The prompt component is treated with DPMJET 2.55 [131], which predicts a prompt component slightly larger than the previous upper limit on atmospheric charm production derived from neutrino measurements [36]. With this relatively large normalization, the prompt muon flux exceeds the conventional one at a few tens of TeV as shown in Figure 5.14.

Figure 5.15 shows the total effective livetime of all penetrating muon simulations used for the background estimate in this analysis. These were cobbled together out of the CORSIKA simulations reflected in Figure 5.12 and MuonGun simulations from three different sources, shown in Table 5.4. The first was a by-product of background studies for [12], while the second was produced during the development of a separate, IC86-only analysis focused on measurement of the conventional atmospheric ν_e flux as a follow-up to [35]. Both of these used a fixed target surface, which made them generally applicable but in the case of the second, severely limited the effective livetime at low energies. The third set was made specifically for this analysis, using a combination of spectrum and an energy-dependent target surface sufficient to fill in the deficit in livetime apparent in the IC79 CORSIKA simulations below 30/TeV nucleon in Figure 5.12. The target surface scales like Equation (4.3), but with $q = E_\mu/4$. This means that the livetime is a function of position in the detector as well as energy, and at low energies drops off dramatically outside the fiducial volume. Since low-energy muons that do not pass through the fiducial volume can't pass the selection by

⁷There is ongoing work to include charm production in a new version of SIBYLL tuned to the latest LHC measurements [135], but as of the time of writing this was not yet finished.




Figure 5.14: Underground muon flux models used for the background estimate. The conventional component is taken from a fit to CORSIKA simulation with SIBYLL 2.1 hadronic interactions [130] and the prompt component from DPMJET 2.55 [131].

Table 5.4: MuonGun simulations used in this analysis

Injected events	Energy		Target surface	Running time
	Spectrum	Range [TeV]		
150×10^6	E^{-1}	10^4-10^7	Fixed	30 CPU-years
270×10^6	E^{-2}	$1-10^5$	Fixed	2.5 GPU-years
500×10^6	E^{-5}	$2-10^4$	Scaled	5 GPU-years

construction, however, this optimization does not degrade the effective livetime of the background estimate in the final sample.




Figure 5.15: Effective livetime of all penetrating muon background simulations used in this analysis for each year of data-taking as a function of lead muon energy at the detector border, evaluated at a point inside the fiducial volume defined by Equation (4.3).

5.5 Observables

5.5.1 Deposited energy

Since the event sample consists of cascades and cascade-dominated starting track events, we use the energy deposited at the neutrino vertex in the initial cascade as reconstructed with **Monopod** as a proxy for the neutrino energy. Figure 5.16 shows the distribution of relative energy reconstruction error for all simulated conventional atmospheric neutrino events that would survive to the final sample, defined as

$$\Delta \log_{10} E \equiv \log_{10} \left(\frac{E_{\text{reco}}}{E_{\text{vis}}} \right) \quad , \quad (5.28)$$

where E_{reco} is the reconstructed energy and E_{vis} is the true deposited electromagnetic-equivalent energy⁸. We define the bias and resolution of the energy reconstruction as the mean and standard deviation, respectively, of a Gaussian fit to the distribution of $\Delta \log_{10} E$. These quantities have an energy dependence, as shown in Figure 5.17. Since the reconstruction error is dominated by statistical fluctuations in the number of collected **PE**, the resolution improves as the deposited energy increases. Some care is required when comparing these resolutions to those presented in [16], however. Whereas the resolutions in [16] were obtained from nearly-unbiased **online filter**-level simulation samples, the neutrino-level sample has a strong selection bias. The charge-dependent fiducial volume cut (cf. Section 4.2.4) ensures that the lowest-energy events are concentrated in **DeepCore**, where the instrumentation is denser, the quantum efficiency of the **PMTs** higher, and the average optical attenuation length of the ice longer. These effects combine to increase the average number of **PE** collected per unit deposited energy, shrinking the statistical error on the reconstructed energy considerably at low energies.

5.5.2 Neutrino flavor discrimination

As already mentioned in Section 4.2.5, we can invert the incoming-track veto presented in Section 4.2.3 to detect tracks emerging from the neutrino interaction vertex. After the neutrino-level event selection we repeat the outgoing track search with all rather than just up-going directions. Any event where the total charge of **HLC pulses** with $\Delta t_{\text{cascade}} < -50$ ns and $-15 \leq \Delta t_{\text{track}} < 1000$ ns is greater than 10 **PE** is considered a starting track event, and all others cascades. Table 5.5 shows how neutrino interactions of

⁸Since the energy bins in the final analysis will be logarithmically spaced, we choose to take the logarithm in order to easily be able to compare the energy resolution to the width of the energy bins.


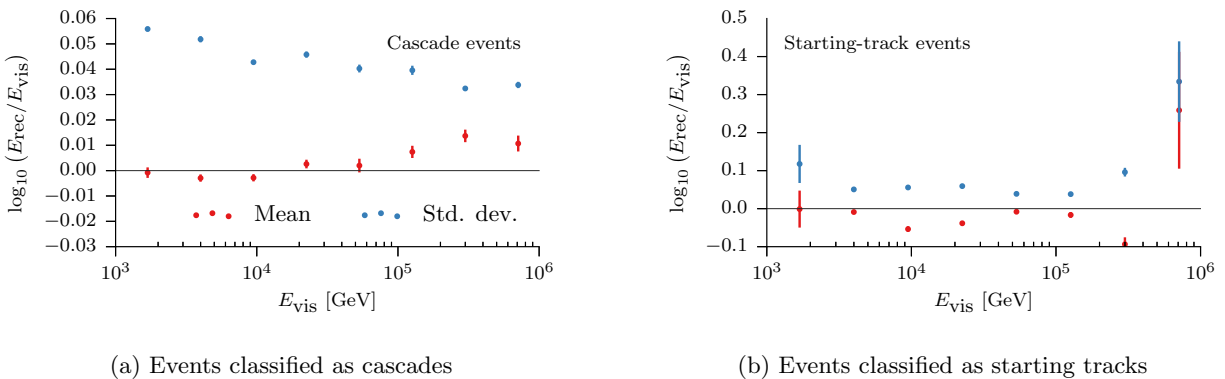


Figure 5.16: Logarithmic error on reconstructed energy for simulated neutrino events that are classified as cascades (cf. Section 5.5.2). The reference energy E_{vis} is the deposited electromagnetic-equivalent energy. The thin blue line shows a Gaussian fit to the distribution.



(a) Events classified as cascades

(b) Events classified as starting tracks

Figure 5.17: Mean and standard deviation of the relative energy reconstruction error as a function of deposited energy, determined with a Gaussian fit as shown in Figure 5.16. For cascade events the energy estimation is nearly unbiased, with a relative error that decreases with increasing energy. For track events the relative error increases slightly with increasing energy, and the cascade energy reconstruction tends to underestimate the total deposited energy in the event.


various flavors are classified. While very few true cascades are mis-classified as tracks, not all tracks can be positively identified. This is expected, especially for interactions near the edge of the detector, as the outgoing muon can leave the detector before it out-strips the slower-moving light front from the initial cascade by significantly more than the required 50 ns. The successful **CC** ν_μ classification rate increases for harder spectra, though this is accompanied by a slightly higher mis-classification rate for cascade events.

5.5.3 Zenith angle


The zenith angle is taken from the **Monopod** reconstruction for events that are classified as cascades. For starting track events, the zenith angle is taken from the best-fit outgoing track, as the large displacement of the associated photons from the neutrino interaction vertex provides a better constraint than the initial cascade. This can be seen in the distribution of zenith angle errors shown in Figure 5.18. Like the energy reconstruction, the energy dependence of the zenith angle reconstruction is strongly affected by selection bias, as shown in Figure 5.19. Rather than degrading at low energies as shown in [16], the median angular error is nearly constant, as the denser instrumentation and longer scattering length in the core of the detector make the Cherenkov cone of the cascade easier to resolve, and tracks leaving the detector easier to pick out.

Table 5.5: Event type classification rates derived from MC simulation for various assumed neutrino fluxes, broken down by neutrino flavor and interaction channel. **CC** ν_e and **NC** interactions always produce true cascades, and **CC** ν_μ interactions always produce true starting tracks. **CC** ν_τ events may be either track- or cascade-like, depending on the decay length of the τ lepton, which on average increases with energy. **Glashow resonance (GR)** $\bar{\nu}_e$ interactions also produce a mixture of track- and cascade-like events, as the W^- may decay to either leptonic or hadronic final states. This interaction channel is only relevant if the flux is very hard (e.g. E^{-2} .)

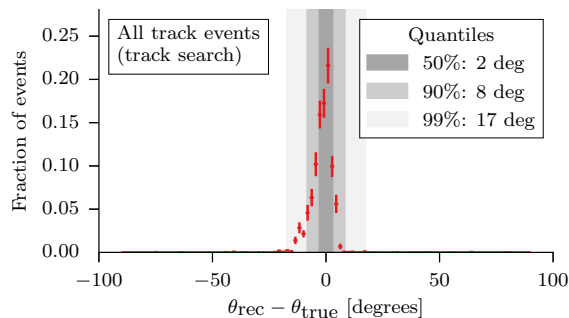
Spectrum	Flavor	Reaction	% of total events	% classified as	
				Cascade	Track
Conventional atmospheric	ν_e	CC	15.6	99.9	0.1
		NC	1.2	100.0	0.0
	ν_μ	CC	48.4	65.6	34.4
		NC	34.7	100.0	0.0
Prompt atmospheric	ν_e	CC	69.0	99.0	1.0
		NC	8.3	99.5	0.5
		GR	0.2	93.1	6.9
	ν_μ	CC	13.9	52.6	47.4
		NC	8.6	99.4	0.6
$E^{-2.5}$ astrophysical	ν_e	CC	42.6	97.8	2.2
		NC	5.7	98.4	1.6
		GR	0.6	48.7	51.3
	ν_μ	CC	9.7	48.5	51.5
		NC	5.9	98.2	1.8
	ν_τ	CC	29.5	96.7	3.3
		NC	5.9	97.8	2.2
E^{-2} astrophysical	ν_e	CC	34.7	90.2	9.8
		NC	6.8	92.2	7.8
		GR	5.0	33.7	66.3
	ν_μ	CC	11.5	37.7	62.3
		NC	6.9	91.8	8.2
	ν_τ	CC	28.1	85.6	14.4
		NC	7.1	92.2	7.8



(a) All events classified as cascades




(b) All events classified as tracks, with directions reconstructed using the joint vertex/time/direction/energy cascade reconstruction described in Section 5.5.1




(c) All events classified as tracks, with directions reconstructed using the simple outgoing-track search described in Section 5.5.2.

Figure 5.18: Zenith angle reconstruction error for simulated neutrino events that pass the final event selection. The grey bands in each panel show the ranges of zenith angle error that contain the given fraction of events. ((a)) shows events classified as cascades, while ((b)) and ((c)) show events classified as tracks. For starting track events the large displacement of PE detections from the neutrino vertex along the track direction provides a much better estimate of the neutrino zenith angle than the cascade reconstruction.



(a) Events classified as cascades



(b) Events classified as tracks

Figure 5.19: Quantiles of the absolute zenith angle distribution (cf. Figure 5.18) as a function of deposited energy. The median error is nearly constant with energy, while the rate of rare failures increases with increasing deposited energy, especially for track events.

5.5.4 Observable distributions for each component

Figure 5.20 shows the distributions of reconstructed zenith angle and deposited energy for conventional and prompt atmospheric neutrinos, astrophysical neutrinos (assuming an E^{-2} flux) and penetrating atmospheric muons. The figure shows the zenith angle in 15 bins, but for purposes of this analysis all zenith angles greater than $\sim 80^\circ$ ($\cos\theta = 0.2$) are lumped together as the “northern sky”. The dividing line is placed at $\cos\theta = 0.2$ rather than the geometric horizon at $\cos\theta = 0$, since the ~ 10 kilometers water-equivalent of overburden at this zenith angle are already sufficient to remove the vast majority of atmospheric muons that would otherwise veto atmospheric neutrino events. The southern sky ($\cos\theta > 0.2$) is further divided into 2 halves at $\cos\theta = 0.6$.

In general the zenith angle and deposited energy distributions reflect the characteristics of the underlying flux, albeit with some distortion due to biases in the event selection. Conventional atmospheric neutrinos are concentrated at deposited energies of a few TeV and in the northern sky around the geometric horizon; the smaller contribution above the horizon in the southern sky is further suppressed by vetoing muons. Since conventional atmospheric neutrinos have the largest fraction of ν_μ , they represent the largest contribution to the track-like portion of the data sample, shown in Figure 5.21. The fraction of events in the track sample provides another constraint on the conventional atmospheric flux in addition to the deposited-energy and zenith distributions. The remaining events are classified as cascades; their observable distributions are shown in Figure 5.22. The astrophysical component, shown as an isotropic flux with a normalization at the best fit of [13], dominates above 100 TeV in the northern and 30 TeV in the southern sky. The prompt atmospheric component, shown with a normalization at the previously published upper limit of 3.8 [36], never provides a dominant contribution to the observed event rate. Instead, a large prompt component appears as an excess over the conventional atmospheric and astrophysical components in the northern sky in the 30–60 TeV region that is not matched in the southern sky. While the exact cross-over energies depend on the normalization and spectral index of the astrophysical component, the ordering of the energy ranges where each component can be constrained is generic.

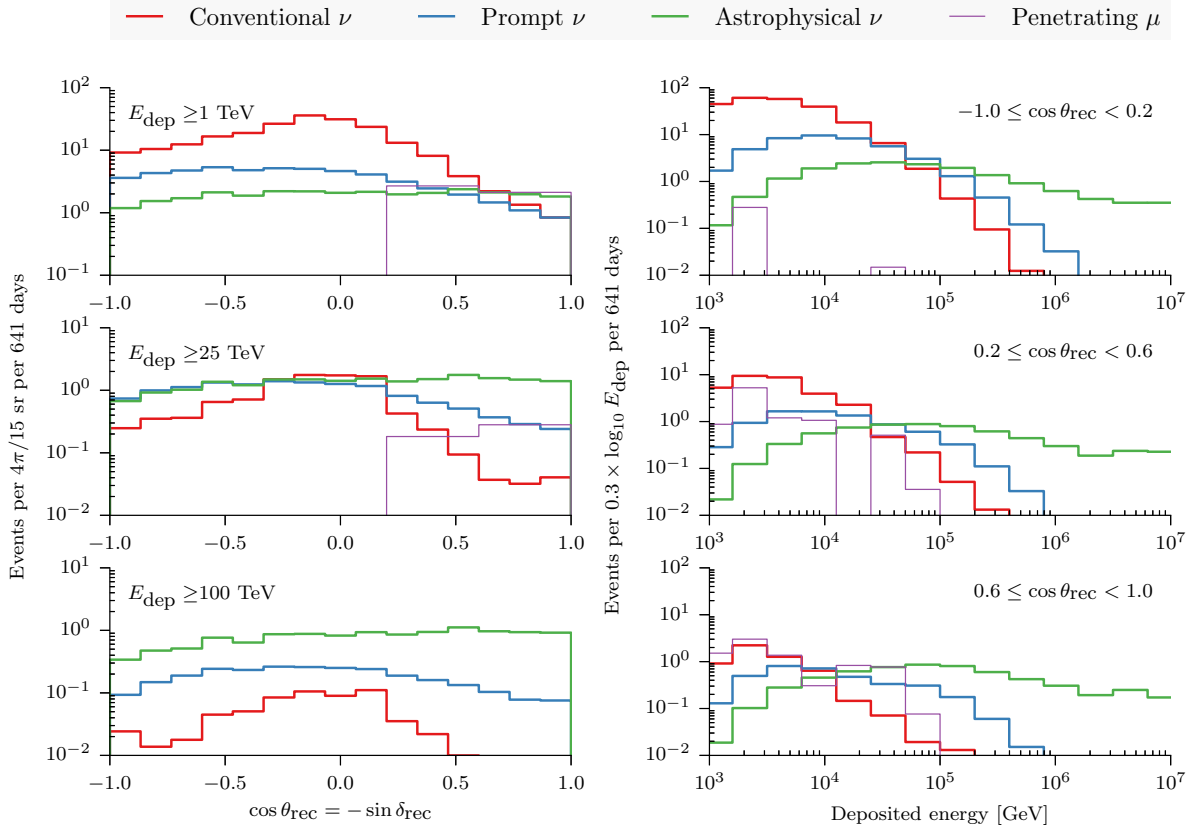


Figure 5.20: Distributions of deposited energy and zenith angle for conventional atmospheric neutrinos ($1 \times \Phi_{\text{HKKMS}}$, cf. Section 5.1.2), prompt atmospheric neutrinos ($3.8 \times \Phi_{\text{ERS}}$, the 90% upper limit from [36], cf. Section 5.1.3), astrophysical neutrinos ($10^{-18} \times (E/100 \text{ TeV})^{-2} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$, cf. Section 5.3), and penetrating atmospheric muons ($1 \times \Phi_{\text{SIBYLL+DPMJET}}$, cf. Section 5.4). The left column shows the distribution of reconstructed zenith angles for all events above the deposited energy threshold given in each panel. The binning is finer than that used in the final analysis in order to show the underlying structure. The penetrating muon component, however, has been averaged over the final zenith angle bins (edges at $\cos \theta_{\text{rec}} = -1.0, 0.2, 0.6, 1.0$). The right column shows the deposited energy spectra in these same zenith angle ranges.

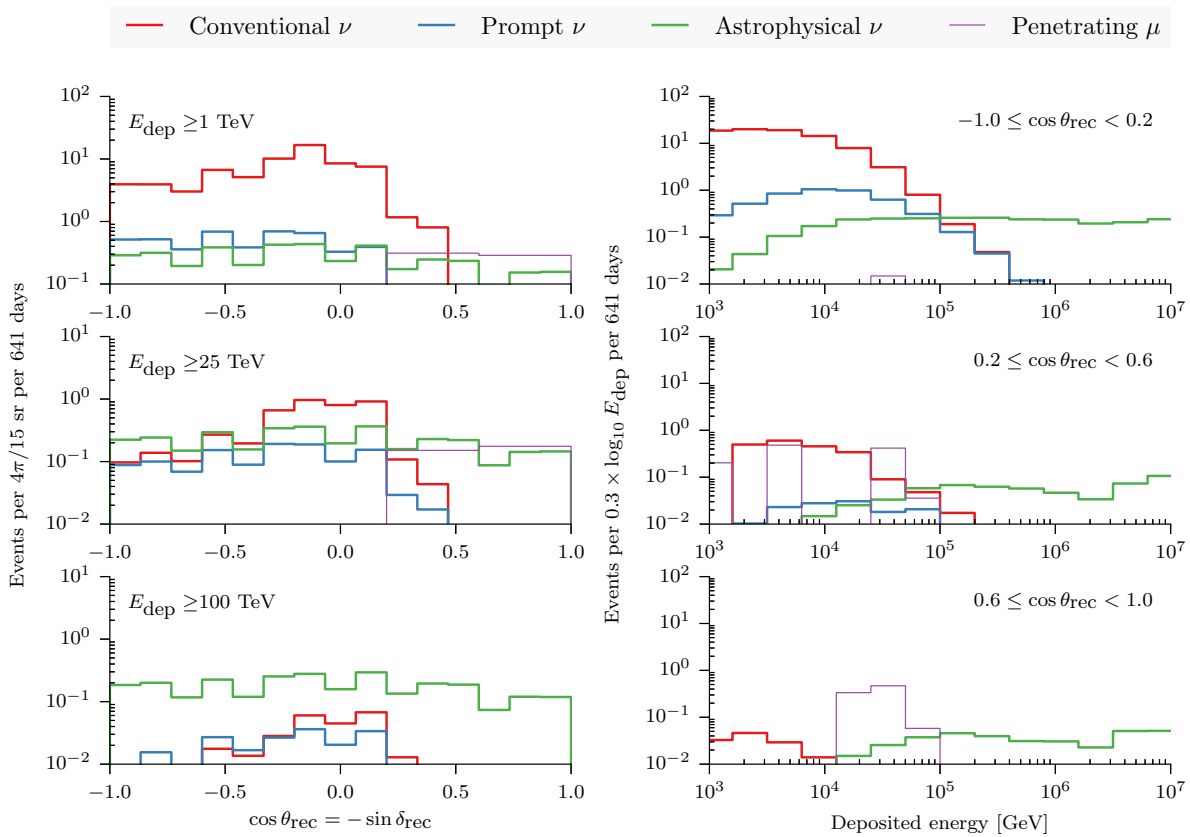


Figure 5.21: Distributions of deposited energy and zenith angle for the four flux components shown in Figure 5.20, but displaying only the events classified as starting tracks. Because of the much lower statistics in both data and simulation, the track sample is binned more coarsely for analysis, with 2 rather than 4 bins per energy decade.

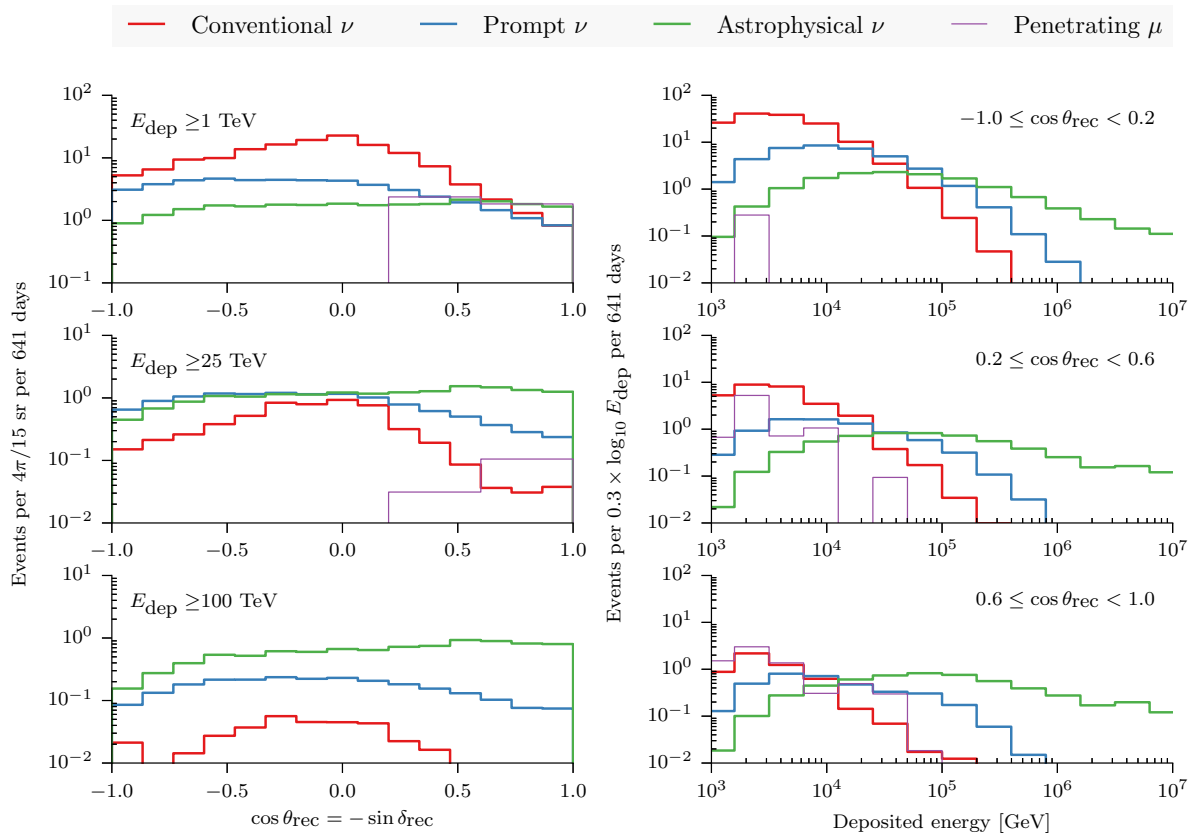


Figure 5.22: Distributions of deposited energy and zenith angle for the four flux components shown in Figure 5.20, but displaying only the events classified as cascades.

5.6 Likelihood fitting

The parameters of the model that best fit the observed data are determined through a binned likelihood fit. In this procedure, the data sample is binned in the three observables: reconstructed deposited energy, reconstructed zenith angle, and presence of a detectable outgoing track. The observed count n_i in each bin i is compared to a model that predicts the mean count rate λ_i in each bin through a Poisson likelihood function

$$L = \prod_{\text{bins } i} \frac{e^{-\lambda_i} \lambda_i^{n_i}}{n_i!}. \quad (5.29)$$

5.6.1 Forward folding

The mean rates λ_i are taken from the template histograms shown in Figures 5.22 and 5.21, which in turn are computed by weighting simulations to the set of flux models under consideration. This convolves the flux model with the detector response to obtain observable distributions, a procedure known as forward folding. The model is fit to the data by varying its parameters until (5.29) is maximized.

5.6.2 Unfolding

The forward-folding approach described above is perfectly adequate if the functional form of the underlying flux model is well known. It can however be useful to relax this restriction, and simply invert the transformation from interesting physics parameters (e.g. neutrino energy) to observables (e.g. deposited energy), in order to recover the distribution of physics parameters from the distribution of observables without assuming any underlying spectrum. This procedure is generally known as unfolding, but is not actually distinct from forward folding. We can turn the forward folding described above into an unfolding by simply using different basis functions⁹ as shown in Figure 5.23. Each basis function is peaked just below the central neutrino energy, but has a long tail to lower energies due to NC interactions that can deposit arbitrarily small fractions of the neutrino energy in the detector. Fitting the amplitudes of these basis functions to the data is equivalent to representing the E^2 -weighted flux $E^2\Phi_\nu$ as a piecewise-constant function.

⁹This is not an original insight; the idea was put forth by N. Whitehorn in the course of discussions about the best way to represent knowledge of the neutrino spectrum for [13].




Figure 5.23: Basis functions for energy spectrum unfolding. The heavy black line shows the deposited energy distribution for an E^{-2} astrophysical flux. The colored lines show the decomposition into smaller basis functions, each covering only a limited range of neutrino energies. Only every third basis function is shown for clarity; the intermediate basis functions are shown in thin lines.

5.6.3 Energy scale correction

One complication arises from the fact that the energy scale assumed in the simulations for **IC86** was 10% higher than those for **IC79**, the result of a calibration study that used the well-known Cherenkov light yield of minimum-ionizing muons to measure the total effective detection efficiency of the **DOM** in situ [16]. The optical efficiency of the real **DOMs** did not change between the data-taking years, however, so it was necessary to correct the **IC79** simulations to assume the same efficiency as the **IC86** simulations.

The energy scale or optical efficiency gives the mean number of **PE** detected per GeV of deposited energy. For events that deposit a fixed amount of energy, an increased energy scale raises the number of detected **PE** as well as the energy reconstructed from the number of **PE**. Since the acceptance of the event selection increases with collected charge, this causes some events that would have previously been excluded to pass the cuts. These two effects are modeled as a linear shift in all charge- and energy-related observables, followed by the threshold correction illustrated in Figure 5.24.

5.6.4 Confidence intervals

The 68% confidence ranges on each model parameter are obtained from a likelihood-ratio test. To test whether a value of one parameter can be rejected at the desired confidence level, the parameter is constrained to that value while all other parameters are varied to maximize the conditional likelihood. The ratio between this conditional likelihood maximum and the global maximum is the **profile likelihood** [137]. It can be used to construct a test statistic

$$-2\Delta \ln L = -2(\ln L - \ln L_{\max}) \quad (5.30)$$

whose distribution approaches that of a χ^2 with 1 degree of freedom in the large-sample limit [138]. If $-2\Delta \ln L > 1$ (2.71), then the tested value of the model parameter is rejected at more than 68% (90%) confidence. The test statistic does not necessarily follow a χ^2 distribution when the sample size is finite or a parameter is close to a bound. In such cases the exact confidence level can be derived numerically from Monte Carlo trials. In this analysis, however, the exact confidence intervals were found to be only slightly smaller than intervals derived from the χ^2 approximation.

We use a slightly different approach to derive the confidence intervals on the amplitudes of the basis functions in the unfolding approach presented in Section 5.6.2. Instead of forming a profile likelihood by allowing the nuisance parameters to float to their conditional best-fit values, we fix them all to their global




Figure 5.24: Energy scale correction for IC79 simulations. The black line in the upper panel shows the neutrino acceptance as a function of deposited energy, expressed as an effective volume. The colored lines show the approximate acceptances that would be obtained if the energy scale, or number of PE detected per GeV of deposited energy, were raised by some fraction. The resulting increase in acceptance, plotted in the lower panel, is applied to the IC79 simulations as a weight that depends on deposited energy.


best-fit values, and adjust the amplitude of each basis function in turn until $\ln L$ changes by 0.5. While this neglects the correlations between the allowed amplitudes of neighboring basis functions, it gives a clearer picture of the constraints on individual basis functions.

5.6.5 Sensitive energy range of the astrophysical spectrum fit


This analysis is only sensitive to neutrinos over a limited range of energies, and that range can influence the results. This is particularly important for the astrophysical component. If the underlying spectrum is exactly a power law, as assumed in the parameterization, then it is scale-free, and the energy range over which it is fit does not matter. If it deviates from a power law, however, then power-law fits over different energy ranges can be incompatible. In order to fully characterize the fit, it is helpful to determine the range of energies over which it applies. We can do this with a variant of the profile likelihood construction presented above. To determine the lower limit of the sensitive energy range, we remove all simulated events with $E_\nu < E_{\nu,\min}$ from the template that describes the astrophysical component, and evaluate the conditional best-fit likelihood, keeping the normalization and spectral index of the astrophysical component fixed. The lower limit of the sensitive energy range is the largest $E_{\nu,\min}$ for which $-2\Delta \ln L < 1$. By keeping the astrophysical flux model fixed, this construction produces the sensitive energy range corresponding to the best fit, and by allowing the other parameters to float, it accounts for degeneracies between the atmospheric components and the astrophysical flux at low energies. A similar procedure is used to derive the upper boundary of the energy range by removing simulated events with $E_\nu > E_{\nu,\max}$ and finding the smallest $E_{\nu,\max}$ for which $-2\Delta \ln L < 1$.

Figure 5.25a shows $-2\Delta \ln L$ as a function of $E_{\nu,\min}$ for the best fit to the experimental data given in Table 6.1. The dip in the test statistic between 10 and 20 TeV is caused by a corresponding under-fluctuation in the data; reducing the event rate predicted from the astrophysical flux in that region momentarily improves the agreement between prediction and data before the missing part of the template begins to degrade the quality of the fit in earnest above 20 TeV. A similar effect can be seen in the scan of $E_{\nu,\max}$ shown in Figure 5.25b. Removing very high-energy events from the template always improves the quality of the fit, since no events were observed there. The fit quality only begins to degrade substantially when the $E_{\nu,\max}$ contributes significantly to the last populated deposited-energy bin. The bounds of the neutrino energy ranges are larger than the deposited energy at which the astrophysical component becomes noticeable

because of **CC** ν_μ interactions and **NC** interactions of all flavors, where only a fraction of the neutrino energy is deposited in the detector.



(a) Minimum sensitive neutrino energy



(b) Maximum sensitive neutrino energy

Figure 5.25: Profile likelihoods used to derive the neutrino energy range to which the astrophysical component fit is sensitive. The dashed lines mark the point where $-2\Delta \ln L = 1$. The best fit astrophysical flux given in Table 6.1 is valid for $25 \text{ TeV} < E_\nu < 1.4 \text{ PeV}$.

Chapter 6

Results

6.1 Forward-folding likelihood fit

The forward-folding likelihood fit approach described above was used to determine the fluxes of neutrinos and muons compatible with the observed events. In the first fit, the normalizations of the penetrating atmospheric muon component, the conventional and prompt atmospheric neutrino components, as well as the normalization Φ_0 and spectral index γ of the astrophysical component were allowed to vary freely, resulting in the best-fit parameters shown in Tab. 6.1. Figure 6.1 shows the deposited energy spectra corresponding to the best-fit model parameters. Figure 6.2 shows the zenith angle distributions of the sample with different energy thresholds.

Table 6.1: Best fit parameters and number of events attributable to each component. The normalizations of the atmospheric fluxes are relative to the models described in Chapter 5. The normalization Φ_0 and spectral index γ of the astrophysical flux are defined in Equation (5.25). The two-sided error ranges given are 68% confidence regions in the χ^2 approximation; upper limits are at 90% confidence. The astrophysical fit applies for $25 \text{ TeV} < E_\nu < 1.4 \text{ PeV}$ (cf. Section 5.6.5). The goodness-of-fit p-value for this model is 0.2.

Parameter	Best-fit value	No. of events
Penetrating μ flux	$1.73 \pm 0.40 \Phi_{\text{SIBYLL+DPMJET}}$	30 ± 7
Conventional ν flux	$0.97^{+0.10}_{-0.03} \Phi_{\text{HKKMS}}$	280^{+28}_{-8}
Prompt ν flux	$< 1.52 \Phi_{\text{ERS}} \text{ (90\% CL)}$	< 23
Astrophysical Φ_0	$2.06^{+0.35}_{-0.26} \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$	87^{+14}_{-10}
Astrophysical γ	2.46 ± 0.12	

This simple model does not describe the data perfectly. There is a notable departure in the southern sky around 30 TeV. However, the excess is not statistically significant; correlated fluctuations of the observed size or greater are expected from a smooth underlying power-law spectrum in 5% of experiments. The events in the energy and zenith region of the excess are overwhelmingly cascade-like and display no signs of early hits from penetrating atmospheric muons. Their rate far exceeds that expected from penetrating muon

background and conventional atmospheric neutrinos (~ 1 event per year), and their distribution in time and within the fiducial volume is compatible with a uniform one. Known sources of systematic uncertainty in the neutrino acceptance of the detector, like the optical properties of the South Pole ice or the optical efficiency of the DOM, are unable to create structure in the observed energy distribution. More details can be found in Appendix B. At present this has to be interpreted as a statistical fluctuation. Future searches using more years of data will help constrain the cause of the excess, either by reducing its significance or by strengthening it enough that definitive statistical statements can be made.

The spectral index of 2.46 needed to explain the low-energy data has implications for the underlying neutrino production mechanism. As pointed out in [139], pp interactions produce neutrinos and γ -rays that follow the same scale-free power-law spectrum, and the γ spectra from pp interactions at \sim GeV energies can be extrapolated to the TeV range where IceCube observes neutrinos. This extrapolation argument does not apply to $p\gamma$ interactions. If the diffuse extragalactic γ background measured by Fermi-LAT is due to extragalactic pp interactions in optically thin regions, then the spectral index of the associated neutrino spectrum must be smaller than 2.2. [139]. The data presented here indicate that the neutrino spectrum is softer than $E^{-2.2}$ with 90% confidence (see Figure 6.4a), implying that one of these assumptions is violated.

All of the parameters in Tab. 6.1 are correlated except for the conventional atmospheric neutrino flux normalization. The latter is determined mostly by the northern-sky data below 10 TeV deposited energy, where the contributions of the other components are negligible, and is compatible with the expected normalization [47] to within statistical errors, providing a useful check of the neutrino acceptance calculated from simulation. Similarly, the low-energy component provides a verification of the atmospheric neutrino veto independent of the observed astrophysical excess, as shown in Figure 6.3. The prompt atmospheric neutrino flux, on the other hand, can provide a significant contribution to the overall event rate between 10 and 100 TeV deposited energy, but has no region where it contributes exclusively. Its inferred normalization depends on assumptions about the astrophysical neutrino flux. The correlations between the astrophysical and prompt atmospheric components are shown in Figure 6.4. Since the power-law index of the astrophysical flux is constrained primarily by the large number of events below the pivot point at 100 TeV, the normalization and index are correlated. Similarly, the prompt normalization is correlated with the astrophysical index; as the index is forced to smaller values, a larger prompt flux is required to explain the data between 10 and 100 TeV deposited energy. The effect of the assumed prompt normalization on the inferred astrophysical flux is much weaker, as shown in Figure 6.5. If the prompt normalization is forced to the value predicted by [14],

for example, the best-fit astrophysical power law index only hardens from 2.46 to 2.44. The normalization of the penetrating muon component is constrained by the excess of southern-sky data over the conventional atmospheric neutrino expectation below a few TeV, and is weakly correlated with it (not shown).




Figure 6.1: Deposited energy spectra from the northern and southern skies (points) with the best-fit combination of atmospheric and astrophysical contributions from Table 6.1. Below 3 TeV, the events observed from the northern sky are adequately explained by conventional atmospheric neutrinos. In the same energy range in the southern sky, penetrating atmospheric muons account for the remaining events. Above 10 TeV, an extra component is required to account for the observed high-energy events, especially those in the southern sky. Since atmospheric neutrinos of any kind are often vetoed by accompanying muons, the excess is best explained by astrophysical neutrinos. We interpret the excess over the best-fit sum around 30 TeV as a statistical fluctuation.




Figure 6.2: Zenith angle distribution of events depositing more than 1, 25, and 100 TeV (points) with the best-fit combination of atmospheric and astrophysical contributions from Table 6.1, using the same color scheme as in Figure 6.1. At the lowest energies the sample is concentrated at the horizon, as expected from conventional atmospheric neutrinos. The astrophysical component contributes significantly to the sample above 25 TeV, and the bulk of the sample is down-going. By 100 TeV only the astrophysical component remains, and the up-going flux is suppressed by absorption in the Earth.






Figure 6.3: Verification of atmospheric neutrino veto with low-energy data. The points show events depositing less than 3 TeV, while the stacked histograms show the expected contributions from conventional atmospheric neutrinos, penetrating muons, and the negligible contribution of astrophysical neutrinos, using the color scheme of Figure 6.1. These match the observed data much better than the dotted line, which shows the number of events that would be collected if atmospheric neutrinos were never vetoed by accompanying muons.



(a) Likelihood profile in astrophysical power-law index γ and normalization $\Phi_0/10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$. $E^{-2.5}$ is strongly favored over E^{-2} .



(b) Likelihood profile in astrophysical power-law index γ and prompt atmospheric neutrino flux normalization [14].

Figure 6.4: Profile likelihood scans showing the correlation between the astrophysical power-law index and the normalizations of the astrophysical and prompt atmospheric components. In each plot, the colors show the test statistic (5.30), obtained by fixing the parameters shown on the axes and varying all others to obtain the conditional best fit. The x shows the best-fit point as in Tab. 6.1 and the contours show confidence regions in the χ^2 approximation [138] with 2 degrees of freedom. The thin dotted line shows the conditional best fit for each value of γ .




Figure 6.5: Best-fit parameters of the astrophysical neutrino flux as a function of the assumed prompt contribution, given in units of the prediction of [14]. The best fit prompt normalization is zero, but the conditional best-fit astrophysical spectral index and normalization do not change significantly if the prompt flux is assumed to be at the level of [14] (1 on the x -axis).

6.2 Descriptive statistics: unfolding

These correlations would not be problematic if the model of the astrophysical flux were exact, but since its sources are not known, any single model will necessarily be an approximation. It is useful to examine how assumptions about astrophysical models affect the upper limit on the prompt atmospheric flux normalization. The first assumption made is that the astrophysical flux must follow a single power-law energy distribution. This assumption can be relaxed by describing the astrophysical neutrino flux with a piecewise constant function of neutrino energy as shown in Figure 6.6. The observed excess in the deposited energy spectrum is reflected in a corresponding excess in the neutrino energy spectrum, and the additional freedom granted to the astrophysical component weakens the 90% upper limit on the prompt atmospheric flux from 1.52 to 1.75 times the prediction of [14]. This remaining limit is driven primarily by the assumption that the astrophysical neutrino flux is isotropic. If this assumption is weakened by allowing the astrophysical fluxes that contribute to the northern- and southern-sky data to vary independently, the limit relaxes further to 3.69. While this limit is not meaningfully smaller than the previously published limit of 3.8 [36], it involves many fewer assumptions about the nature of the astrophysical neutrino background.




Figure 6.6: Unfolding the non-atmospheric excess as piecewise-constant per-flavor fluxes $E^2\Phi$. The horizontal error bars show the range of primary neutrino energies that contribute to each bin, while the vertical error bars show the range of $E^2\Phi$ that change the $-2\Delta \ln L$ test statistic by less than 1. The black points show the fit to the data sample presented here; the light grey data points are from the 3-year data sample of [13], shifted slightly to the right for better visibility. Above the highest observed energy, the error bars provide upper limits on the flux; these are less constraining than the upper limits of [140] above 10 PeV. The thin lines show models for the diffuse astrophysical neutrino background: the upper bound from the total luminosity of EeV cosmic rays from [8], the AGN core emission model of [141], and the starburst galaxy model of [133].

Chapter 7

Conclusion and future directions

The analysis presented used a veto-based technique to isolate 388 events starting in the IceCube instrumented volume and depositing more than 1 TeV from 641 days of data, of which 92% were neutrino events. Astrophysical neutrino candidates were observed in the southern sky with deposited energies as low as 10 TeV, far below the threshold of the previous high-energy starting event analysis [12, 13] and in a region inaccessible to the traditional up-going track analysis [36]. The analysis characterized the contributions of penetrating atmospheric muons, conventional and prompt atmospheric neutrinos, and astrophysical neutrinos to the data sample using a likelihood fit to the distributions of deposited energy and zenith angle for cascade and starting-track events.

The analysis yielded new information about the behavior of the neutrino spectrum between 10 and 100 TeV. If the energy spectrum of the astrophysical neutrinos is a single power law, then it must have a spectral index of 2.46 ± 0.12 , softer than the typical E^{-2} benchmark spectrum. The $\gamma = 2$ hypothesis can be rejected with 99% confidence under this assumption. The new constraint on the spectral index is due primarily to the lower deposited-energy threshold of this analysis. If the threshold is raised to 60 TeV (corresponding to sensitivity for $E_\nu > 100$ TeV), then the best-fit spectral index hardens to 2.26 ± 0.35 , compatible with the previous high-energy result [13]. The statistically insignificant excess that appeared in the down-going data near 30 TeV, a region where atmospheric leptons are heavily suppressed, had only a minor influence on the inferred spectral index of the astrophysical neutrinos. If the spectral index is forced to $\gamma = 2$, then the per-flavor normalization Φ_0 (cf. Equation (5.25)) drops to $1.22 \pm 0.5 \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$, consistent with the previously published 90% C.L. upper limit of 1.44 derived from northern-sky ν_μ events [36]. At the same time, the analysis searched for atmospheric neutrinos from charmed meson decay. No such component was observed, resulting in an upper limits on their flux. These limits depend strongly on assumptions about the astrophysical neutrino background, and range from 1.52 times the prediction from

perturbative QCD [14] at 90% confidence when the astrophysical flux is assumed to follow a single isotropic power-law distribution to 3.69 times the prediction when it is described with a piecewise constant function of energy and zenith angle.

The constraints on the astrophysical flux are currently limited by the small number of observed high-energy events, making it difficult to draw strong inferences about the classes of cosmic-ray accelerators from the characteristics of the associated neutrino spectrum. Beyond astrophysical considerations, the inability to model the astrophysical flux precisely and reliably extrapolate its angular and energy distribution to lower energies impedes any attempt to measure the level of charmed-meson production in air showers via high-energy neutrinos. Both of these problems may be approached with more and different data. IceCube will continue to collect data, and future iterations of this analysis will be able to use at least twice as many high-energy neutrino events to constrain the energy spectrum and eventually possible anisotropies of the astrophysical neutrino flux.

Beyond simply collecting more events, a worthwhile follow-up analysis will need to be able to meaningfully constrain the flavor composition of the astrophysical neutrino flux. While the event selection used in this analysis was primarily sensitive to cascade events, it retained some sensitivity to **CC** ν_μ track events, and included methods for identifying them as such. Nonetheless, the number of positively-identified ν_μ events was too small to draw any conclusions about the flavor ratio of the observed astrophysical neutrino flux. This could be improved by increasing the selection acceptance for starting track events, improving track/cascade separation, or adding a dedicated through-going **CC** ν_μ sample. Ideally, all three improvements would be included. The final improvement to the characterization of the flavor ratio would come from the observation of positively-identifiable **CC** ν_τ interactions through the “double-bang” signature of a ν_τ interaction followed by the decay of the τ lepton after on average $50 \cdot \frac{E_\tau}{1\text{PeV}}$ m [48]. While quite sophisticated techniques have been developed to identify such events (see, for example, [142]), no ν_τ candidates have been observed to date, and more work is needed to make the identification robust against systematic errors in the modeling of photon transport in the South Pole ice.

Constraints on charmed meson production in the atmosphere can, in principle, be improved by analyzing penetrating muon events jointly with neutrino events. These are produced in the same decays as prompt muon neutrinos, but can be detected in far greater numbers and have no astrophysical background to contend with. The muon flux from charmed meson decay does, however, have to contend with the muon flux from rare ($\sim 10^{-5}$ branching fraction) muonic decays of light, unflavored mesons like the η , ρ , and ω [2]. While

these rare decay channels are unimportant at TeV energies where π and K mesons can decay in flight, the extremely short lifetimes of the parent mesons ($\sim 10^{-20}$ s) mean that their decays are not suppressed at high energies, and that they can become the dominant source of muons with energies greater than 1 PeV [128]. Aside from these modeling issues, there are also significant experimental difficulties to overcome, some more and some less severe than those encountered in traditional underground muon experiments (cf. Section 2.2.1). Unlike a detector buried in a deep mine whose overburden is composed of a variety of different rock strata, the ice at the South Pole has a constant, known chemical composition and density, reducing the inherent systematic error in extrapolating from energies measured at depth to surface energies. At the same time, the overburden is relatively thin, so single, high-energy muons from prompt decays are often embedded in bundles of muons from decays in other branches of the shower, especially near the zenith, where the conventional muon flux is most suppressed, but where the overburden is also the thinnest. While it is in principle possible to identify high-energy muons inside a bundle of lower-energy muons by the pattern of stochastic energy losses [143], much more work is needed to understand the resolution and systematic error of joint energy/multiplicity inference.

Several genuinely novel methods were developed in the course of this work, including the multidimensional B-spline histogram interpolation method used to smoothly parameterize photon transport in ice [15], the cascade vertex, direction, and energy reconstruction method based on it [16], the calculation of the fraction of atmospheric neutrinos that are vetoed by muons [17], the extension of the neutrino event simulation to treat full air showers (cf. Section 5.2.4), and the targeted penetrating muon simulation scheme (MuonGun, cf. Appendix A). These will be useful tools in the future research program outlined above.

APPENDIX

MuonGun

A.1 Introduction

In diffuse neutrino searches it is almost always necessary to estimate the background due to atmospheric muons from simulation. In IceCube this has typically been done directly, by simulating air showers to ground level with **CORSIKA**, propagating the muons in the shower through the firn and ice with **MMC** to a cylindrical sampling surface surrounding the detector, and then weighting the simulated events to an assumed cosmic-ray flux. Though this method offers the highest possible precision available from the chosen simulation software, it suffers from two key inefficiencies. First, since the simulation starts with cosmic-ray primaries rather than in-ice muons, one has only loose control over the characteristics of the muon bundles that actually reach the detector. For example, if one were interested only in single muons with a few TeV of energy, one would spend quite a lot of time simulating both showers whose muons never reached the detector as well as those that result in high-multiplicity bundles. Second, the direct approach makes it necessary to repeat the entire simulation chain in order to change aspects of the air shower simulation such as atmospheric profile or hadronic model.

An alternative approach is to de-couple the air shower simulation and muon propagation from the remainder of the simulation by constructing a parametrization of the muon flux under the ice and drawing muon bundles from the parameterized distribution. This allows one to generate specific bundle configurations and weight them properly, and also to re-weight existing simulated events to a muon flux associated with different assumptions about interactions in the atmosphere.

The parametric approach is already used heavily by **ANTARES** in the form of their MUPAGE event generator [134]. MuonGun is an application of the same technique to IceCube simulation.

A.2 Constructing the parameterization

A.2.1 Propagation

The initial set of fits was made using sets of **dCORSIKA Format 2000 (F2k)** files that were stored as an intermediate step of IC79 simulation production. Approximately 25 billion air showers were simulated on an E^{-2} spectrum from 600 GeV/nucleon to 100 PeV, with a 10:5:3:2:1 ratio of proton/He/N/Al/Fe primaries. The simulations were split evenly between atmosphere parametrizations 11 to 14 and used SIBYLL [130] to simulate hadronic interactions. The muons from the shower were then propagated through firn and ice with **PROPOSAL** [73] and recorded at a set of vertical depths spaced every 100 m between 1 and 2.8 km.

A.2.2 Variables

At each vertical depth and zenith angle the number of surviving muons was tabulated, along with the energy of each and its distance from the shower axis. This is a complete description of the bundle under 4 approximations:

1. The flux of cosmic-ray primaries that reach the atmosphere and their daughter muons is independent of the azimuthal arrival direction.
2. The bundle is azimuthally symmetric around the shower axis.
3. All muons in the bundle are perfectly parallel to each other and to the shower axis.
4. The shower front has no curvature.

The first assumption is violated by deflection in the Earth's magnetic field, but the same approximation is used whenever **dCORSIKA** showers are randomized in azimuth before being fed in to IceTray (i.e. nearly always). The remaining three approximations are important only if it is possible to measure very detailed properties of the bundle structure over relatively short (~ 1 km) observation baselines.

The variables were filled into a set of histograms with increasing dimensionality:

1. Multiplicity: 3 dimensions (zenith, depth, multiplicity)
2. Radius: 4 dimensions (zenith, depth, multiplicity, radius)
3. Energy: 5 dimensions (zenith, depth, multiplicity, radius, energy)

A.2.3 Spline tables

In their parametrization, [134] fit the single-muon flux, ratio of single- to n-muon bundles, normalized radial distributions, and normalized energy distributions to a stable of relatively simple functional forms (polynomials and power laws in depth and zenith angle). While this resulted in a small number of parameters that could be printed in a journal article, re-fitting those forms to **CORSIKA** simulation turned out to require a great deal of manual fiddling and yield unsatisfactory representations in the range of depths important for IceCube simulation.

Instead of simple functions, this parameterization uses tensor-product B-splines from the Photospline package [15] to smoothly interpolate between the bins of the histograms. The resulting surfaces fit much more closely to the underlying histograms than the functional forms given in [134], but without requiring a great deal of creativity in inventing new terms to fix mismatches that appear only in small regions of the parameter space. With appropriately chosen regularization they can even be used to reliably extrapolate into unpopulated regions of the histograms.


Since B-splines are piecewise polynomials of some fixed order, they can only perfectly represent polynomial data. Some additional transformations were applied to make the distributions more polynomial-like:

- Flux: $\ln \Phi$ vs. $\cos \theta$, depth, multiplicity
- Radius: $\ln(dP/dr^2)$ vs. $\cos \theta$, depth, multiplicity, radius
- Energy: $\ln(dP/dE)$ vs. $\cos \theta$, depth, multiplicity, radius, $\ln E$


All were fit with an order-2 regularization that penalizes any curvature in the fit surface, that is, it prefers straight lines in the absence of strong evidence from the data. The transformations applied to the flux and radial distributions turns this into exponential extrapolation; the energy distributions are extrapolated as power laws.

A.3 Fits


Figures A.1–A.5 show examples of fits to **CORSIKA** simulation using hadronic interactions from SIBYLL 2.1 [130], weighted to the spectrum of Hoerandel [144].



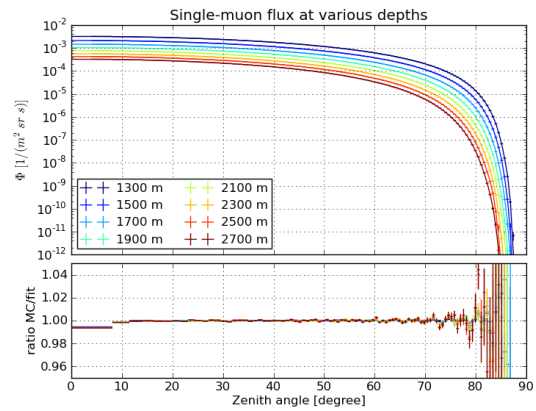
(a) March 31st atmosphere



(b) July 1st atmosphere




(c) October 1st atmosphere




(d) December 31st atmosphere


Figure A.1: Total single-muon flux as a function of depth and zenith angle for various South Pole atmospheres.




(a) Zenith 9 degrees



(b) Zenith 56 degrees

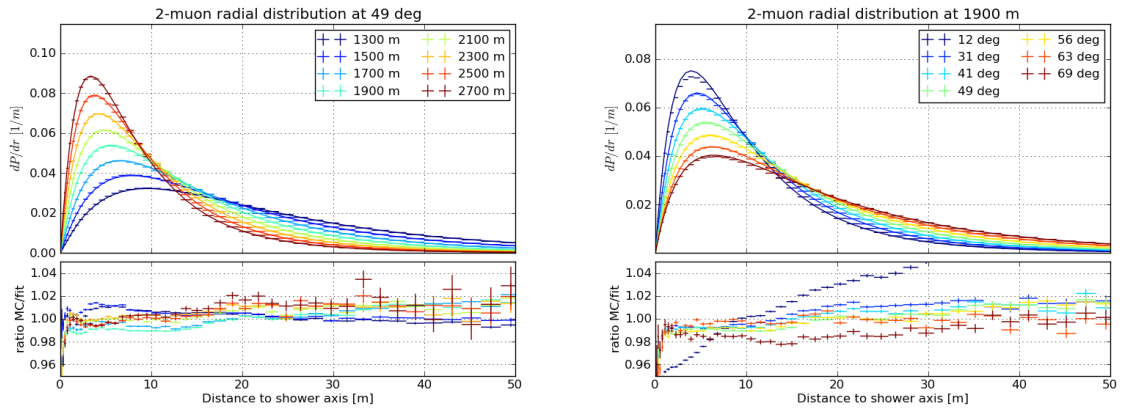


(c) Zenith 66 degrees




(d) Zenith 75 degrees

Figure A.2: Total flux as a function of depth and bundle multiplicity for various zenith angles with the March 31st atmosphere.




(a) Varied depth

(b) Varied zenith



(c) Varied multiplicity

Figure A.3: Slices of the normalized radial distribution, March 31st atmosphere



(a) Varied depth

(b) Varied zenith

Figure A.4: Slices of the normalized single-muon energy distribution, March 31st atmosphere




Figure A.5: Slices of the normalized multi-muon energy distribution, March 31st atmosphere

A.4 Sampling

Having access to the underlying distributions opens up a number of possibilities for combining unbiased and weighted sampling techniques for different dimensions of the problem.

A.4.1 Fixed surface injection

The default method implemented in MuonGun is very similar to the one described in [134], where the bundles are injected on a fixed, upright cylinder. It consists of 3 phases: first, a choice of the core position and multiplicity, followed by a choice of radial distance from the shower axis for each muon in the bundle based on the core position and multiplicity, followed by a choice of energy for each muon.

1. Choose an arrival direction (θ, ϕ) uniformly from the upper hemisphere. Because the sampling surface is an upright cylinder rather than a sphere, the zenith angles are not distributed uniformly in $\cos \theta$. Instead, they are distributed proportional to the projected area of an upright cylinder,

$$\frac{dA_{\perp}}{d \cos \theta} = \pi r^2 |\cos \theta| + 2rl \sin \theta \quad , \quad (\text{A.1})$$

where r is the radius of the cylinder and l is its full height.

2. Choose a core position uniformly on the surface of the cylinder projected along the arrival direction, and find the vertical depth d of the intersection point.
3. Choose a multiplicity m uniformly from a set range.
4. Evaluate the flux at the chosen zenith angle, intersection depth, and multiplicity, and accept the configuration with a probability given by the likelihood ratio

$$L = \frac{\Phi(\theta, d, m) dA_{\perp}(\theta)}{\Phi(\theta_{min}, d_{min}, 1) \max_{\theta} dA_{\perp}} \quad . \quad (\text{A.2})$$

5. For single muons, the muon can be assumed to be colinear with the shower axis. For $m \geq 2$, draw samples from the (normalized) radial distribution $\frac{dP}{dr}(\theta, d, m, r)$ by rejection sampling. Then, draw a uniformly-distributed azimuthal angle and rotate the offset track around the shower axis.

6. For energy, draw a sample from a fixed offset power law $\frac{dP}{dE} \propto (E + b)^{-\gamma}$. This is done to avoid costly rejection sampling on a steeply falling spectrum as well as allow better control over the generated energy distribution.

A.4.2 Energy-dependent surface injection

When simulating background for a selection whose fiducial volume scales with deposited energy, it can be much more efficient to sample in different order by first choosing the energy of the muon first, and then choosing an injection surface based on the energy of the muon under consideration. This mode allows the generator to inject low-energy muons only into the core of the detector, while spreading high-energy muons over larger and larger sampling surfaces. Since effective livetime is proportional to the number of events injected per unit area, this allows for much larger livetimes at low energies (where underground muon fluxes are larger) without needing to simulate the large fraction of low-energy muons that pass only through the edges of the detector and will be removed by construction in the event selection.

A.4.3 Weighting

To actually make use of the simulation, one has to calculate weights to turn the event counts into rates. These are given by

$$w = \frac{\Phi(\theta, d, m) dA_{\perp}(\theta) \prod_{i=1}^m \frac{dP}{dt}(\theta, d, m, r_i) \prod_{i=1}^m \frac{dP}{dt}(\theta, d, m, E_i)}{N_{generated}(\theta, d, m, r_1, \dots, r_m, E_1, \dots, E_m)} . \quad (\text{A.3})$$

Based on its configuration the generator knows how to calculate $N_{generated}$ for any given muon bundle, so one can always re-weight to a different muon flux model based on other assumptions about the nucleon flux, hadronic model, or atmospheric profile.

A.5 Combining with CORSIKA simulation

Event samples generated using this method can be combined with conventional samples generated directly using **CORSIKA**. To calculate the correct weight, one simply has to parameterize the distribution of muons produced with the given **CORSIKA** configuration (just as one would for a physical flux) and add it to the denominator Equation (A.3). Figure A.6 shows a parameterization of the muon spectrum generated by the ‘‘cascade-optimized’’ **CORSIKA** configuration used to simulate 1 year of livetime above 30 TeV/nucleon for

IC79. Figure A.7 shows how these existing simulations can be combined with single muons from MuonGun with appropriate weights.


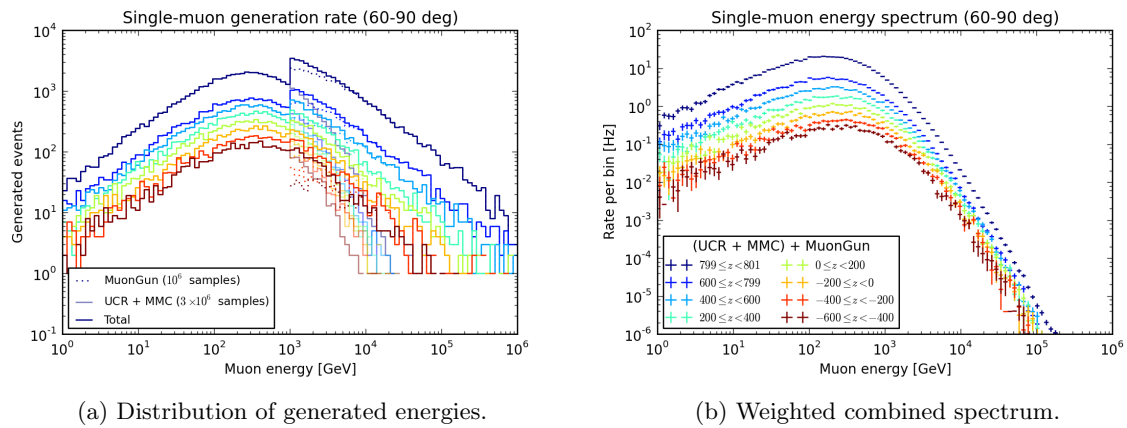


Figure A.6: Distribution of zenith angles and intersection depths produced by the “cascade-optimized” **CORSIKA** configuration (points with error bars) and a parameterization thereof (solid lines). The parameterization provides a good description of the generated distribution, and can be used to calculate weights for combining the sample with MuonGun simulation.



(a) Distribution of generated energies.

(b) Weighted combined spectrum.

Figure A.7: A demonstration of how **CORSIKA** and MuonGun simulations can be combined with appropriate weights. **CORSIKA** was run for 3 million primaries with the “cascade-optimized” settings, and single muons were generated from 1 TeV to 1 PeV as $\frac{dP}{dE} \propto (E + 800 \text{ GeV})^{-2}$ with MuonGun. The weighted distribution connects the two generation regions smoothly, with greatly reduced statistical errors above 1 TeV.

A.6 Comparison to conventional CORSIKA simulation

A.6.1 Distributions

Figure A.8 shows a comparison of the muon spectrum from conventional **CORSIKA** simulation and one generated with MuonGun, both weighted to the same primary flux model. The sampling and weighting procedures reproduce the distributions that the parameterization was fit to.

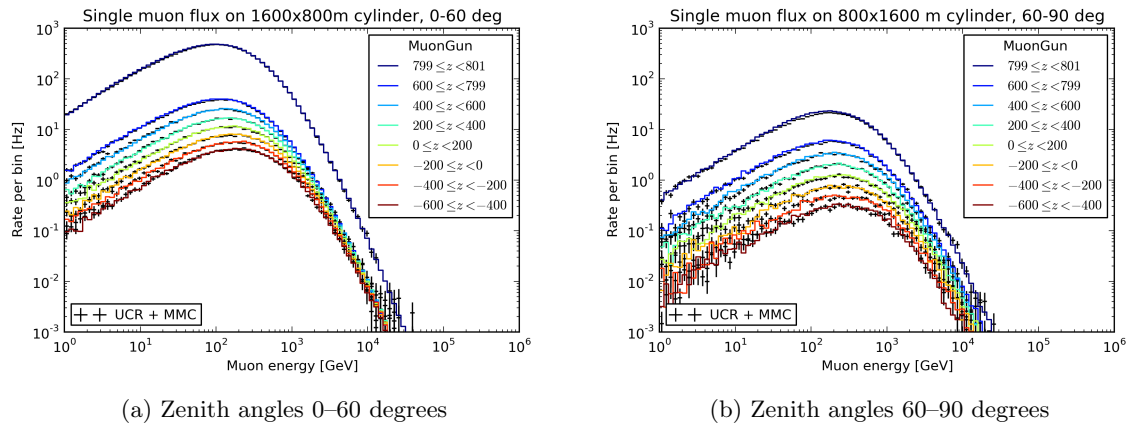


Figure A.8: A comparison of the muon spectra generated by **CORSIKA** (black crosses) and MuonGun (colored lines), both weighted to the primary flux parameterization of [144]

A.6.2 Performance

Table A.1 shows a comparison of the time budget for generating 1 million events where at least one muon reaches the 800×1600 m cylindrical sampling surface. **dCORSIKA** was run for 1 million primaries on an E^{-2} spectrum from 600 GeV to 100 EeV with a 10:5:3:2:1 primary ratio. Of these showers, only 29018 had at least one muon reach the sampling surface, so the run times for the various components were multiplied by 35. The aforementioned comparison shows **CORSIKA** in the worst possible light; with an energy threshold low enough to avoid distortions in the in-ice muon spectrum, much of the time is spent simulating showers where no muons make it to ground level, let alone to a kilometer below the surface. A more favorable comparison would be against the optimized settings used for the mass production of background simulation for IC40 and IC59 cascade analyses. With an energy threshold of 30 TeV/nucleon and an $E^{-2.6}$ spectrum, 344450 out of 1 million showers reach the detector.

Table A.1: Time required to put 1 million single muons at the sampling surface for two configurations of **CORSIKA** as well as MuonGun.

Step	5-comp CORSIKA	“Cascade-opt.”	MuonGun
Generate input muons	16610 s	11116 s	21 s
Read air shower file	861 s	203 s	N/A
Propagate muons	3101 s	1335 s	260 s
Total	5.7 hours	3.5 hours	5 minutes
MuonGun speed-up	73	45	1

A.7 Open issues

The variables chosen for the parameterization of the underground muon flux in Section A.2.2 are sufficient for single muons. For muon bundles, however, there is some loss of information, since the distribution of the energies and transverse displacements of N muons (i.e. $2N$ parameters) must be captured in only 5 dimensions. Following [134], we have assumed that the joint energy and radius distribution of the muon bundle is a function of its multiplicity only. While this is nearly true, it does not quite capture the full distribution.

We can see this by parameterizing the muon flux produced by **CORSIKA** as in Section A.5 and using it to calculate weights for **CORSIKA** events based on the muon bundle configuration at depth, as in Section A.4.3. If the parameterization correctly describes the muon flux in all its details, then the weights calculated this way will be asymptotically equivalent to weights calculated from the primary in each shower. Figure A.9 shows how the bundle-based weights are equivalent to the primary-based weights for single muons, but systematically too small for high-multiplicity bundles. The disagreement becomes worse with increasing energy. This is unfortunate, but not terribly surprising. Because the model neglects hidden correlations (e.g. between the energies of all the muons in bundles from a high-energy shower), many generated bundle configurations appear either more or less unlikely than they should, leading to a weight that is too small or too large, respectively.

Since this problem does not affect the single-muon events that are the overwhelming majority of the background for veto-based event selections, its solution is outside the scope of this work. While it is unlikely that the distortion of a bundle distribution generated from the parameterization would be noticeable in observable distributions, the lack of a complete description does make it impossible to combine MuonGun simulation with the large body of **CORSIKA** simulation that has already been produced in IceCube in the general case. Fixing this deficiency will involve adding at least one dimension to the parameterization, most likely the energy of the lead muon or the total energy in the bundle.

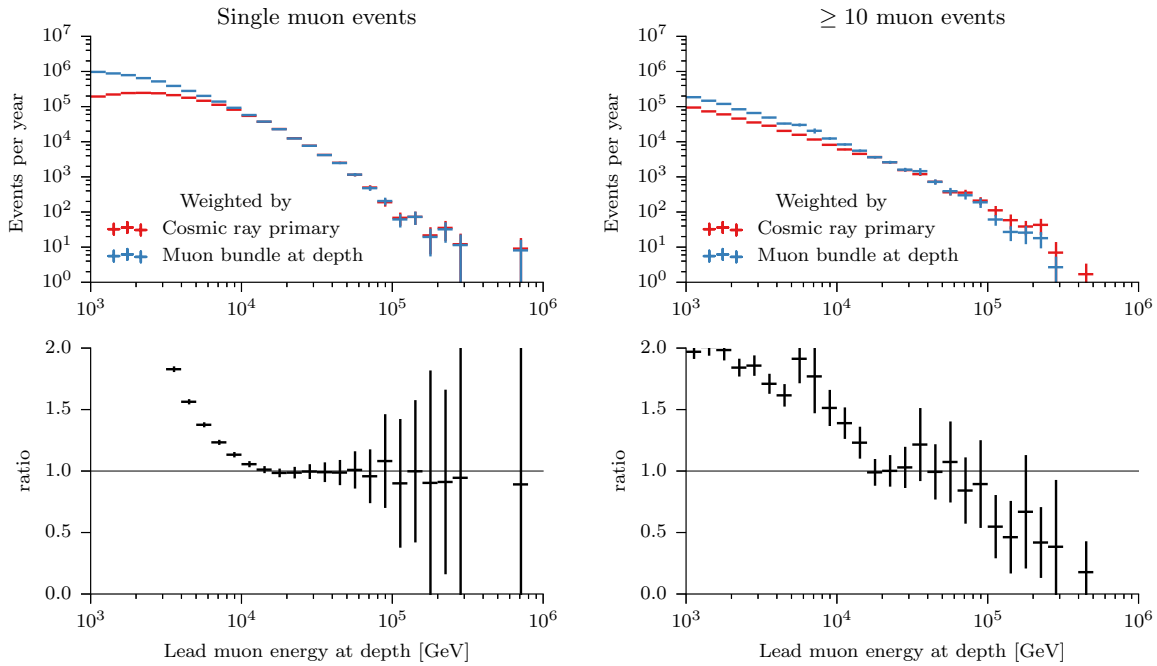


Figure A.9: An illustration of deficiencies in the muon bundle parameterization. The upper panel in each sub-figure shows the spectrum of the highest-energy muon in each event from a high-threshold (30 TeV/nucleon) **CORSIKA** air shower simulation (that passed the offline pre-selection described in Section 4.1.2. For the red histograms, the weight was calculated in the usual fashion, based on the primary of each shower. For the blue histograms, the weight was calculated from based on the configuration of the muon bundle at depth as in Section A.4.3. The lower panel shows the ratio of the two spectra. Below 10 TeV, the weighting scheme based on the cosmic ray primary underestimates the rate because of missing simulated primaries with < 30 TeV/nucleon. Above 10 TeV, both weighting schemes are equivalent for single muons. For muon bundles, however, the weights based on the muon bundle are systematically smaller at high energies than the weights based on the cosmic ray primary.

APPENDIX

The 30 TeV down-going excess

The excess of events observed in the down-going region with deposited energies around 30 TeV warrants some further investigation. An excess of this sort could be explained by a routine statistical fluctuation, systematic errors in the estimated penetrating muon background, atmospheric neutrino veto probability, modeling of the neutrino acceptance of the detector (e.g. through the optical properties of the ice); or lastly, an astrophysical neutrino flux that does not follow a single, isotropic power law.

While the excess around 30 TeV appears quite significant by itself, the probability of obtaining such an excess anywhere in the histograms used for this analysis can actually be quite large, a phenomenon commonly known as the “look elsewhere” effect [145]. To control for this effect, we can ask a simple question: if true underlying means matched the best fit exactly, how often would we observe a mismatch between the data and the model that is more extreme than the one observed? As a measure of the mismatch we use a saturated Poisson test statistic [146],

$$-2 \ln L_{\text{sat}} = 2 \sum_{i=1}^N \left(\lambda_i - n_i + n_i \ln \frac{\lambda_i}{n_i} \right), \quad (\text{B.1})$$

where λ_i is the mean the best-fit model predicts in bin i and n_i is the count observed in that bin. Since the contribution of a term in the Poisson likelihood function (5.29) is maximized when the predicted mean is equal to the observed counts, the test statistic compares the specific model prediction to the best possible model for the observed data. If the sample size is infinite, then (B.1) is distributed like a χ^2 with N degrees of freedom, but for smaller data sets like the one presented here, its empirical distribution must be derived from repeated sampling from the assumed underlying distribution [146]. For this sampling, we use the same binning scheme as for the analysis, but combine the three bins around the 30 TeV bin into one to account for the observed correlation. Figure B.1 shows the distribution of the test statistic obtained from 10000 samples from the best-fit hypothesis. Of the sampled realizations, 4.7% yielded a test statistic greater than

the one observed in the experimental data. From this, we can conclude that the excess is compatible with a statistical fluctuation.




Figure B.1: Distribution of the test statistic (B.1) on the southern-sky energy spectrum obtained from 10000 samples from the best-fit hypothesis. 4.7% of realizations produced a test statistic more extreme than the one observed in data.

Even though the excess is not significant, it is useful to examine sources of systematic error that could be responsible for a distortion in the average energy spectrum like the one observed in data.

First, we can explore the possibility that the excess is due to un-modeled penetrating muon background. This is not a viable explanation. While penetrating muons do contribute to southern-sky spectrum around 30 TeV, the predicted rate of penetrating muon events in that region is 10 times too low to explain the excess events. The estimates of both the penetrating muon component’s overall normalization and of its shape are robust. The normalization is determined from the fit, and is influenced primarily by the southern-sky event rate around 2 TeV deposited energy, where the contribution from penetrating muons is comparable to the contribution from atmospheric neutrinos. The shape is fixed, but matches the shape of the observed muon distribution. This can be verified by isolating a sample of known penetrating muons that are otherwise similar to the events in the neutrino sample. This was done by selecting events that fail the outer-layer veto (cf. Section 4.2.2). The outer-layer and incoming-track vetoes (cf. Section 4.2.3) were then re-applied in a reduced geometry that excluded the outer layer of the DOMs, and was therefore not influenced by the hits on the outer layer. The same selection was applied to experimental data and simulated penetrating muons. While only 63 events from the experimental data survive this selection, the simulation predicts their number of collected-charge distribution correctly as shown in Figure B.2.

Similarly, the excess is unlikely to be caused by an error in the calculation of the fraction of atmospheric neutrinos that are accompanied by muons and thus excluded from the data sample. This can be demonstrated




Figure B.2: Control sample of penetrating muon events that fail the outer-layer veto but pass vetoes in the inner detector. The MC simulation correctly predicts the normalization and shape of the collected-charge distribution.

by varying the neutrino-energy dependence of the self-veto probability in ways that could transform the expected falling energy spectrum into one with a peak at 30 TeV. Figure B.3 shows 3 such modifications, none of which were able to create such a peak in the average spectrum.

Another possibility is a systematic error in the neutrino acceptance calculated from simulation. One possible source of error is the model of photon transport in ice assumed in the simulation. In order to test whether such an effect could produce the observed feature, we can compare simulations that use the nominal model with simulations that use modified models that are still allowed by the fit described in [91]. The 3 chosen variants were

1. “more absorption:” all absorption coefficients increased by 10%
2. “more scattering:” all scattering coefficients increased by 10%
3. “less absorption, less scattering:” all scattering and absorption coefficients reduced by 7.1%

Figure B.4 shows the southern-sky deposited energy spectra calculated from the nominal simulation as well as the 3 variants. The most prominent features are the large statistical fluctuations in the variant simulations; because of resource constraints these each contain only 1/10 the number of simulated events as the nominal set. Nonetheless, no peak appears in the deposited-energy spectrum. From this we can only conclude that it is unlikely that a uniform systematic error in the photon transport model could be responsible for a feature that appears in the energy spectrum in data but not in simulation.

Having examined the distortion in the deposited-energy distribution in detail, we can move on to examine other properties of the events that make up the excess. Two interesting properties are the distance from the reconstructed vertex to the nearest IceCube string and to the border of the detector. Real neutrino interactions should be spread uniformly around the nearest string, while detector artifacts like sparking in the PMT base (a phenomenon observed in AMANDA, but never in IceCube) would lead to vertices clustered within a few meters of the string. Figure B.5 shows the cumulative distributions of the string distance for the observed events and for neutrino simulation. We can gauge whether the observed distribution is a realization of the one from simulation using a Kolmogorov-Smirnov test [147]. The observed deviation between the cumulative distributions would be expected in 14% of experiments, so the mild clustering around the nearest string is again not significant. The distance to the detector border is another useful quantity: if the excess is due to un-modeled penetrating muon background, then the excess events should cluster near the border rather than deep inside the detector. Figure B.6 shows the cumulative distributions in the same scheme as

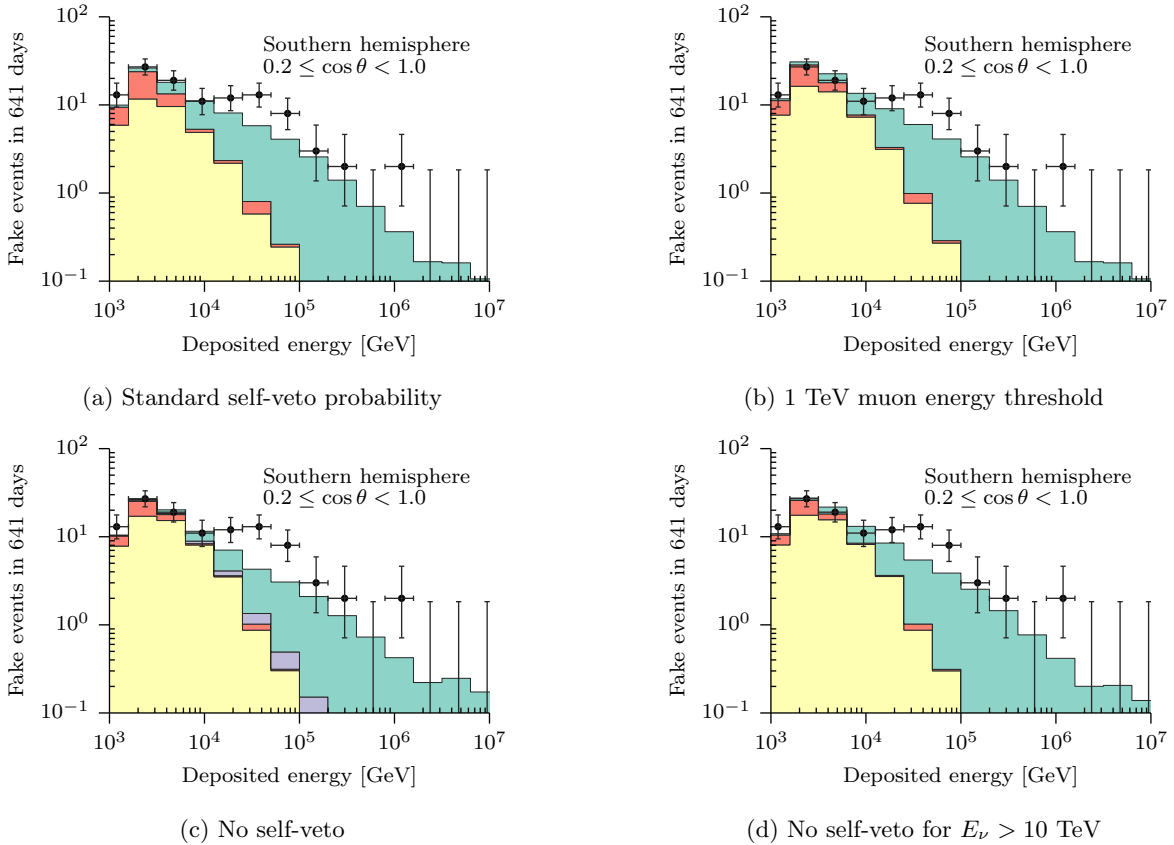


Figure B.3: Effect of variations in the self-veto probability calculation on the southern-sky energy spectrum. (a) shows the nominal calculation, assuming that 100 GeV muons can be rejected in events depositing fewer than 6000 PE, and 1 TeV muons above 6000 PE. (b) uses a 1 TeV muon energy threshold everywhere, which over-predicts the event rate at low energies. (c) assumes that atmospheric neutrinos can never be vetoed, and requires that the conventional atmospheric neutrino flux be reduced by nearly 10% to remain compatible with the event rate at low energies. (d) assumes that atmospheric neutrinos with energies greater than 10 TeV can never be vetoed. None of these modifications are able to produce a bump in the average spectrum.

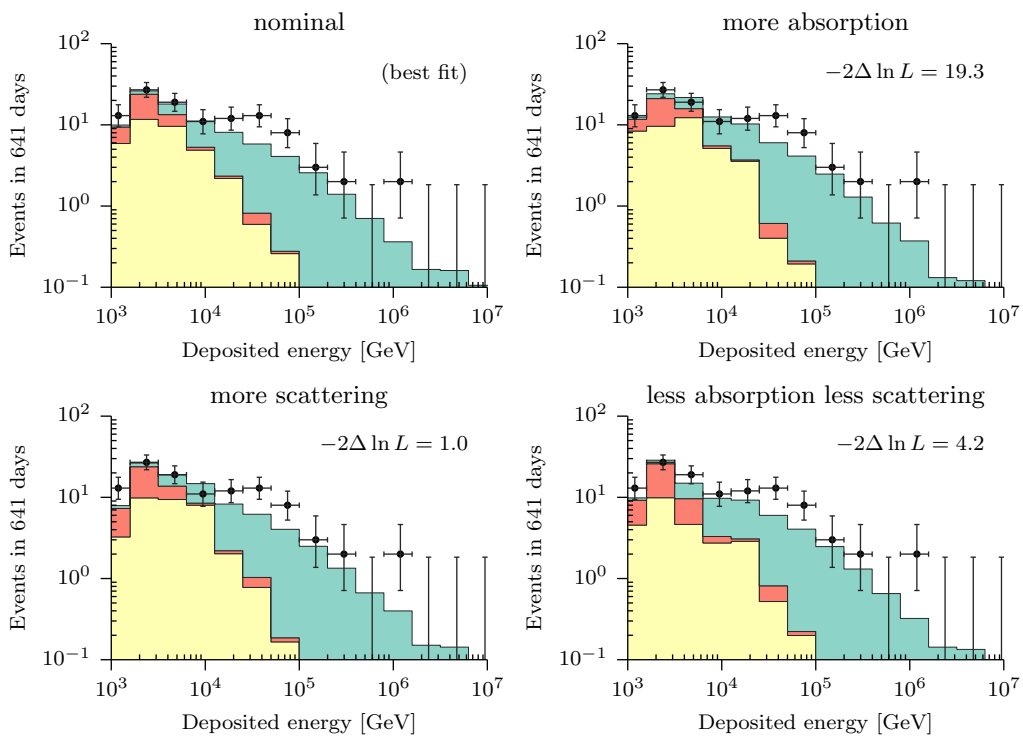


Figure B.4: Effect of systematic variations in the modeling of the optical properties of the glacial ice on the southern-sky energy spectrum. None of the variations are able to produce a non-statistical bump in the average spectrum.

Figure B.5. In this case the distributions are quite compatible; the observed deviation is expected in 53% of experiments.




Figure B.5: Cumulative distribution of the distance from the reconstructed vertex to the nearest IceCube string for the 9 events observed from southern sky in 2011–2012, and for the ensemble of simulated neutrino events expected over the same period. The mismatch between the two is expected in 14% of experiments with this sample size.

Up to this point we have only examined the time-integrated sample. It is also useful to examine the distribution of event times for deviations from uniformity that, if significant, would indicate strong contributions from a single astrophysical source in a flaring state. Figure B.7 shows the number of events observed from the southern sky between 15 and 50 TeV per 90 days of detector livetime from 2010–2012. No significant deviation from the average rate was observed.

Finally, as an exercise, we can construct a model of an extraterrestrial neutrino flux that would explain the excess. The two most prominent features are that it is concentrated around 30 TeV and near $\cos \theta \sim 0.5$. The structure of the excess in energy can be described by a hard energy spectrum with a cutoff, for example

$$\Phi = \Phi_0 \left(\frac{E_\nu}{E_0} \right)^{-\gamma} \exp[-E_\nu/E_b] \quad , \quad (\text{B.2})$$

with $\gamma = 1$ and $E_b = 30$ TeV. The excess in the zenith angle distribution coincides roughly with the declination of the galactic center, and a relatively narrow angular distribution around the galactic center (for example, a Fisher distribution [148] with $\kappa = 100$) would lead to events concentrated at a local zenith angle of $\cos \theta \sim 0.5$, as illustrated in Figure B.8.

When fit to the data, this ad-hoc model with the parameters shown in Table B.1 is preferred over the isotropic extraterrestrial fit shown Table 6.1 by 3.7σ . It should be noted, however, that this model has no




Figure B.6: Cumulative distribution of the distance from the reconstructed vertex to the border of the detector the 9 events observed from southern sky in 2011–2012, and for the ensemble of simulated neutrino events expected over the same period. The mismatch between the two is expected in 53% of experiments with this sample size.




Figure B.7: Arrival times of neutrino events in the vicinity of the 30 TeV excess. The error bars on each point are $1/\sqrt{N}$. There is no significant deviation from the average rate.




Figure B.8: Effect of a source at the galactic center on the observed zenith angle distribution. The upper left panel shows of 30 events drawn from a Fisher angular distribution [148] with $\kappa = 100$ centered at the galactic core, in galactic coordinates. The upper right panel shows the same distribution in equatorial coordinates. In both of these panels, the galactic plane is marked with a blue line. The lower histogram shows the corresponding zenith angle distribution in IceCube detector coordinates.

a priori motivation, and was constructed to describe the observed data, as shown in Figures B.9 and B.10. It is not surprising that this hand-tailored model fits the data quite well.

Table B.1: Best fit parameters for a galactic-center model, defined as in Table 6.1. The ad-hoc “galactic-center” flux is defined in Equation (B.2).

Parameter	Best-fit value
Penetrating μ flux	$1.67 \pm 0.47 \Phi_{\text{SIBYLL+DPMJET}}$
Conventional ν flux	$1.07 \pm 0.07 \Phi_{\text{HKKMS}}$
Prompt ν flux	$< 2.28 \Phi_{\text{ERS}} \text{ (90\% CL)}$
Astrophysical Φ_0	$1.10^{+0.47}_{-0.45} \times 10^{-18} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$
Astrophysical γ	$2.22^{+0.26}_{-0.22}$
Galactic center Φ_0	304^{+106}_{-96}

If this model is correct, then many of the excess events should come from the direction of the galactic center. Figure B.11 shows the arrival directions of events between 15 and 50 TeV in galactic coordinates. There are no indications of significant clustering at the galactic center. This is at best inconclusive, given the moderate angular resolution for the cascade events that make up the majority of the sample. The resolution of the reconstructed azimuthal angle (which is not used in the primary analysis) is significantly worse as shown in Figures B.12 and B.13; this is primarily due to the azimuthal symmetry of the DOMs and their large horizontal spacing (more than 1 effective optical scattering length everywhere except in DeepCore). Despite the coincidence between the declinations of the excess events and the galactic center, it will be quite difficult to search for concentrated sources of neutrinos with an event sample like the one presented here until the angular resolution (in particular, the azimuthal resolution) is improved significantly.

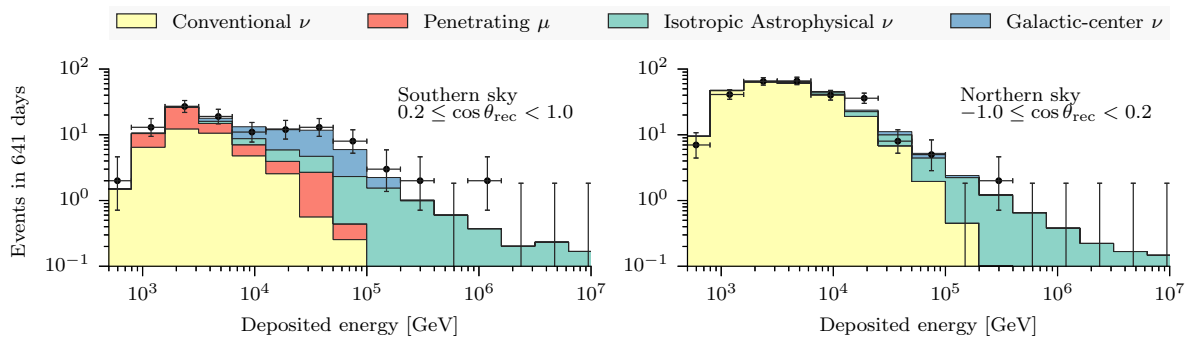


Figure B.9: Energy spectra fit with a galactic-center component.

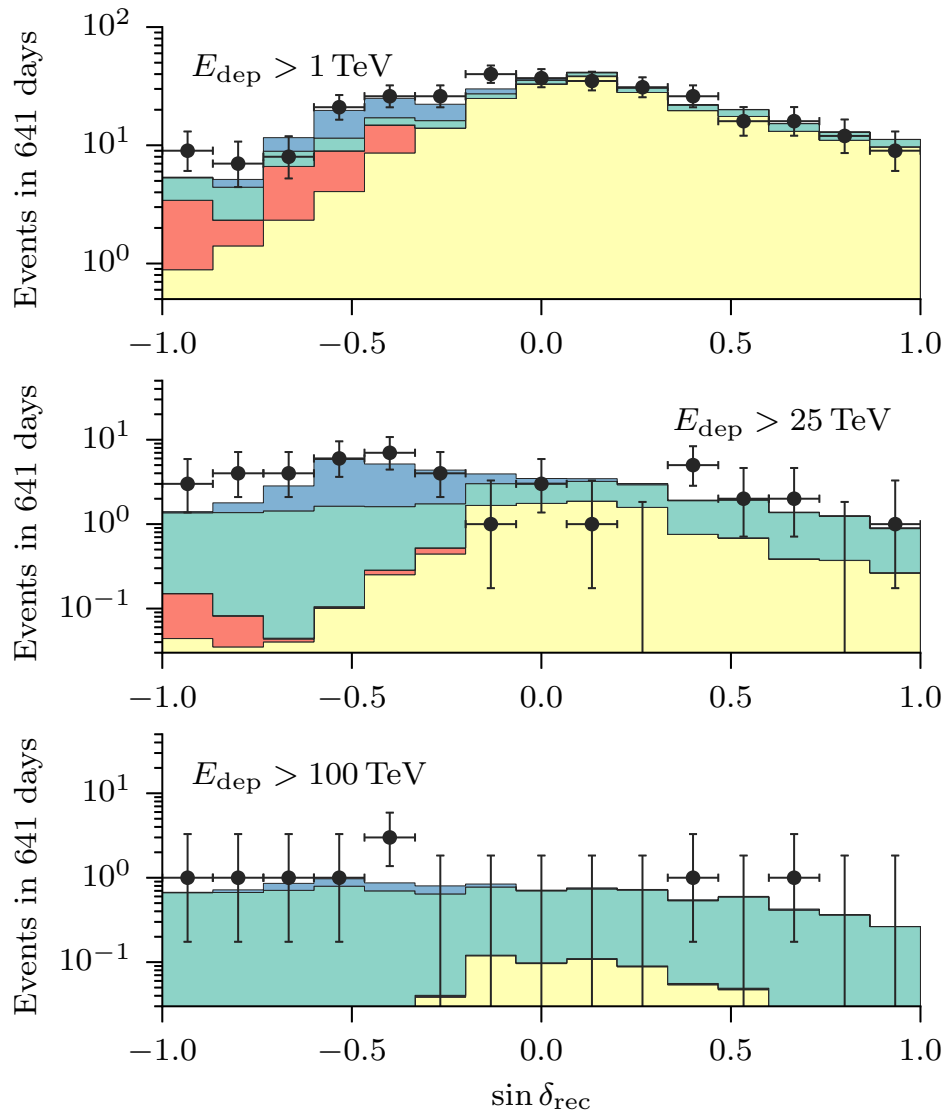


Figure B.10: Zenith angle distributions fit with a galactic-center component. The color scheme is the same as in Figure B.9.







Figure B.11: Arrival directions of neutrino events in the vicinity of the 30 TeV excess in galactic coordinates and Mollweide projection. Events from 2010–2011 are shown in red and events from 2011–2012 are shown in blue. There is no obvious clustering in the arrival directions.




(a) All events classified as cascades


(b) All events classified as tracks, with directions reconstructed with **Monopod**.

(c) All events classified as tracks, with directions reconstructed using the simple outgoing-track search described in Section 5.5.2.

Figure B.12: Azimuth angle reconstruction error for simulated neutrino events that pass the final event selection. The grey bands in each panel show the ranges of azimuth angle error that contain the given fraction of events. ((a)) shows events classified as cascades, while ((b)) and ((c)) show events classified as tracks. For starting track events the large displacement of PE detections from the neutrino vertex along the track direction provides a much better estimate of the neutrino direction than the cascade reconstruction.



(a) Events classified as cascades



(b) Events classified as tracks

Figure B.13: Quantiles of the absolute azimuth angle difference distribution (cf. Figure B.12) as a function of deposited energy. The median error is nearly constant with energy, while the rate of rare failures increases with increasing deposited energy, especially for track events. The azimuth angle resolution for cascade events is much worse than the zenith angle resolution because of the detector geometry.

APPENDIX

Glossary

AMANDA Antarctic Muon And Neutrino Detector Array, IceCube’s predecessor neutrino telescope at the South Pole. AMANDA construction began in 1996, and it was decommissioned in 2009 [149]. [143](#), [151](#), [153](#)

ANTARES Astronomy with a Neutrino Telescope and Abyss environmental RESearch, a water-Cherenkov neutrino detector deployed on the bed of the Mediterranean Sea near Toulon. The instrumented volume of ANTARES is approximately 100 times smaller than that of IceCube [150]. [32](#), [121](#), [143](#)

ATWD Analog Transient Waveform Digitizer, the fast 10-bit digitizer on the mainboard of the **DOM**. The ATWD has 4 channels. Channels 0–2 are used for waveform readout, and have preamplifier gain factors of 16, 2, and 1/4 respectively, for an effective dynamic range of 16 bits. At the typical 10^7 operating **PMT** gain, a single-**PE** pulse is 40 counts high. Each channel consists of a bank of 128 sampling capacitors; when the discriminator fires, the input signal is rippled across the capacitor bank at ~ 300 MHz for a total sampling time of approximately 425 nanoseconds. The voltage on each capacitor is then digitized in sequence, which takes 29 microseconds per channel [151]. The DOM mainboard has two ATWD chips configured in “ping-pong” mode; when one ATWD shuts off, the remaining ATWD can be ready to digitize in 50 ns. [143](#), [152](#)

cascade filter The online filter specifically tuned for cascade-like events. [46](#), [47](#), [143](#), *see* [online](#) & [online filter](#)

CC “Charged-current,” **DIS** by exchange of a W boson, where the incoming neutrino is transformed into a charged lepton of the same flavor. [20](#), [64](#), [94](#), [96](#), [108](#), [119](#), [143](#)

center of gravity The average position of hit DOMs, each weighted by the photoelectron charge it collects. [36](#), [143](#), *see* [tensor of inertia](#)

CORSIKA COsmic Ray SIMulations for KAscade, the most widely-used Monte Carlo simulation program for extensive air showers [152]. 77, 79–81, 84–87, 89, 91, 121–123, 130–134, 143, 151

DAQ data acquisition. 31, 143, 151, 154

dCORSIKA A version of **CORSIKA** optimized for penetrating muon background simulation in **AMANDA** and still used for centrally-produced simulation in IceCube. While many of the modifications were integrated upstream, some differences to mainline **CORSIKA** remain. The most important of these are an option to generate “natural rate” simulation by drawing cosmic ray primaries from the spectrum of [144] instead of a single power law for a single chemical element and an option to skip the simulation of showers that have too little energy to produce a muon capable of reaching the depth where **AMANDA** was deployed. See: [153, Chapter 2]. 121, 122, 132, 143

DeepCore IceCube’s low-energy infill array, more densely spaced and instrumented with **PMTs** with $\sim 35\%$ higher quantum efficiency than the standard IceCube **PMT**. 43, 48, 59, 94, 143, 152

degrees of freedom Number of data points minus number of fitted parameters. 143, *see reduced log-likelihood*

DIS Deep inelastic scattering is a weak interaction of neutrinos with a target nucleus where enough momentum is transferred to break the nucleus apart and create jets. 20, 143, 150, 153

DOM Digital Optical Module. 31, 35, 40, 42, 48, 50, 51, 54, 59, 104, 106, 138, 143, 150–152, 154

ERS Enberg, Reno, and Sarcevic. 5, 143

F2k Format 2000, the ASCII file format used for **AMANDA** data processing. Vestiges of F2k are left in the parts of the IceCube simulation software that were adopted directly from **AMANDA**. 121, 143

fADC Fast analog-to-digital converter, the 10-bit pipelined ADC on the mainboard of the **DOM**. It samples constantly at 40 MHz. 143, 152

fill ratio Fraction of **DOMs** within some distance of a reconstructed cascade vertex that trigger in an event [99]. The search radius is proportional to the average distance of all triggered **DOMs** from the vertex. Cascade events typically have fill ratios close to 1, while extended tracks have fill ratios close to 0. 48, 143

- frame** The top-level data-container in IceTray. Each frame is a collection of key-value pairs holding data associated with a single global trigger or a subset thereof. [32](#), [143](#), [153](#)
- global trigger** In the IceCube DAQ system, the union of the [readout windows](#) of triggers. [31](#), [32](#), [143](#)
- GR** “Glashow Resonance,” resonant $\bar{\nu}_e + e^-$ scattering near lab-frame energies of 6.3 PeV (80 GeV center-of-mass). The intermediate W^- can decay leptonically or hadronically. [96](#), [143](#)
- hit** Detection by a [DOM](#) of one or more photons. [48](#), [49](#), [143](#), *see* [hit record](#)
- hit record** Basic element of IceCube data assembled by the software running on the [DOM](#) mainboard transmitted to the surface after the PMT discriminator fires, consisting of a [DOM](#) identifier, timestamp, and 1 or more digitized waveforms. The records comes in two varieties, SLC and HLC. [143](#), *see* [HLC](#) & [HLC](#)
- HLC** Hard local coincidence, the [DOM](#) readout mode used when the discriminator attached to the [PMT](#) output of one [DOM](#) fires within 1 microsecond of one of its nearest or next-to-nearest neighbors on the same [string](#). If this condition is satisfied, then 128 samples from each of the $N + 1$ channels of the [Analog Transient Waveform Digitizer \(ATWD\)](#) and 256 samples from the [fast analog-to-digital converter \(fADC\)](#) are transmitted to the surface, where N is the number of [ATWD](#) channels that are nearly saturated, for a total sampling time of 6.4 microseconds. [31](#), [48](#), [50](#), [59](#), [94](#), [143](#), [152](#)
- homogenized total charge** Total charge of all [HLC pulses](#) detected on non-DeepCore [DOMs](#) that do not contribute more than 50% of the total. This quantity has a smaller variance with respect to deposited energy than a naïve sum over pulse charges. [59](#), [143](#)
- IC79** The nearly complete configuration of IceCube, run from May 2010 to May 2011 with 79 strings. [43](#), [47](#), [48](#), [50](#), [104](#), [106](#), [107](#), [143](#), [154](#)
- IC86** The first complete configuration of IceCube, run from May 2011 to May 2012, with 86 strings. Configurations run after May 2012 also had 86 strings, but are generally referred to by year rather than by configuration. [43](#), [47](#), [48](#), [50](#), [104](#), [106](#), [143](#), [154](#)
- IceTray** IceCube’s modular data-processing framework. [32](#), [143](#), [153](#)

- line fit** A first-guess track reconstruction algorithm that treats the hit pattern as a plane wave moving through the detector with speed v_{LineFit} . 35, 143
- line fit speed** The speed of a plane wave moving through the detector that best describes the observed hit pattern. Cascade-like events have line fit speeds close to 0, while extended tracks have line fit speeds close to c . 35, 47, 50, 143, *see line fit*
- local coincidence** correlated PMT discriminator crossings in neighboring or next-to-neighboring DOMs on the same string. 54, 143, *see HLC*
- log-likelihood** The natural logarithm of a likelihood function. This usually refers to the value of $-\ln L$ at its minimum. 38, 47, 143
- MMC** Muon Monte Carlo, a Monte Carlo simulation of muon energy losses in matter, developed for **AMANDA** and implemented in Java [72]. 86, 121, 143, 154
- module** The basic processing unit in **IceTray**. The modules are arranged in a directed graph, where each module receives **frames** from the previous module in the chain, manipulates it, and pushes it to the next module. Modules may add or delete objects from the **frame**, or discard the frame entirely to drop an event from the selection. 32, 143
- Monopod** A maximum-likelihood reconstruction of the vertex, time, direction, and energy of a single cascade. This is the single-source specialization of the more general Millipede reconstruction described in [87, Section 4.4], which itself is a generalization of the earlier Credo reconstruction [98]. 42, 54, 59, 94, 97, 143, 148
- MPE** “Multi-photoelectron,” a formulation of the timing likelihood that compares the time residual of each pulse to the expected distribution of the time residual of the first photon, given N observed. This formulation is valid when the likelihood is calculated using only the first pulse on each DOM, and approaches a delta function at $\Delta t = 0$ for large N [88]. 38, 143, *see Pandel & SPE*
- NC** “Neutral-current,” **DIS** by exchange of a Z boson, where the incoming neutrino remains a neutrino. 20, 96, 104, 108, 143

offline In the context of IceCube data handling, all stages of processing done on pre-filtered data after they have been transmitted from the South Pole to the IceCube data center in Madison, WI. The delay between data-taking and processing can be anywhere from weeks to years. [46](#), [47](#), [143](#), *see online*

online In the context of **PnF**, processed in nearly real time at the South Pole. [32](#), [46](#), [47](#), [143](#)

online filter Any event selection run as part of **PnF**. Only the events that pass one of the online filters are transmitted over the satellite; events that fail the filters can only be recovered from magnetic tapes shipped from the Pole once a year. [32](#), [50](#), [94](#), [143](#), *see online*

Pandel A family of timing-based likelihood functions based loosely on the analytic approximation to the time-residual distribution derived in [94]. [38](#), [47](#), [48](#), [143](#)

PE photoelectron. [31](#), [32](#), [40](#), [42](#), [50](#), [51](#), [94](#), [98](#), [106](#), [107](#), [143](#), [148](#), [150](#)

PMT photomultiplier tube. [31](#), [40](#), [42](#), [94](#), [143](#), [150–152](#), [154](#)

PnF Processing and Filtering. [31](#), [32](#), [143](#), [153](#)

profile likelihood Likelihood function formed by maximizing out nuisance parameters [137]. [106](#), [143](#)

PROPOSAL PPropagator with Optimal Precision and Optimized Speed for All Leptons, a C++ translation of **MMC** [73]. [73](#), [122](#), [143](#)

pulse A photoelectron pulse reconstructed from the digitized **PMT** signal, defined as any basis function in the waveform unfolding described in [87, Section 3.4.1] with non-zero amplitude. Each pulse has a time, amplitude, and width. The leading-edge time (or simply, time) of the pulse is the time when the photoelectron was ejected from the photocathode, and the width is the time delay between the leading edge of the basis function and the next later one. The amplitude is the best-fit number of photons detected during the width of the basis function. [32](#), [35](#), [50](#), [51](#), [54](#), [59](#), [94](#), [143](#), [152](#)

readout window In the IceCube **DAQ** system, the period surrounding a trigger when all **DOMs** are read out. For **SMT-8** in **IC79** and **IC86**, this was $-4/+6$ microseconds from the start and end of the trigger. [31](#), [32](#), [143](#), [151](#)

reduced log-likelihood log-likelihood divided by number of degrees of freedom: $-\ln L/n_{\text{dof}}$. 38, 47, 48, 143, *see* log-likelihood & degrees of freedom

SMT simple multiplicity trigger. 31, 143

SMT-8 simple multiplicity trigger with multiplicity 8. 31, 143, 154

SPE “Single-photoelectron,” a formulation of the timing likelihood that compares the time residual of each pulse to the expected distribution for all pulses. Since the Pandel timing likelihoods are typically calculated using only the first pulse in each DOM, the likelihood is only valid when there is only one photoelectron, hence the name [88]. 38, 47, 143, *see* Pandel & MPE

string The 2.5 km long cable providing power and surface communication to 60 DOMs. The term usually refers to the set of DOMs hosted on the cable rather than to the cable itself. The full IceCube configuration has 86 strings. 43, 44, 47, 48, 50, 51, 143, 152

tensor of inertia A first-guess directional reconstruction algorithm that treats the hit pattern as a rigid body with “masses” given by the total photoelectron charge collected by each DOM. 36, 47, 143

tensor-of-inertia eigenvalue ratio The size of the smallest eigenvalue of the tensor of inertia relative to the sum of all 3. For perfectly spherical events this is close to 1/3, while for an exact line of hits it is 0. 36, 47, 50, 143, *see* tensor of inertia

time residual Difference between photon detection time and the earliest possible detection time given a cascade or track hypothesis: $\Delta t \equiv t - t_{\text{geo}}$. 38, 42, 54, 143

APPENDIX

Bibliography

- [1] Viktor F. Hess. Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Physikalische Zeitschrift*, 13:1084–1091, 1912. URL <http://www.mpi-hd.mpg.de/hfm/HESS/public/HessArticle.pdf>.
- [2] J. Beringer et al. (Particle Data Group). The Review of Particle Physics. *Phys. Rev. D*, 86:010001, 2012.
- [3] R. Abbasi et al. (IceCube Collaboration). Observation of Anisotropy in the Arrival Directions of Galactic Cosmic Rays at Multiple Angular Scales with IceCube. *Astrophys. J.*, 740:16, 2011. doi: [10.1088/0004-637X/740/1/16](https://doi.org/10.1088/0004-637X/740/1/16), arXiv: [1105.2326](https://arxiv.org/abs/1105.2326).
- [4] Pierre Auger, P. Ehrenfest, R. Maze, J. Daudin, and Robley A. Fréon. Extensive cosmic-ray showers. *Rev. Mod. Phys.*, 11:288–291, Jul 1939. doi: [10.1103/RevModPhys.11.288](https://doi.org/10.1103/RevModPhys.11.288).
- [5] James J. Beatty and Stefan Westerhoff. The Highest-Energy Cosmic Rays. *Annual Review of Nuclear and Particle Science*, 59:319–345, 2009. ISSN 0163-8998. doi: [10.1146/annurev.nucl.58.110707.171154](https://doi.org/10.1146/annurev.nucl.58.110707.171154).
- [6] Julia K. Becker. High-energy neutrinos in the context of multimessenger astrophysics. *Physics Reports*, 458(4–5):173 – 246, 2008. ISSN 0370-1573. doi: [10.1016/j.physrep.2007.10.006](https://doi.org/10.1016/j.physrep.2007.10.006).
- [7] R. U. Abbasi, M. Abe, T. Abu-Zayyad, et al. Indications of Intermediate-scale Anisotropy of Cosmic Rays with Energy Greater Than 57 EeV in the Northern Sky Measured with the Surface Detector of the Telescope Array Experiment. *ApJ. Lett.*, 790:L21, August 2014. doi: [10.1088/2041-8205/790/2/L21](https://doi.org/10.1088/2041-8205/790/2/L21).
- [8] Eli Waxman and John Bahcall. High energy neutrinos from astrophysical sources: An upper bound. *Phys. Rev. D*, 59:023002, Dec 1998. doi: [10.1103/PhysRevD.59.023002](https://doi.org/10.1103/PhysRevD.59.023002).

- [9] Thomas K. Gaisser, Francis Halzen, and Todor Stanev. Particle astrophysics with high energy neutrinos. *Physics Reports*, 258(3):173 – 236, 1995. ISSN 0370-1573. doi: [10.1016/0370-1573\(95\)00003-Y](https://doi.org/10.1016/0370-1573(95)00003-Y).
- [10] John G. Learned and Karl Mannheim. High-energy neutrino astrophysics. *Annual Review of Nuclear and Particle Science*, 50(1):679–749, 2000. doi: [10.1146/annurev.nucl.50.1.679](https://doi.org/10.1146/annurev.nucl.50.1.679).
- [11] M. Aartsen et al. (IceCube Collaboration). Searches for Extended and Point-like Neutrino Sources with Four Years of IceCube Data. 2014, arXiv: [1406.6757](https://arxiv.org/abs/1406.6757).
- [12] M. Aartsen et al. (IceCube Collaboration). Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector. *Science*, 342(6161), 2013. doi: [10.1126/science.1242856](https://doi.org/10.1126/science.1242856).
- [13] M. Aartsen et al. (IceCube Collaboration). Observation of high-energy astrophysical neutrinos in three years of icecube data. *Phys. Rev. Lett.*, 113:101101, Sep 2014. doi: [10.1103/PhysRevLett.113.101101](https://doi.org/10.1103/PhysRevLett.113.101101).
- [14] Rikard Enberg, Mary Hall Reno, and Ina Sarcevic. Prompt neutrino fluxes from atmospheric charm. *Phys. Rev. D*, 78(4):043005, Aug 2008. doi: [10.1103/PhysRevD.78.043005](https://doi.org/10.1103/PhysRevD.78.043005).
- [15] Nathan Whitehorn, Jakob van Santen, and Sven Lafebre. Penalized splines for smooth representation of high-dimensional Monte Carlo datasets. *Computer Physics Communications*, 184(9):2214 – 2220, 2013. ISSN 0010-4655. doi: [10.1016/j.cpc.2013.04.008](https://doi.org/10.1016/j.cpc.2013.04.008).
- [16] M. Aartsen et al. (IceCube Collaboration). Energy reconstruction methods in the IceCube neutrino telescope. *Journal of Instrumentation*, 9(03):P03009, 2014. doi: [10.1088/1748-0221/9/03/P03009](https://doi.org/10.1088/1748-0221/9/03/P03009).
- [17] Thomas K. Gaisser, Kyle Jero, Albrecht Karle, and Jakob van Santen. Generalized self-veto probability for atmospheric neutrinos. *Phys. Rev. D*, 90:023009, Jul 2014. doi: [10.1103/PhysRevD.90.023009](https://doi.org/10.1103/PhysRevD.90.023009), arXiv: [1405.0525](https://arxiv.org/abs/1405.0525).
- [18] M. Aartsen et al. (IceCube Collaboration). Atmospheric and Astrophysical Neutrinos above 1 TeV Interacting in IceCube. 2014, arXiv: [1410.1749](https://arxiv.org/abs/1410.1749). Accepted for Phys. Rev. D.
- [19] B. Louis, V. Sandberg, G. Garvey, et al. The Evidence for Oscillations. *Los Alamos Science*, 25:16–27, 1997.

- [20] A. E. Brenner, D. C. Carey, J. E. Elias, et al. Experimental study of single-particle inclusive hadron scattering and associated multiplicities. *Phys. Rev. D*, 26:1497–1553, Oct 1982. doi: [10.1103/PhysRevD.26.1497](https://doi.org/10.1103/PhysRevD.26.1497).
- [21] Thomas K. Gaisser. *Cosmic rays and particle physics*. Cambridge University Press, Cambridge and New York, 1990.
- [22] Francis Halzen and Alan Martin. *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. John Wiley & Sons, New York, 1984.
- [23] Paolo Lipari. Lepton spectra in the earth’s atmosphere. *Astroparticle Physics*, 1(2):195 – 227, 1993. ISSN 0927-6505. doi: [10.1016/0927-6505\(93\)90022-6](https://doi.org/10.1016/0927-6505(93)90022-6).
- [24] F. Reines, M. F. Crouch, T. L. Jenkins, et al. Evidence for high-energy cosmic-ray neutrino interactions. *Phys. Rev. Lett.*, 15:429–433, Aug 1965. doi: [10.1103/PhysRevLett.15.429](https://doi.org/10.1103/PhysRevLett.15.429).
- [25] C.V. Achar, M.G.K. Menon, V.S. Narasimham, et al. Detection of muons produced by cosmic ray neutrinos deep underground. *Physics Letters*, 18(2):196 – 199, 1965. ISSN 0031-9163. doi: [10.1016/0031-9163\(65\)90712-2](https://doi.org/10.1016/0031-9163(65)90712-2).
- [26] G. Danby, J-M. Gaillard, K. Goulianos, et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962. doi: [10.1103/PhysRevLett.9.36](https://doi.org/10.1103/PhysRevLett.9.36).
- [27] K. Daum, W. Rhode, P. Bareyre, et al. Determination of the atmospheric neutrino spectra with the Fréjus detector. *Zeitschrift für Physik C Particles and Fields*, 66:417–428, 1995. ISSN 0170-9739. doi: [10.1007/BF01556368](https://doi.org/10.1007/BF01556368).
- [28] M.C. Gonzalez-Garcia, M. Maltoni, and J. Rojo. Determination of the atmospheric neutrino fluxes from atmospheric neutrino data. *JHEP*, 0610:075, 2006. doi: [10.1088/1126-6708/2006/10/075](https://doi.org/10.1088/1126-6708/2006/10/075), arXiv: [hep-ph/0607324](https://arxiv.org/abs/hep-ph/0607324).
- [29] G. Giacomelli and A. Margiotta. Macro results on atmospheric neutrinos. *Nucl.Phys.Proc.Suppl.*, 145: 116–119, 2005. doi: [10.1016/j.nuclphysbps.2005.03.047](https://doi.org/10.1016/j.nuclphysbps.2005.03.047), arXiv: [hep-ex/0504029](https://arxiv.org/abs/hep-ex/0504029).
- [30] Fukuda, Y. et al (Super-KamiokaNDE Collaboration). Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998. doi: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).

- [31] T. Kajita. Atmospheric neutrinos. *New J. Phys.*, 6:194, 2004.
- [32] S. Adrian-Martinez et al (ANTARES Collaboration). Measurement of the atmospheric ν_μ energy spectrum from 100 GeV to 200 TeV with the ANTARES telescope. *Eur.Phys.J.*, C73:2606, 2013. doi: [10.1140/epjc/s10052-013-2606-4](https://doi.org/10.1140/epjc/s10052-013-2606-4), arXiv: [1308.1599](https://arxiv.org/abs/1308.1599).
- [33] R. Abbasi et al. (AMANDA and IceCube Collaborations). Determination of the atmospheric neutrino flux and searches for new physics with AMANDA-II. *Phys. Rev. D*, 79:102005, May 2009. doi: [10.1103/PhysRevD.79.102005](https://doi.org/10.1103/PhysRevD.79.102005).
- [34] R. Abbasi et al. (AMANDA and IceCube Collaborations). The energy spectrum of atmospheric neutrinos between 2 and 200 TeV with the AMANDA-II detector. *Astroparticle Physics*, 34(1):48 – 58, 2010. ISSN 0927-6505. doi: [10.1016/j.astropartphys.2010.05.001](https://doi.org/10.1016/j.astropartphys.2010.05.001).
- [35] M. Aartsen et al. (IceCube Collaboration). Measurement of the Atmospheric ν_e Flux in IceCube. *Phys. Rev. Lett.*, 110:151105, Apr 2013. doi: [10.1103/PhysRevLett.110.151105](https://doi.org/10.1103/PhysRevLett.110.151105).
- [36] M. Aartsen et al. (IceCube Collaboration). Search for a diffuse flux of astrophysical muon neutrinos with the IceCube 59-string configuration. *Phys. Rev. D*, 89:062007, 2014, arXiv: [1311.7048](https://arxiv.org/abs/1311.7048).
- [37] Y. M. Andreyev, V. I. Gurentsov, and I. M. Kogai. Muon Intensity from the Baksan Underground Scintillation Telescope. *International Cosmic Ray Conference*, 6:200, 1987.
- [38] Berger, Ch. et al (Fréjus Collaboration). Experimental study of muon bundles observed in the fréjus detector. *Phys. Rev. D*, 40:2163–2171, Oct 1989. doi: [10.1103/PhysRevD.40.2163](https://doi.org/10.1103/PhysRevD.40.2163).
- [39] M. Aglietta et al. (NUSEX Collaboration). Primary cosmic ray composition in the 10^{13} – 10^{17} eV energy range from the analysis of multiple muon events in the NUSEX experiment. *Nuclear Physics B - Proceedings Supplements*, 14(2):193 – 203, 1990. ISSN 0920-5632. doi: [10.1016/0920-5632\(90\)90380-D](https://doi.org/10.1016/0920-5632(90)90380-D).
- [40] Ambrosio, M. et al (MACRO Collaboration). Vertical muon intensity measured with macro at the gran sasso laboratory. *Phys. Rev. D*, 52:3793–3802, Oct 1995. doi: [10.1103/PhysRevD.52.3793](https://doi.org/10.1103/PhysRevD.52.3793).
- [41] M. Aglietta et al. Upper limit on the prompt muon flux derived from the LVD underground experiment. *Phys.Rev.*, D60:112001, 1999. doi: [10.1103/PhysRevD.60.112001](https://doi.org/10.1103/PhysRevD.60.112001), arXiv: [hep-ex/9906021](https://arxiv.org/abs/hep-ex/9906021).

- [42] Aglietta, M. et al (LVD Collaboration). Muon “depth-intensity” relation measured by the lvd underground experiment and cosmic-ray muon spectrum at sea level. *Phys. Rev. D*, 58:092005, Oct 1998. doi: [10.1103/PhysRevD.58.092005](https://doi.org/10.1103/PhysRevD.58.092005).
- [43] E. V. Bugaev, A. Misaki, V. A. Naumov, et al. Atmospheric muon flux at sea level, underground, and underwater. *Phys. Rev. D*, 58:054001, Jul 1998. doi: [10.1103/PhysRevD.58.054001](https://doi.org/10.1103/PhysRevD.58.054001).
- [44] E. Bugaev, V. Naumov, S. Sinegovsky, and E. Zaslavskaya. Prompt leptons in cosmic rays. *Il Nuovo Cimento C*, 12:41–73, 1989. ISSN 1124-1896. doi: [10.1007/BF02509070](https://doi.org/10.1007/BF02509070).
- [45] John F Beacom and Julián Candia. Shower power: isolating the prompt atmospheric neutrino flux using electron neutrinos. *Journal of Cosmology and Astroparticle Physics*, 2004(11):009, 2004. URL <http://stacks.iop.org/1475-7516/2004/i=11/a=009>.
- [46] A.D. Martin, M.G. Ryskin, and A.M. Stasto. Prompt neutrinos from atmospheric c-cbar and b-bbar production and the gluon at very small x. arXiv: [hep-ph/0302140v2](https://arxiv.org/abs/hep-ph/0302140v2).
- [47] M. Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki. Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data. *Phys. Rev. D*, 75(4):043006, Feb 2007. doi: [10.1103/PhysRevD.75.043006](https://doi.org/10.1103/PhysRevD.75.043006).
- [48] John G Learned and Sandip Pakvasa. Detecting oscillations at pev energies. *Astroparticle Physics*, 3(3):267 – 274, 1995. ISSN 0927-6505. doi: [10.1016/0927-6505\(94\)00043-3](https://doi.org/10.1016/0927-6505(94)00043-3).
- [49] Sandhya Choubey and Werner Rodejohann. Flavor composition of ultrahigh energy neutrinos at source and at neutrino telescopes. *Phys. Rev. D*, 80:113006, Dec 2009. doi: [10.1103/PhysRevD.80.113006](https://doi.org/10.1103/PhysRevD.80.113006).
- [50] Matthew G. Baring. Diffusive shock acceleration of high energy cosmic rays. *Nuclear Physics B - Proceedings Supplements*, 136(0):198 – 207, 2004. ISSN 0920-5632. doi: [10.1016/j.nuclphysbps.2004.10.008](https://doi.org/10.1016/j.nuclphysbps.2004.10.008). CRIS 2004 Proceedings of the Cosmic Ray International Seminars: GZK and Surroundings.
- [51] Tamar Kashti and Eli Waxman. Astrophysical neutrinos: Flavor ratios depend on energy. *Phys. Rev. Lett.*, 95:181101, Oct 2005. doi: [10.1103/PhysRevLett.95.181101](https://doi.org/10.1103/PhysRevLett.95.181101).

- [52] T. K. Gaisser, F. Halzen, and T. Stanev. Particle astrophysics with high energy neutrinos. *Physics Reports*, 258:173–236, July 1995. doi: [10.1016/0370-1573\(95\)00003-Y](https://doi.org/10.1016/0370-1573(95)00003-Y), arXiv: [hep-ph/9410384](https://arxiv.org/abs/hep-ph/9410384).
- [53] F. Halzen and D. Hooper. High-energy neutrino astronomy: the cosmic ray connection. *Reports on Progress in Physics*, 65:1025–1078, July 2002. doi: [10.1088/0034-4885/65/7/201](https://doi.org/10.1088/0034-4885/65/7/201), arXiv: [astro-ph/0204527](https://arxiv.org/abs/astro-ph/0204527).
- [54] Christian Spiering. Towards high-energy neutrino astronomy. *The European Physical Journal H*, 37(3):515–565, 2012. ISSN 2102-6459. doi: [10.1140/epjh/e2012-30014-2](https://doi.org/10.1140/epjh/e2012-30014-2).
- [55] S. Adrian-Martinez et al. Searches for clustering in the time integrated skymap of the ANTARES neutrino telescope. *JCAP*, 1405:001, 2014. doi: [10.1088/1475-7516/2014/05/001](https://doi.org/10.1088/1475-7516/2014/05/001), arXiv: [1402.2809](https://arxiv.org/abs/1402.2809).
- [56] R. Abbasi et al. (IceCube Collaboration). An absence of neutrinos associated with cosmic-ray acceleration in γ -ray bursts. *Nature*, 484:351–353, 2012. doi: [10.1038/nature11068](https://doi.org/10.1038/nature11068), arXiv: [1204.4219](https://arxiv.org/abs/1204.4219).
- [57] Julia Schmid. Search for High Energy GRB Neutrino Emission with ANTARES. pages 327–330, 2013, arXiv: [1311.4069](https://arxiv.org/abs/1311.4069).
- [58] M. Aartsen et al. (IceCube Collaboration). Search for neutrino-induced particle showers with icecube-40. *Phys. Rev. D*, 89:102001, May 2014. doi: [10.1103/PhysRevD.89.102001](https://doi.org/10.1103/PhysRevD.89.102001).
- [59] M. Aartsen et al. (IceCube Collaboration). First Observation of PeV-Energy Neutrinos with IceCube. *Phys. Rev. Lett.*, 111:021103, Jul 2013. doi: [10.1103/PhysRevLett.111.021103](https://doi.org/10.1103/PhysRevLett.111.021103).
- [60] Paolo Lipari. Establishing the astrophysical origin of a signal in a neutrino telescope. 2013, arXiv: [1308.2086](https://arxiv.org/abs/1308.2086).
- [61] Kai Zuber. *Neutrino physics*. Institute of Physics Publishing, Philadelphia, 2003.
- [62] Sheldon L. Glashow. Resonant scattering of antineutrinos. *Phys. Rev.*, 118:316–317, Apr 1960. doi: [10.1103/PhysRev.118.316](https://doi.org/10.1103/PhysRev.118.316).
- [63] Raj Gandhi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. Ultrahigh-energy neutrino interactions. *Astroparticle Physics*, 5(2):81 – 110, 1996. ISSN 0927-6505. doi: [10.1016/0927-6505\(96\)00008-4](https://doi.org/10.1016/0927-6505(96)00008-4).

- [64] Raj Gandhi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. Neutrino interactions at ultrahigh energies. *Phys. Rev. D*, 58(9):093009, Sep 1998. doi: [10.1103/PhysRevD.58.093009](https://doi.org/10.1103/PhysRevD.58.093009).
- [65] George Sterman et al. (CTEQ Collaboration). Handbook of perturbative qcd. *Rev. Mod. Phys.*, 67(1):157–248, Jan 1995. doi: [10.1103/RevModPhys.67.157](https://doi.org/10.1103/RevModPhys.67.157).
- [66] H.L. Lai, J. Huston, S. Kuhlmann, et al. Global qcd analysis of parton structure of the nucleon: Cteq5 parton distributions. *The European Physical Journal C - Particles and Fields*, 12(3):375–392, 2000. ISSN 1434-6044. doi: [10.1007/s100529900196](https://doi.org/10.1007/s100529900196).
- [67] Bernhard Voigt. *Sensitivity of the IceCube Detector for Ultra-High Energy Electron-Neutrino Events*. PhD thesis, Humboldt-Universität zu Berlin, April 2008. URL <http://edoc.hu-berlin.de/dissertationen/voigt-bernhard-2008-07-16/PDF/voigt.pdf>.
- [68] Christiaan Huygens. *Traité de la Lumiere*. Pierre Vander Aa, Leyden, 1690.
- [69] John David Jackson. *Classical Electrodynamics*. John Wiley & Sons, 3 edition, 1999.
- [70] Stephen G. Warren. Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt.*, 23(8):1206–1225, Apr 1984. doi: [10.1364/AO.23.001206](https://doi.org/10.1364/AO.23.001206).
- [71] R. M. Sternheimer and R. F. Peierls. General expression for the density effect for the ionization loss of charged particles. *Phys. Rev. B*, 3:3681–3692, Jun 1971. doi: [10.1103/PhysRevB.3.3681](https://doi.org/10.1103/PhysRevB.3.3681).
- [72] Dmitry Chirkin and Wolfgang Rhode. Muon Monte Carlo: A High-precision tool for muon propagation through matter. 2004, arXiv: [hep-ph/0407075](https://arxiv.org/abs/hep-ph/0407075).
- [73] J.-H. Koehne, K. Frantzen, M. Schmitz, et al. PROPOSAL: A tool for propagation of charged leptons. *Computer Physics Communications*, 184(9):2070 – 2090, 2013. ISSN 0010-4655. doi: [10.1016/j.cpc.2013.04.001](https://doi.org/10.1016/j.cpc.2013.04.001).
- [74] Donald E. Groom, Nikolai V. Mokhov, and Sergei I. Striganov. Muon Stopping Power and Range Tables 10 MeV–100 TeV. *Atomic Data and Nuclear Data Tables*, 78(2):183 – 356, 2001. ISSN 0092-640X. doi: [10.1006/adnd.2001.0861](https://doi.org/10.1006/adnd.2001.0861).

- [75] Christopher Wiebusch. *The Detection of Faint Light in Deep Underwater Neutrino Telescopes*. PhD thesis, RWTH Aachen, 1995. URL <http://web.physik.rwth-aachen.de/~wiebusch/Publications/VariouS/phd.pdf>.
- [76] Marek Kowalski. *Search for Neutrino-Induced Cascades with the AMANDA-II Detector*. PhD thesis, Humboldt-Universität zu Berlin, January 2004.
- [77] Marek Kowalski. On the cherenkov light emission of hadronic and electro-magnetic cascades. Internal Report 20020803, AMANDA Collaboration, August 2002. URL <http://internal.icecube.wisc.edu/reports/amanda/data/20020803-track.pdf>.
- [78] Leif Rädcl and Christopher Wiebusch. Calculation of the cherenkov light yield from electromagnetic cascades in ice with geant4. *Astroparticle Physics*, 44(0):102 – 113, 2013. ISSN 0927-6505. doi: [10.1016/j.astropartphys.2013.01.015](https://doi.org/10.1016/j.astropartphys.2013.01.015).
- [79] S. Agostinelli et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003. ISSN 0168-9002. doi: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [80] T. A. Gabriel, Donald E. Groom, P. K. Job, N. V. Mokhov, and G. R. Stevenson. Energy dependence of hadronic activity. *Nucl. Instrum. Meth.*, A338:336–347, 1994. doi: [10.1016/0168-9002\(94\)91317-X](https://doi.org/10.1016/0168-9002(94)91317-X).
- [81] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. Pythia 6.4 physics and manual. *Journal of High Energy Physics*, 2006(05):026, 2006. URL <http://stacks.iop.org/1126-6708/2006/i=05/a=026>.
- [82] R. Abbasi et al. (IceCube Collaboration). The IceCube data acquisition system: Signal capture, digitization, and timestamping. *Nuclear Instruments and Methods A*, 601(3):294 – 316, 2009. ISSN 0168-9002. doi: [10.1016/j.nima.2009.01.001](https://doi.org/10.1016/j.nima.2009.01.001).
- [83] R. Abbasi et al. (IceCube Collaboration). The design and performance of icecube deepcore. *Astroparticle Physics*, 35(10):615 – 624, 2012. ISSN 0927-6505. doi: [10.1016/j.astropartphys.2012.01.004](https://doi.org/10.1016/j.astropartphys.2012.01.004).
- [84] R. Abbasi et al. (IceCube Collaboration). Calibration and characterization of the icecube photomultiplier tube. *Nuclear Instruments and Methods A*, 618(1–3):139 – 152, 2010. ISSN 0168-9002. doi: [10.1016/j.nima.2010.03.102](https://doi.org/10.1016/j.nima.2010.03.102).

- [85] Peter Redl. *A Search for Muon Neutrinos Coincident with Gamma-Ray Bursts with the IceCube 59-String Detector*. PhD thesis, University of Maryland, College Park, 2011. URL http://drum.lib.umd.edu/bitstream/1903/11964/1/Redl_umd_0117E_12524.pdf.
- [86] Thomas Eberl. The seatray software framework. In *KM3Net Collaboration Meeting*, Marseille, 2013. URL <https://indico.in2p3.fr/getFile.py/access?contribId=49&sessionId=14&resId=0&materialId=slides&confId=7690>.
- [87] Nathan Whitehorn. *A Search for High-Energy Neutrino Emission from Gamma-Ray Bursts*. PhD thesis, University of Wisconsin-Madison, April 2012. URL <https://docushare.icecube.wisc.edu/dsweb/Get/Document-60879/thesis.pdf>.
- [88] J. Ahrens et al. (AMANDA Collaboration). Muon track reconstruction and data selection techniques in AMANDA. *Nuclear Instruments and Methods A*, 524(1–3):169 – 194, 2004. ISSN 0168-9002. doi: [10.1016/j.nima.2004.01.065](https://doi.org/10.1016/j.nima.2004.01.065).
- [89] P. B. Price, K. Woschnagg, and D. Chirkin. Age vs depth of glacial ice at south pole. *Geophysical Research Letters*, 27(14):2129–2132, 2000. ISSN 1944-8007. doi: [10.1029/2000GL011351](https://doi.org/10.1029/2000GL011351).
- [90] M. Ackermann et al (AMANDA Collaboration). Optical properties of deep glacial ice at the South Pole. *J. Geophys. Res.*, 111(D13):1–26, Jan 2006. doi: [10.1029/2005JD006687](https://doi.org/10.1029/2005JD006687).
- [91] M. Aartsen et al. (IceCube Collaboration). Measurement of South Pole ice transparency with the IceCube LED calibration system. *Nucl. Instrum. Meth. A*, 711(0):73 – 89, 2013. ISSN 0168-9002. doi: [10.1016/j.nima.2013.01.054](https://doi.org/10.1016/j.nima.2013.01.054).
- [92] Dmitry Chirkin. Photon tracking with GPUs in IceCube. *Nucl. Instrum. Meth. A*, 725(0):141 – 143, 2013. ISSN 0168-9002. doi: [10.1016/j.nima.2012.11.170](https://doi.org/10.1016/j.nima.2012.11.170). VLVT 11, Erlangen, Germany, 12 - 14 October, 2011 5th International Workshop on Very Large Volume Neutrino Telescopes, The future of high-energy neutrino astronomy.
- [93] Johan Lundberg, Predrag Miočinović, Kurt Woschnagg, et al. Light tracking through ice and water-scattering and absorption in heterogeneous media with photonics. *Nucl. Inst. Meth. A*, 581(3):619 – 631, 2007. ISSN 0168-9002. doi: [10.1016/j.nima.2007.07.143](https://doi.org/10.1016/j.nima.2007.07.143), arXiv: [arXiv:astro-ph/0702108v2](https://arxiv.org/abs/astro-ph/0702108v2).

- [94] Dirk Pandel. Bestimmung von Wasser- und Detektorparametern und Rekonstruktion von Myonen bis 100 TeV mit dem Baikal-Neutrino-Teleskop NT-72. Diploma thesis, Humboldt-Universität zu Berlin, 1996.
- [95] George Japaridze and Mathieu Ribordy. Realistic arrival time distribution from an isotropic light source. 2005. URL <http://arxiv.org/abs/astro-ph/0506136>, arXiv: astro-ph/0506136.
- [96] N. van Eijndhoven, O. Fadiran, and G. Japaridze. Implementation of a gauss convoluted pandel {PDF} for track reconstruction in neutrino telescopes. *Astroparticle Physics*, 28(4–5):456 – 462, 2007. ISSN 0927-6505. doi: 10.1016/j.astropartphys.2007.09.001.
- [97] Jakob van Santen. Markov-Chain Monte Carlo Reconstruction for Cascade-like Events in IceCube. Master’s thesis, Humboldt-Universität zu Berlin, January 2010. URL http://www.icecube.wisc.edu/~jvansanten/lit/jvs_diplom.pdf.
- [98] Eike Middell. Reconstruction of cascade-like events in icecube. Diploma thesis, Humboldt-Universität zu Berlin, July 2008.
- [99] Douglas Rutledge. *A Search for Neutrino-Induced Electromagnetic Showers in the 2008 Combined IceCube and AMANDA Detectors*. PhD thesis, Pennsylvania State University, July 2011.
- [100] Dmitry Chirkin and Christopher Weaver. TopologicalSplitter. August 2012. URL <http://code.icecube.wisc.edu/svn/projects/TopologicalSplitter>.
- [101] K. M. Górski, E. Hivon, A. J. Banday, et al. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005. URL <http://stacks.iop.org/0004-637X/622/i=2/a=759>.
- [102] Benedikt Riedel. *Modeling and Understanding Supernova Signals in the IceCube Neutrino Observatory*. PhD thesis, University of Wisconsin-Madison, October 2014.
- [103] Aya Ishihara and Kotoyo Hoshina. NeutrinoGenerator. April 2005. URL <http://code.icecube.wisc.edu/svn/projects/NeutrinoGenerator>.
- [104] A. Gazizov and M. Kowalski. ANIS: High energy neutrino generator for neutrino telescopes. *Computer Physics Communications*, 172(3):203 – 213, 2005. ISSN 0010-4655. doi: 10.1016/j.cpc.2005.03.113.

- [105] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa. New calculation of the atmospheric neutrino flux in a three-dimensional scheme. *Phys. Rev. D*, 70:043008, Aug 2004. doi: [10.1103/PhysRevD.70.043008](https://doi.org/10.1103/PhysRevD.70.043008).
- [106] L. V. Volkova. Energy Spectra and Angular Distributions of Atmospheric Neutrinos. *Soviet Journal of Nuclear Physics*, 31:784, 1980.
- [107] T. K. Gaisser, S. A. Bludman, H. Lee, and T. Stanev. Flux of atmospheric neutrinos. *Phys. Rev. Lett.*, 51:223–226, July 1983. doi: [10.1103/PhysRevLett.51.223](https://doi.org/10.1103/PhysRevLett.51.223).
- [108] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa. Calculation of the flux of atmospheric neutrinos. *Phys. Rev. D*, 52:4985–5005, November 1995. doi: [10.1103/PhysRevD.52.4985](https://doi.org/10.1103/PhysRevD.52.4985), arXiv: [hep-ph/9503439](https://arxiv.org/abs/hep-ph/9503439).
- [109] G. D. Barr, T. K. Gaisser, P. Lipari, Simon Robbins, and T. Stanev. A three-dimensional calculation of atmospheric neutrinos. *Phys. Rev.*, D70:023006, 2004. doi: [10.1103/PhysRevD.70.023006](https://doi.org/10.1103/PhysRevD.70.023006), arXiv: [astro-ph/0403630](https://arxiv.org/abs/astro-ph/0403630).
- [110] C. Gonzalez-Garcia, M. Maltoni, and J. Rojo. Determination of the atmospheric neutrino fluxes from atmospheric neutrino data. *Journal of High Energy Physics*, 10:075, October 2006. doi: [10.1088/1126-6708/2006/10/075](https://doi.org/10.1088/1126-6708/2006/10/075), arXiv: [hep-ph/0607324](https://arxiv.org/abs/hep-ph/0607324).
- [111] P. Lipari, M. Lusignoli, and D. Meloni. Flavor composition and energy spectrum of astrophysical neutrinos. *Phys. Rev. D*, 75(12):123005, June 2007. doi: [10.1103/PhysRevD.75.123005](https://doi.org/10.1103/PhysRevD.75.123005), arXiv: [0704.0718](https://arxiv.org/abs/0704.0718).
- [112] A. Fedynitch, J. Becker Tjus, and P. Desiati. Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux. *Phys. Rev. D*, 86(11):114024, December 2012. doi: [10.1103/PhysRevD.86.114024](https://doi.org/10.1103/PhysRevD.86.114024), arXiv: [1206.6710](https://arxiv.org/abs/1206.6710).
- [113] S. Haino et al (BESS Collaboration). Measurements of primary and atmospheric cosmic-ray spectra with the bess-tev spectrometer. *Physics Letters B*, 594(1–2):35 – 46, 2004. ISSN 0370-2693. doi: [10.1016/j.physletb.2004.05.019](https://doi.org/10.1016/j.physletb.2004.05.019).
- [114] T. Sanuki et al (BESS Collaboration). Measurements of atmospheric muon spectra at mountain altitude. *Physics Letters B*, 541(3–4):234 – 242, 2002. ISSN 0370-2693. doi: [10.1016/S0370-2693\(02\)02265-7](https://doi.org/10.1016/S0370-2693(02)02265-7).

- [115] P. Achard et al (L3 Collaboration). Measurement of the atmospheric muon spectrum from 20 to 3000 gev. *Physics Letters B*, 598(1-2):15 – 32, 2004. ISSN 0370-2693. doi: [10.1016/j.physletb.2004.08.003](https://doi.org/10.1016/j.physletb.2004.08.003).
- [116] R. Abbasi et al. (IceCube Collaboration). Search for a diffuse flux of astrophysical muon neutrinos with the icecube 40-string detector. *Phys. Rev. D*, 84:082001, Oct 2011. doi: [10.1103/PhysRevD.84.082001](https://doi.org/10.1103/PhysRevD.84.082001).
- [117] Dmitry Chirkin. Fluxes of atmospheric leptons at 600-GeV - 60-TeV. 2004, arXiv: [hep-ph/0407078](https://arxiv.org/abs/hep-ph/0407078).
- [118] Christine Lewis and Teresa Montaruli. NeutrinoFlux. November 2007. URL <http://code.icecube.wisc.edu/svn/projects/neutrinoflux>.
- [119] Christopher Weaver. NewNuFlux. May 2013. URL <http://code.icecube.wisc.edu/svn/sandbox/cweaver/NewNuFlux>.
- [120] L. V. Volkova. On uncertainties in prompt atmospheric neutrino flux calculations. *Physics Letters B*, 462:211–216, September 1999. doi: [10.1016/S0370-2693\(99\)00786-8](https://doi.org/10.1016/S0370-2693(99)00786-8).
- [121] V. S. Berezinsky, D. Cline, and D. N. Schramm. Prompt atmospheric neutrino production of W-bosons. *Physics Letters B*, 78:635–637, October 1978. doi: [10.1016/0370-2693\(78\)90658-5](https://doi.org/10.1016/0370-2693(78)90658-5).
- [122] L. V. Volkova. Prompt Lepton Production - Atmospheric Muon and Neutrino Spectra at High Energies. *International Cosmic Ray Conference*, 7:22, August 1983.
- [123] G. Gelmini, P. Gondolo, and G. Varieschi. Prompt atmospheric neutrinos and muons: Dependence on the gluon distribution function. *Phys. Rev. D*, 61(5):056011, March 2000. doi: [10.1103/PhysRevD.61.056011](https://doi.org/10.1103/PhysRevD.61.056011), arXiv: [hep-ph/9905377](https://arxiv.org/abs/hep-ph/9905377).
- [124] A. D. Martin, M. G. Ryskin, and A. M. Stasto. Prompt Neutrinos from Atmospheric $c\bar{c}$ Production and the Gluon at Very Small x . *Acta Physica Polonica B*, 34:3273, June 2003, arXiv: [hep-ph/0302140](https://arxiv.org/abs/hep-ph/0302140).
- [125] Thomas K. Gaisser. Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio. *Astroparticle Physics*, 35(12):801 – 806, 2012. ISSN 0927-6505. doi: [10.1016/j.astropartphys.2012.02.010](https://doi.org/10.1016/j.astropartphys.2012.02.010).
- [126] J. W. Elbert. Multiple Muons Produced by Cosmic Ray Interactions. In *Proceedings of the DUMAND Summer Workshop (Scripps Institution of Oceanography, La Jolla, CA)*, 1979.

- [127] M. Thunman, G. Ingelman, and P. Gondolo. Charm production and high energy atmospheric muon and neutrino fluxes. *Astroparticle Physics*, 5(3–4):309 – 332, 1996. ISSN 0927-6505. doi: [10.1016/0927-6505\(96\)00033-3](https://doi.org/10.1016/0927-6505(96)00033-3).
- [128] Anatoli Fedynitch. Calculation of conventional and prompt lepton fluxes at very high energy. Presentation at ISVHECRI 2014, August 2014. URL <https://indico.cern.ch/event/287474/session/10/contribution/50/material/slides/0.pdf>.
- [129] Stefan Schönert, Thomas K. Gaisser, Elisa Resconi, and Olaf Schulz. Vetoing atmospheric neutrinos in a high energy neutrino telescope. *Phys. Rev. D*, 79:043009, Feb 2009. doi: [10.1103/PhysRevD.79.043009](https://doi.org/10.1103/PhysRevD.79.043009).
- [130] Eun-Joo Ahn, Ralph Engel, Thomas K. Gaisser, Paolo Lipari, and Todor Stanev. Cosmic ray interaction event generator SIBYLL 2.1. *Phys. Rev. D*, 80:094003, Nov 2009. doi: [10.1103/PhysRevD.80.094003](https://doi.org/10.1103/PhysRevD.80.094003).
- [131] P Berghaus, T Montaruli, and J Ranft. Charm production in dpmjet. *Journal of Cosmology and Astroparticle Physics*, 2008(06):003, 2008. doi: [10.1088/1475-7516/2008/06/003](https://doi.org/10.1088/1475-7516/2008/06/003).
- [132] Fred James. Monte carlo phase space. Lectures given at the academic training programme at CERN 15, CERN, 1968. URL <http://cds.cern.ch/record/275743/files/CERN-68-15.pdf>.
- [133] Abraham Loeb and Eli Waxman. The cumulative background of high energy neutrinos from starburst galaxies. *Journal of Cosmology and Astroparticle Physics*, 2006(05):003, 2006. URL <http://stacks.iop.org/1475-7516/2006/i=05/a=003>.
- [134] Y. Becherini, A. Margiotta, M. Sioli, and M. Spurio. A parameterisation of single and multiple muons in the deep water or ice. *Astroparticle Physics*, 25(1):1 – 13, 2006. ISSN 0927-6505. doi: [10.1016/j.astropartphys.2005.10.005](https://doi.org/10.1016/j.astropartphys.2005.10.005).
- [135] Felix Riehn. Charm Production in SIBYLL. Presentation at ISVHECRI 2014, Geneva, Switzerland, August 2014. URL <https://indico.cern.ch/event/287474/session/15/contribution/48/material/slides/0.pdf>.

- [136] Anatoli Fedynitch, Julia Becker Tjus, and Paolo Desiati. Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux. *Phys. Rev. D*, 86: 114024, Dec 2012. doi: [10.1103/PhysRevD.86.114024](https://doi.org/10.1103/PhysRevD.86.114024).
- [137] S. A. Murphy and A. W. van der Vaart. On profile likelihood. *Journal of the American Statistical Association*, 95(450):pp. 449–465, 2000. ISSN 01621459. URL <http://www.jstor.org/stable/2669386>.
- [138] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 03 1938. doi: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360).
- [139] K. Murase, M. Ahlers, and B. C. Lacki. Testing the hadronuclear origin of PeV neutrinos observed with IceCube. *Phys. Rev. D*, 88(12):121301, December 2013. doi: [10.1103/PhysRevD.88.121301](https://doi.org/10.1103/PhysRevD.88.121301), arXiv: [1306.3417](https://arxiv.org/abs/1306.3417).
- [140] M. Aartsen et al. Probing the origin of cosmic rays with extremely high energy neutrinos using the icecube observatory. *Phys. Rev. D*, 88:112008, Dec 2013. doi: [10.1103/PhysRevD.88.112008](https://doi.org/10.1103/PhysRevD.88.112008).
- [141] F. W. Stecker. Note on high-energy neutrinos from active galactic nuclei cores. *Phys. Rev. D*, 72: 107301, Nov 2005. doi: [10.1103/PhysRevD.72.107301](https://doi.org/10.1103/PhysRevD.72.107301).
- [142] Patrick Hallen. On the Measurement of High-Energy Tau Neutrinos with IceCube. Master’s thesis, RWTH Aachen, November 2013. URL <https://internal.icecube.wisc.edu/reports/details.php?type=report&id=icecube%2F201312001>.
- [143] Berghaus, Patrick. Atmospheric muons as icecube signal. *EPJ Web of Conferences*, 52:09006, 2013. doi: [10.1051/epjconf/20125209006](https://doi.org/10.1051/epjconf/20125209006).
- [144] Joerg R. Hoerandel. On the knee in the energy spectrum of cosmic rays. *Astroparticle Physics*, 19(2): 193 – 220, 2003. ISSN 0927-6505. doi: [10.1016/S0927-6505\(02\)00198-6](https://doi.org/10.1016/S0927-6505(02)00198-6).
- [145] Louis Lyons. Open statistical issues in particle physics. *The Annals of Applied Statistics*, 2(3):887–915, 09 2008. doi: [10.1214/08-AOAS163](https://doi.org/10.1214/08-AOAS163).
- [146] Steve Baker and Robert D. Cousins. Clarification of the use of chi-square and likelihood functions in fits to histograms. *Nuclear Instruments and Methods in Physics Research*, 221(2):437 – 442, 1984. ISSN 0167-5087. doi: [10.1016/0167-5087\(84\)90016-4](https://doi.org/10.1016/0167-5087(84)90016-4).

- [147] W. Feller. On the kolmogorov-smirnov limit theorems for empirical distributions. *The Annals of Mathematical Statistics*, 19(2):177–189, 06 1948. doi: [10.1214/aoms/1177730243](https://doi.org/10.1214/aoms/1177730243).
- [148] Ronald A. Fisher. Dispersion on a sphere. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 217(1130):295–305, May 1953. URL <http://www.jstor.org/stable/99186>.
- [149] E. Andres et al (AMANDA Collaboration). Observation of high-energy neutrinos using Cherenkov detectors embedded deep in Antarctic ice. *Nature*, 410:441–443, 2001. doi: [10.1038/35068509](https://doi.org/10.1038/35068509).
- [150] M. Ageron et al. (ANTARES Collaboration). Antares: The first undersea neutrino telescope. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 656(1):11 – 38, 2011. ISSN 0168-9002. doi: [10.1016/j.nima.2011.06.103](https://doi.org/10.1016/j.nima.2011.06.103).
- [151] Thorsten Stetzlberger. DOMApp firmware timing. IceCube Internal Report, 2007. URL <https://docushare.icecube.wisc.edu/dsweb/Get/Document-28424/DOMAPptiming.pdf>.
- [152] D. Heck, J. Knapp, J.N. Capdevielle, G. Schatz, and T. Thouw. CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers. Technical Report FZKA 6019, Forschungszentrum Karlsruhe, 1998.
- [153] Dmitry Chirkin. *Cosmic Ray Energy Spectrum Measurement with the Antarctic Muon and Neutrino Detector Array (AMANDA)*. PhD thesis, University of California at Berkeley, 2003. URL <http://dima.lbl.gov/work/BKP/DCS/THESIS/main.pdf>.