

Evidence for Astrophysical Muon Neutrinos from the Northern Sky

By

Christopher Weaver

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2015

Date of final oral examination: February 2, 2015

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

Albrecht Karle, Professor, Physics

Francis Halzen, Professor, Physics

John Gallagher, Professor, Astronomy

Jennifer Thomas, Professor, Physics (University College London)

Justin Vandenbroucke, Assistant Professor, Physics

TABLE OF CONTENTS

	Page
LIST OF TABLES	iii
LIST OF FIGURES	iv
1 Cosmic Rays	1
1.1 Observation	1
1.2 Acceleration	4
1.3 Possible Sources	8
1.4 Showering Processes	12
2 Neutrinos	14
2.1 Neutrino Physics	14
2.1.1 History	14
2.1.2 Current Understanding	16
2.2 Astrophysical Production of Neutrinos	20
2.2.1 Bound on the Total Diffuse Neutrino Flux	21
2.2.2 Neutrinos from Active Galactic Nuclei	21
2.2.3 Neutrinos from Starburst Galaxies	22
2.2.4 Neutrinos from Gamma Ray Bursts	24
2.2.5 Cosmogenic Neutrinos	25
2.3 Atmospheric Production of Neutrinos	26
2.3.1 Conventional Atmospheric Neutrinos	26
2.3.2 Prompt Atmospheric Neutrinos	30
3 Detector	32
3.1 Detection Principle	32
3.2 The IceCube Detector	37
3.2.1 Ice	39
3.3 Data Acquisition and Processing	42
3.4 Event Reconstruction	45
3.4.1 Waveform Feature Extraction	45
3.4.2 Directional Reconstruction of Track-like Events	47
3.4.3 Event Splitting	49
3.4.4 Muon Energy Reconstruction	51
3.5 Production of Simulated Data	53
3.5.1 Neutrino Interaction Generation	53
3.5.2 Air-shower Simulation with CORSIKA	54
3.5.3 Muon Propagation and Cascade Generation	55
3.5.4 Photon Propagation	55

	Page
3.5.5 Hardware Simulation	56
3.5.6 Triggering Simulation	59
4 Data Selection	61
4.1 Online Filter	62
4.1.1 2010 Muon Filter	62
4.1.2 2011 Muon Filter	64
4.2 Data Volume Reduction	64
4.3 High Level Reconstruction	65
4.4 Final Selection Criteria	69
4.4.1 Zenith Angle Restriction	70
4.4.2 Track Reconstruction Quality Selection	70
4.4.3 Bayesian Background Rejection	74
4.4.4 Direction Dependent Brightness Cut	76
4.4.5 Overlapping Muon Rejection	78
4.5 Muon and Neutrino Effective Areas	82
4.6 Final Data Properties	84
5 Analysis Method	87
5.1 Likelihood Fit	87
5.1.1 Error Estimation	89
5.1.2 Note on Biased Observables	90
5.2 Model Construction	91
5.2.1 Note on Weighting	91
5.2.2 Atmospheric Parameters	92
5.2.3 Astrophysical Parameters	93
5.2.4 Detector Parameters	93
5.3 Analysis Sensitivity	96
6 Results	98
6.1 Observed Events	98
6.2 E^{-2} Fit	102
6.2.1 Separate Fits of Data-taking Periods	104
6.2.2 Observation Domain	106
6.2.3 Observation Significance	108
6.3 Free Power law Fit	109
6.4 Model Fits	113
6.5 Signal Probability of Individual Events	114
7 Conclusions	116
7.1 Comparison to Other Recent Results	116
7.2 Outlook for Distinguishing Flux Models	119
7.3 Implications for Source Identification	120
APPENDIX Details of Observed Events	125

LIST OF REFERENCES 135

LIST OF TABLES

Table	Page
4.1 Fraction of coincident track events which were correctly separated by the TopologicalTrigger algorithm for various choices of its parameters.	68
4.2 Rates and fractions of data surviving by type as a function of the level of selection applied . . .	80
6.1 Basic properties of the 20 observed events with the highest energy proxy values	99
6.2 Best fit parameters when an E^{-2} astrophysical flux is included	103
6.3 Best fit parameters when an astrophysical flux with the form of an arbitrary power law is included	110
6.4 Best fit normalizations resulting from fitting an astrophysical flux given by various representative theoretical model predictions	113
A.1 Images of observed events with highest estimated energies	125
A.2 Inferred event energies	129

LIST OF FIGURES

Figure	Page
1.1	The combined energy spectrum of cosmic ray particles 3
1.2	The high-energy cosmic ray spectrum observed by the Pierre Auger Observatory 9
1.3	Comparison of object sizes and magnetic field strengths to cosmic rays production requirements 11
1.4	Analytic calculation of the electromagnetic component of an air-shower 12
1.5	Schematic of important portions of a cosmic ray induced particle shower 13
2.1	The energy spectrum of electrons emitted from a mixture of ^{214}Pb and ^{214}Bi 15
2.2	The time distribution of events observed by the Kamiokande II detector, in coincidence with supernova 1987A 17
2.3	Common interactions between neutrinos and hadronic matter 18
2.4	High energy neutrino cross sections for interactions with isoscalar material 19
2.5	Attenuation of neutrinos passing through the Earth as a function of zenith angle 19
2.6	The Waxman Bahcall diffuse neutrino flux bound 22
2.7	The Stecker AGN core neutrino flux model 23
2.8	The Loeb-Waxman starburst galaxy neutrino flux model 24
2.9	The Waxman-Bahcall GRB neutrino flux model 25
2.10	The conventional atmospheric neutrino flux predicted by the HKKM calculation 28
2.11	The pionic and kaonic components of the conventional atmospheric neutrino flux 29
2.12	The correction to the HKKM flux calculation to include the effects of the CR knee 30
2.13	Comparison of the prompt atmospheric neutrino flux to the conventional flux 31
3.1	Schematic of the Čerenkov detection process 35
3.2	Cross sections for muon energy loss mechanisms in ice as a function of energy 36

Figure	Page
3.3 Cut-away view of the IceCube detector	38
3.4 Schematic of the IceCube Digital Optical Module	39
3.5 Two measurements of the optical properties of the South Pole ice	41
3.6 Schematic of the detector module readout electronics	43
3.7 Example of waveform unfolding	46
3.8 The expected distribution of light amplitude as a function of distance from a muon track	52
3.9 The Preliminary Reference Earth Model	54
3.10 The probability density function from which amplified photomultiplier pulse charges are sampled	57
3.11 The simulated time distributions of different photomultiplier output pulse types	58
3.12 The simulated treatment of photomultiplier saturation for the IceCube photomultipliers	59
4.1 The cut used by the Muon Filter for down-going events during the 2010 data taking period	63
4.2 The overburden above the IceCube detector as a function of zenith angle	70
4.3 The effects of the zenith angle restriction cut	71
4.4 The ad-hoc correction to the paraboloid sigma	73
4.5 The effects of the reconstruction quality cut	75
4.6 The prior used on the zenith angle of the reconstructed track for the Bayesian reconstruction	76
4.7 The effects of the bayesian background rejection cut	77
4.8 The zenith angle distribution of observed data after the quality and Bayesian background rejection cuts	78
4.9 The effects of the direction dependent brightness cut	79
4.10 The effects of the overlapping muon rejection cuts	81
4.11 Neutrino effective area for this data sample	83
4.12 The muon effective area for this data sample	84
4.13 Median angular resolution for neutrino events in the final data sample for two representative spectra	85
4.14 Energy resolution for this data sample	86
5.1 Partial representation of one of the six DOM efficiency parameterizations needed in this analysis	95

Figure	Page
5.2 The expected sensitivity of this analysis estimated by Monte Carlo trials	97
6.1 Visualization from side and top perspectives of the observed event with the largest energy proxy value	99
6.2 Results of the detailed reconstruction study for the observed event with the highest energy proxy	101
6.3 Distributions of reconstructed event energies and zenith angles compared to the best fit model for an E^{-2} power law astrophysical flux	103
6.4 Behavior of all other fit parameters as a function of the normalization of an E^{-2} astrophysical flux	105
6.5 Comparison of the fit likelihoods as functions of E^{-2} astrophysical flux normalization for each year of data	106
6.6 Best fits for an E^{-2} astrophysical flux with hard cutoffs below 500 TeV and above 1.2 PeV . . .	108
6.7 Distribution of test statistic values obtained from realizations of the atmospheric-only hypothesis	109
6.8 Distribution of reconstructed event energies compared to the best fit model for an $E^{-\gamma}$ power law astrophysical flux	110
6.9 Comparisons of profile likelihood and prompt atmospheric normalization as functions of possible astrophysical power law fluxes	112
6.10 The fraction of the total flux arising from each fitted component	114
6.11 The predicted rate of events to be observed in this analysis as a function of neutrino energy . . .	115
7.1 Comparison of allowed astrophysical spectrum properties to other recent results	117
7.2 Comparison of the result of this work to other recent results and various flux models	118
7.3 Projected significance for distinguishing a flux which follows the $E^{-2.3}$ case of the starburst model from an unbroken E^{-2} power law	120
7.4 The sensitivity of searches for point sources, compared to the level of the diffuse neutrino flux .	121
7.5 Minimum number of point-like neutrino sources which could produce the observed diffuse flux while evading detection by dedicated point source searches	122
7.6 The distribution of the highest energy events observed by this analysis on the sky in equatorial coordinates	123
Appendix	
Figure	

Chapter 1

Cosmic Rays

1.1 Observation

Cosmic rays are generally defined as charged particles which reach the Earth from outer space. They have been a topic of study since they were first identified as a source of radiation at ground level on Earth distinct from the radioactive decay of unstable isotopes in the the Earth's crust. The first definite evidence that this radiation in fact came from outside the atmosphere was provided by Viktor Hess, who carried electroscopes (the typical means of measuring ambient radiation at the time) to altitudes of up to 5.2 km on balloon flights in 1912 [1]. Further experiments of the same type conducted by other researchers solidified the conclusion that there was some source of radiation reaching the atmosphere from above. Further investigation showed that the radiation typically observed at ground level did not consist of the original (primary) particles, but of secondary particles produced during passage through the atmosphere (it was assumed that the primary particles were gamma rays since no other radiation then known came close to having sufficient penetrating power, while secondary electrons produced by Compton scattering would provide the substantial amounts of observed ionization), and the study of these secondaries lead to both the theory of particle shower development (discussed in Section 1.4) and the experimental discovery of a number of particles including the positron and the muon.

After a century of observation there are both a number of salient properties of cosmic rays which have emerged, and some major open questions. First, the overall spectrum of the cosmic rays, shown in Fig. 1.1 is close to being a power law over a vast range in energies. A few features are visible as changes in the power law index at approximately 10^{15} eV (known as the 'knee') and at approximately 10^{18} eV (known as

the ‘ankle’). The spectrum is otherwise quite smooth until an apparent cutoff at around 10^{20} which may be evidence of the suppression predicted by Greisen, Zatsepin, and Kuz’min [2] due to the interaction of cosmic ray particles with the cosmic microwave background. At low to medium energies (up to about 100 TeV), where detectors small enough to be lifted above the bulk of the atmosphere (via balloons or rockets, for example [3]) can capture significant numbers of primary particles directly, the composition of the cosmic rays is dominated by protons up to about 10 TeV, after which Helium nuclei become dominant. Heavier nuclei are observed as well, in smaller quantities. At high energies, the low flux becomes a substantial challenge; the required collecting areas are so large that indirect detection techniques must be used in order to obtain a useful number of events, but this comes at the cost of detailed information about the individual events. In particular, it becomes generally impossible to directly observe the primary particle itself, so its identity must be inferred from the particle shower produced when it interacts, leading to considerable uncertainty about the composition of the high energy cosmic rays. The cosmic rays are generally found to be quite isotropic, which is consistent with the interpretation that they diffuse randomly through complex magnetic fields in space. Some anisotropy is observed, but few concrete conclusions can be drawn from it thus far. As a result, except for a single instance of direct evidence for emission from energetic hadrons in two supernova remnants [4], the production sites of the cosmic rays still remain unknown. One possible approach to determining these production sites observationally is by searching for neutrinos produced in conjunction with the cosmic rays. Neutrinos are a good tracer of cosmic ray acceleration since unlike gamma rays they cannot be produced electromagnetically and so are not easily produced by systems in which only leptons are present; their production is a strong indicator of hadronic particle interactions. Furthermore, since neutrinos have no electromagnetic charge and very small interaction cross-sections, they can reach Earth and still be tracked back to their points of origin.




Figure 1.1: A composite of the results from several cosmic ray experiments, illustrating the combined energy spectrum of all particle types over most of its range [5]. The ‘knee’ and ‘ankle’ are changes in the observed spectral index of the flux, consisting of the spectrum becoming steeper (‘softer’) and less steep (‘harder’), respectively.

1.2 Acceleration

As the central questions regarding cosmic rays are the locations at which they are produced and the mechanisms of their production, it is worthwhile to review the existing theoretical proposals we seek to test experimentally. Besides extending to extremely high energies, the spectrum of cosmic rays is notably in the form of a power law (with some features), and so attempts to explain the production mechanism typically hinge on deriving such a power law.

We will begin by examining a set of basic assumptions which, in general, give rise to such spectra, as follows: First assume that the particle undergoes some acceleration interactions at some average rate $1/\tau_a$, which on average increase its energy by a factor:

$$E' = \alpha E$$

Clearly, the energy of an individual particle with initial energy E_0 will on average grow exponentially with time

$$E(t) = E_0 \alpha^{\frac{t}{\tau_a}} = E_0 e^{(\frac{t}{\tau_a}) \ln \alpha} \quad (1.1)$$

Next, assume that each particle has a constant probability per unit time to escape from the acceleration region, carrying with it whatever final energy it had achieved. This means that the probability of a particle remaining for a time t while undergoing acceleration is distributed exponentially according to the average time to escape, τ_e :

$$P(t) = \frac{1}{\tau_e} e^{-\frac{t}{\tau_e}}$$

The distribution of energies can then be derived by solving Eq. 1.1 for the expected time required for a particle to reach a given energy E and substituting this into the time distribution:

$$t = \tau_a \frac{\ln \frac{E}{E_0} - \ln E_0}{\ln \alpha} \quad (1.2)$$

$$P(E) = \frac{1}{\tau_e} e^{-\frac{\tau_a}{\tau_e \ln \alpha} \ln \left(\frac{E}{E_0} \right)} = \frac{1}{\tau_e} \left(\frac{E}{E_0} \right)^{-\frac{\tau_a}{\tau_e \ln \alpha}} \quad (1.3)$$

As desired, this is a power law whose index $\frac{\tau_a}{\tau_e \ln \alpha}$. This has several logically-consistent properties: Increasing τ_a , which is to say making acceleration steps less frequent, makes the spectrum softer due to particles gaining less energy in the same period of time. Increasing τ_e hardens the spectrum, and this is expected since it is the same as requiring the particles to be confined and so accelerated for a longer time on average. Finally, increasing α , or causing each particle to gain more energy on average in each acceleration step also hardens the spectrum, since particles gain more energy over the same period of time.

This simple mechanism for generating a power law spectrum is termed a statistical bottom-up approach to producing high energy cosmic rays, as it produces them from a larger population of lower energy particles over a large number of randomized iterations. Other production schemes exist: direct acceleration in which particles' energies are drastically raised in a single step, such as by a strong electric field; or top-down, in which some (as yet unknown) particles which contain large energy by virtue of having large masses decay to produce the cosmic rays. These types of mechanisms are not necessarily ruled out by current observational data, but this work will mostly concern itself with possible shock acceleration locations and scenarios.

Having a general prescription for the conditions on a system which can produce power law spectra of accelerated particles, we must then identify physical systems in which those conditions are realized. The main proposed mechanism, upon which essentially all subsequent models have been based, put forward by Fermi in 1949 [6], is collisionless interaction of charged particles with ambient, variable magnetic fields. It should be noted that Fermi's original idea of acceleration of particles by random motions of magnetized gas clouds in the Galaxy (along with other contemporary models such as those favored by Teller and Alfvén in which cosmic rays were produced by the Sun and confined to the neighborhood of the solar system [7]) has been largely replaced by acceleration at shock fronts in moving plasma, as the latter is expected to accelerate with higher efficiency. A full treatment of Fermi acceleration at shocks (such as [8]) must deal with a number of details which will be omitted here for brevity such as the acceleration of particles to high enough initial energy to reach the regime considered here, the magnetohydrodynamics which implement the scattering, and specifics of the power law indices which may be produced. This discussion, based substantially on the derivation given in Chapter 11 of [9], will deal only with verifying the conditions such that a power law spectrum may be produced at all.

The important features of a shock are that it involves flow of gas at two different speeds on opposite sides of a planar interface, and that turbulent magnetic fields can exist within both regions. For the purposes

of this discussion, let the difference in the bulk motion of gas in front of and behind the shock be a non-relativistic velocity $\beta \ll 1$, and consider only particles which are substantially relativistic (in a reference frame motionless with respect to either region of the shock) so their energies may be reasonably approximated as $E = pc$.

The key insight suggested by Fermi is that, in the presence of a magnetic field, a charged particle may have its direction changed while neither gaining nor losing energy in the rest frame of that magnetic field, unlike in the case of colliding directly with other particles in the region, which would tend to have lower energies, thus sapping what energy the particle under consideration has thus far gained. Conservation of kinetic energy may still appear unhelpful for producing overall acceleration, but it provides the possibility for a particle to shuttle back and forth among multiple local magnetic fields which may be in motion with respect to each other. Although the particle gains no energy in the frame of one field while interacting with that field, it may gain (or lose) energy with respect to another field, and if each interaction gives an energy gain on average with respect to the other regions (or some larger reference frame in which all of the moving fields reside), overall acceleration will be achieved. A shock gives a convenient environment for this type of behavior, since there are two regions of magnetic fields in close proximity, making it easy for a particle to move back and forth between them.

Specifically, let us assume that each time our particle of interest enters one of the two regions of the shock, its direction will be isotropized in the rest frame of the gas in that region. This simple picture leads to a constant probability of escape: With a distribution of directions which is isotropic, the particle is equally likely to leave the region through the side on which the shock is present (and thus re-enter the other region), or to leave through the side away from the shock, in which case it ceases to be accelerated and presumably escapes into the surrounding space. Next, we must show that on average the particle gains energy proportionally to its instantaneous energy. This can be done by considering the boosts undergone by the particle when it transitions between sides of the shock. The timelike component of the boost between the two frames is

$$E' = \gamma(E \pm \beta p_x)$$

where the sign depends on the direction of the boost and p_x is the component of the particle's momentum perpendicular to the shock (and γ and β both refer to the relative velocity of the shock, rather than the particle's momentum p). Since we have assumed that the particle is relativistic we may rewrite the momentum

term as the energy to obtain

$$E' = \gamma E(1 \pm \beta \cos \theta)$$

where θ is the angle of the particle's direction with respect to the shock (so that $p_x = p \cos \theta$). Since we assume that the particle potentially crosses the shock a number of times, take one 'acceleration interaction' to consist of the cycle of crossing the shock twice, so that the particle undergoes two boosts with opposite signs:

$$\begin{aligned} E'_1 &= \gamma E_1(1 - \beta \cos \theta_1) \\ E_2 &= \gamma E'_1(1 + \beta \cos \theta'_2) \end{aligned}$$

Due to the conservation of kinetic energy in the primed frame $E_2 = E'_1$, so the relative change in energy for the whole cycle is

$$\frac{E_2 - E_1}{E_1} = \frac{\gamma^2 E_1(1 - \beta \cos \theta_1)(1 + \beta \cos \theta'_2) - E_1}{E_1} = \gamma^2(1 - \beta \cos \theta_1)(1 + \beta \cos \theta'_2) \quad (1.4)$$

We want to know the average gain in energy, so we must compute the expectation value of the terms which depend on the angles (which may vary with each crossing). Because of the symmetry of the problem both have the same value (in each case the particle approaching the shock will be doing so according to an isotropic distribution in the frame on that side), so we merely need to obtain $\langle \cos \theta \rangle$ for particles with an isotropic angular distribution impinging on a plane (the shock). The isotropic distribution itself is $P(\theta) = \frac{2}{\pi} \cos \theta$ so the expectation value is

$$\langle \cos \theta \rangle = \frac{2}{\pi} \int_0^{2\pi} \int_0^\pi \cos^2 \theta \sin \theta \, d\theta \, d\phi = \frac{2}{3}$$

Plugging this result into Eq. 1.4, and using the assumption that $\beta \ll 1$ to neglect the β^2 terms yields

$$\frac{\langle \Delta E \rangle}{E} = \frac{4}{3} \beta$$

or

$$\alpha = 1 + \frac{\langle \Delta E \rangle}{E} = 1 + \frac{4}{3} \beta \quad (1.5)$$

Thus, the magnetic fields in the vicinity of a shock can meet the requirements to accelerate particles into a power law spectrum. It is then of interest to consider in what locations such shocks might exist which would be able to both accelerate particles to the observed high energies, and also provide enough total flux of cosmic rays to fill the Galaxy or universe.

1.3 Possible Sources

Candidate sources of the cosmic rays (and thus also neutrinos) must meet two major requirements: They must be capable of producing particles at energies up to the highest which have been observed, and they must be both numerous enough and energetic enough to produce the total quantity of cosmic rays which are observed. Naturally, one class of source need not shoulder this entire burden, but the overall simplicity of the cosmic ray spectrum as it has been observed so far suggests that the number of populations with distinct properties which make dominant contributions to the flux is probably not large.

Current large cosmic ray observatories such as the Pierre Auger Observatory and the Telescope Array project routinely observe cosmic rays with energies up to 10^{20} eV, as shown in Fig 1.2, so at least one class of sources must be able to accelerate particles to this energy. In any model (such as the shock acceleration discussed in Section 1.2) in which particles undergo gradual acceleration over a substantial period of time it is necessary to keep the particles inside the acceleration region (with only a limited probability of escape) long enough for them to gain appreciable energy. A simple constraint, due to Hillas, is that the acceleration region must be large enough and the magnetic fields present strong enough to reverse the direction of a particle before it leaves the region. The constraint is then based on the Larmor radius, which is the radius with which a particle with energy E , which we assume moves at the speed of light, and charge Z orbits within a field of strength B :

$$r_L = \frac{E}{ZBc}$$

This turns out, however, to be only an approximate constraint on the object size when we assume that the magnetic fields are disorderly, and so the particle's direction must be changed by a diffusive process. A more complete constraint on the size of the acceleration region, or equivalently on the maximum attainable particle energy is then [11]:




Figure 1.2: The high-energy cosmic ray spectrum observed by the Pierre Auger Observatory between 2005 and 2010 [10]. The spectral break known as the ‘ankle’ is visible at around 4×10^{18} eV as well as an apparent cutoff at around 10^{20} eV.

$$\begin{aligned}
 r &> \frac{E}{ZB\beta c} \\
 E &< ZB\beta cr
 \end{aligned}
 \tag{1.6}$$

where β is the characteristic velocity of the magnetic scattering centers.

Armed with this information, it is then possible to consider known astrophysical objects' sizes and estimated magnetic fields in order to assess their capability to accelerate cosmic rays, although substantial uncertainties remain due to lack of knowledge about properties of the shocks (β) and the composition of the cosmic rays at high energies (Z). A schematic view of such a comparison is shown in Fig. 1.3. Based on this examination, only a few classes of possible sources remain as likely candidates for accelerating the highest energy cosmic rays; among these are Gamma Ray Bursts (GRBs), galaxy clusters, and Active Galactic Nuclei (AGNs), along with the jets associated with some types of AGN. Although they appear unlikely to be responsible for the highest energy production, Supernova Remnants (SNR) are potentially interesting as the producers of lower energy cosmic rays.

Aside from the question of whether a given class of sources is capable of producing cosmic rays at a given energy, there is still the question of whether the production rate in such sources can match the observed flux. This is more challenging to address, as it requires knowledge of both the number of sources (or their frequency of occurrence for transient phenomena like GRBs), as well as the fraction of their emitted energy which they put into accelerating particles.




Figure 1.3: Comparison of estimated object sizes and magnetic field strengths to the requirements for producing the highest energy cosmic rays, based on a figure from [12]. Diagonal lines indicate the requirements for accelerating particles of different charges under different assumptions of the shock properties; for each line, objects below and to the left are excluded.

1.4 Showering Processes

Although the cosmic ray primary particles are, in general, stable (since many of them traverse large distances in space between emission and arrival at Earth) they cannot pass unaffected through the Earth's atmosphere. Instead, they tend to interact with the nuclei present in air, producing secondary particles, which themselves often interact, creating still more particles. This process continues, dividing the energy of the original particle more finely and increasing the total number of particles until the individual particle energies become low enough that production of further secondaries becomes impossible (or at least inefficient), and the size of the shower decreases again as particles are absorbed by the medium. Initial treatments of this process, such as [13] dealt only with the component of these showers consisting of high energy photons, electrons, and positrons, known now as the 'electromagnetic' component of the shower. This treatment, as shown in Fig. 1.4 is sufficient to describe the numerically dominant portion of the shower, but was recognized by its authors to neglect the other types of particles which exist in the shower. At the time this 'penetrating component' was essentially the muons produced in the shower, but is now more broadly understood as the 'hadronic' portion of the shower containing mostly heavy, short-lived hadrons, and the 'muonic' component produced by the decays of various of the hadrons.




Figure 1.4: The electromagnetic component of an air-shower according to the calculation of [13] from 1937. The vertical axis is the relative number of particles, while the horizontal axis is the number of electron interaction lengths traversed. The model (solid line) is compared to observed data, and agrees well until most of the electrons have been attenuated and are no longer the dominant component.

The electromagnetic component is of considerable interest to cosmic ray air-shower detectors, but as it consists solely of particles with low ability to penetrate material, it is of little importance for neutrino experiments (which are usually protected by considerable shielding). The muonic component is interesting

for two reasons, however. Firstly, the great penetration capability of high energy muons means that they are likely to reach even well shielded neutrino experiments, meaning that they form a major background which must be identified or eliminated in order to effectively study neutrinos. Secondly, this portion of the air-shower is itself a source of neutrinos, which must be produced in order for mesons to decay into charged leptons. The simplest case of this is shown in the schematic of an air-shower in Figure 1.5, in which a charged pion decays to a muon, producing a muon neutrino as well. At a later point, the muon itself decays to an electron, yielding another muon neutrino and an electron neutrino. Not only pions, but also heavier mesons, can decay producing neutrinos, and this actually represents the strongest source of neutrinos at GeV-TeV energies on Earth. Their properties will be discussed more quantitatively in Section 2.3.




Figure 1.5: Schematic of important portions of a cosmic ray induced particle shower [14]. Specifically, the decay chains of charged pions, which give rise to the muonic component of the shower as well as neutrinos, and neutral pions, which produce the ‘electromagnetic’ component of the shower are shown. Not shown in detail are the complex behaviors of heavier hadrons which exist in the early stages of the shower.

Chapter 2

Neutrinos

2.1 Neutrino Physics

2.1.1 History

The early history of neutrino physics is strongly connected to the investigation of β -decay, as this was the first easily accessible system in which neutrinos play a role. Around 1910 measurements suggested that the electrons emitted by β -decays of particular isotopes did not have single energies, but there was considerable controversy over why this should be the case and what the form of the spectrum actually was. Experiments such as those of Meitner, Hahn, and Bayer provided indirect evidence of multiple energies being present by showing the complexity of the reduction of β -radiation from a source after passing through differing thicknesses of absorber [15], but more conclusive were direct measurements using magnetic deflection to select decay electrons by energy. By 1913, Rutherford and Robison had performed such an experiment, and concluded that the energy spectrum was a complex combination of discrete lines [16], but in 1914 the experiment was repeated by Chadwick using a more precise measurement technique providing the first clear and direct evidence that the spectrum is a continuum [17]; this result is shown in Fig. 2.1. (Rutherford's result was likely an artifact of the sensitivity of the photographic plates which were used to details of experimental handling.)

The continuity of the β -decay spectrum did not lead immediately to a recognition of the need for the neutrino. A number of discussions and measurements were first required to cement the idea that the nucleus is in fact a quantum system (containing no electrons) and that therefore the energy released in a decay is actually constant [18]. Once the issue was clear, two proposed resolutions were that of Bohr, who argued for relaxing the assumption of energy conservation, and Pauli, who postulated an as yet undetected particle with which the electron could share the decay energy but which interacted so rarely that it was never seen in




Figure 2.1: The energy spectrum of electrons emitted from a mixture of ^{214}Pb and ^{214}Bi as measured by Chadwick in 1914 [17]. Curve A, corresponding to the right axis, shows the number of electrons counted using Geiger’s then-new technique, while Curve B, corresponding to the left axis shows the same result using an older, ionization measurement technique. In both cases the energy spectrum is continuous, and with the exception of the structure in the low energy region it is clearly not composed of the lines claimed previously by Rutherford (whose positions are shown in the small inset beneath the curves).

β -decay experiments [19]. He observed that this particle would be required to be spin $\frac{1}{2}$, electrically neutral, have a mass “not more than 0.01 of the proton mass”, and named it the ‘neutron’ originally, although this name was changed to ‘neutrino’ by Fermi after the 1932 observation of the particle now known as the neutron by Chadwick [20]. With the discovery of the neutron the theory of a nucleus consisting only of protons and neutrons became viable, and on this basis Fermi composed a detailed theory of β -decay with neutrinos which was in fair agreement with various observational results and predicted that the neutrino should have very small or zero mass [21].

After Fermi’s early theory of neutrinos in β -decay, the existence of this particle became commonly accepted in the particle-physics community. However, it was generally considered impractical or impossible to ever observe experimentally. It was not until 1951, after the invention of the nuclear bomb, that Reines proposed that the strong flux of neutrinos from a bomb might be possible to observe with a modestly sized detector [22]. After beginning work on the experiment with Cowan, the two realized that a reactor would actually provide a better source, and so the experiment was eventually conducted in 1955 using the Savannah

River reactor. The result was a clear detection of electron anti-neutrinos at a rate of 2.88 ± 0.22 events per hour [23], for a total of a few thousand anti-neutrino interactions observed.

With the combination of the acceptance of the theory that the Sun produces energy via nuclear reactions [24] and the successful observation of neutrinos from human-made reactors, it was then a natural question whether the neutrinos produced by the sun could be measured in order to learn more about its internal properties. This was undertaken by Davis in his experiment at the Homestake mine, using a large quantity of ^{37}Cl to capture electron neutrinos, but the result was that the observed rate of neutrino interactions was substantially lower than expected from solar models at the time [25]. One solution to this anomaly was that while the reactions in the Sun were expected to produce only one neutrino flavor, they might transform probabilistically into another flavor or flavors (of which the muon neutrino had already been observed [26]) in transit [27]. Various evidence in favor of this hypothesis accumulated, but a clear confirmation was obtained by the SNO experiment in which both charged-current interactions of electron-neutrinos and neutral-current interactions of all neutrinos were measured. Critically, this observation was able to distinguish these types of interactions, showing that the electron neutrino flux was consistent with measurements like those of Davis, while the total flux was at the same time consistent with solar models, as expected when a fraction of the neutrinos have oscillated to other flavors [28].

One other result of note is that low energy (tens of MeV) neutrinos have been detected in one instance from a distant astrophysical source. As supernovae involve nuclear processes they can be expected to produce neutrinos, and in fact calculations suggest that this is the form in which they radiate most of their energy. There has been one opportunity to observe this phenomenon, namely supernova 1987A, which, occurring in the Large Magellanic Cloud, was close enough that low energy neutrinos were detected by multiple neutrino experiments [29, 30], as shown in part in Figure 2.2.

2.1.2 Current Understanding

The neutrino is now understood to be a family of fermions with small, but non-zero mass, with no electromagnetic or color charge, but with a weak hypercharge such that they are affected by the weak force, mediated by the W^\pm and Z^0 bosons. The interactions mediated by these bosons are known as ‘charged current’ (CC) and ‘neutral-current’ (NC), respectively, and the charged current interaction has the interesting capability to change the flavors of interacting quarks, as shown in Figure 2.3.




Figure 2.2: The time distribution of events observed by the Kamiokande II detector, in coincidence with supernova 1987A. The vertical scales are proportional to the energies deposited, with the left scale corresponding to the neutrino-induced electron events, and the right scale corresponding to the background muon events (labeled μ_i).

The total number of neutrino flavors which couple to the Z boson is constrained by collider measurements of the width of the Z resonance in electron-positron collisions, as this is controlled by the number of final states into which the resonance can decay. The result from a combination of experiments is 2.9840 ± 0.0082 types of neutrinos [31], so it is reasonable to take the number of neutrinos as 3, corresponding as expected to the three types of charged leptons, the third neutrino flavor having been first observed in 2000 by the dedicated DONuT experiment [32]. Other types of neutrinos may still exist, but they must either have masses greater than half the Z boson mass, or not couple via the usual weak interaction (‘sterile’ neutrinos).

Currently, the masses of all neutrinos are still unknown, but the most stringent limit on any single neutrino mass is $m_{\bar{\nu}_e} < 2.05$ eV at 95% CL [33], but a tighter limit can be placed on the sum of masses of all active neutrinos using cosmological data; depending on the exact combination of data used the result is in the range $\sum_i m_i \lesssim (0.3 - 1.3)$ eV at 95% CL [34].

Since at least two of the neutrinos have non-zero masses, it is possible that the neutrino may be a Majorana particle, in which case it would be its own anti-particle. This could be identified by observing the phenomenon of neutrinoless double β -decay, in which the two neutrinos which would normally be produced in double β -decay are replaced by a single particle internal to the decay diagram. However, this process has not been conclusively observed, so it remains unclear whether it is extremely rare or forbidden [35].




Figure 2.3: Common interactions between neutrinos and hadronic matter. Interactions mediated by the W^\pm are termed ‘charged-current’ while those involving the Z are ‘neutral-current’. In charged-current interactions a neutrino is transformed into a lepton of the same flavor (or vice versa) while the flavor of the struck quark is changed. In neutral-current interactions the identities of both particles are left unchanged. In both cases energy will be transferred to the quark, which recoils, producing a shower of hadronic particles.

At high energies (> 100 GeV), neutrinos interacting with atomic matter do so primarily via deep inelastic scattering with a single quark, as depicted in Figure 2.3. The cross section for the charged-current interaction of a neutrino with quarks and antiquarks is given by [36]

$$\frac{d^2\sigma}{dxdy} = \frac{2G_F^2 M_q E_\nu}{\pi} \left(\frac{M_W^2}{Q^2 + M_W^2} \right)^2 [xq(x, Q^2) + x\bar{q}(x, Q^2)(1-y)^2] \quad (2.1)$$

where G_F is the Fermi constant, M_W is the W mass, M_q is the quark mass, E_ν is the neutrino energy, $\Delta = E_\nu - E_\mu$ is difference in lepton energies in the lab frame, $x = Q^2/2M_q\Delta$ and $y = \Delta/E_\nu$ are the Bjorken scaling variables, and q and \bar{q} are the parton distribution functions for quarks and antiquarks. (Note that even in normal matter the contribution from antiquarks need not be zero, due to sea quarks.) The neutral-current cross section has the same general form, which rises roughly linearly with the neutrino energy up to 10^4 GeV, above which the Q^2 term in the denominator of the propagator becomes large compared to M_W^2 . A detailed calculation of neutrino-nucleon cross sections from [37] is shown in Figure 2.4, which is based on the HERAPDF1.5 [38] calculations of the parton distribution functions (this is the cross section calculation used for the results in this work).

The growth of the neutrino cross section has important implications for detection. On one hand, it partially combats the typical decrease of fluxes with energy, making detection of high energy neutrinos slightly easier. On the other hand, it also means that at sufficiently high energies the Earth becomes opaque to neutrinos; in practice this effect begins to be noticeable at $\sim 10^5$ GeV. Experiments seeking such neutrinos should therefore expect a partial blind spot in the direction of the Earth’s core, which becomes more severe at higher energies, as shown in Figure 2.5.




Figure 2.4: The cross sections calculated in [37] for high energy neutrinos interacting with isoscalar material (composed equally of up- and down-quarks). On the left the charged-current and neutral current cross sections are compared, while on the right the neutrino and antineutrino charged current cross sections are shown.

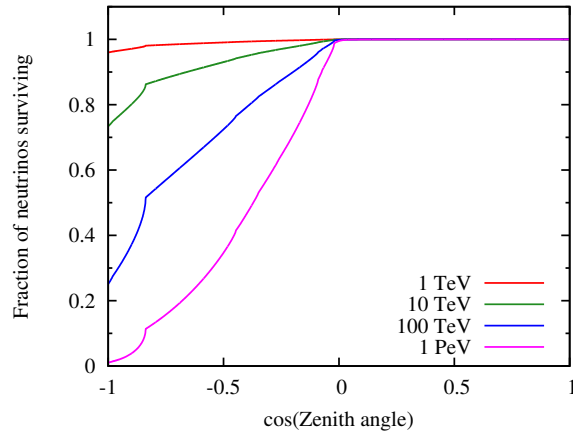




Figure 2.5: Attenuation of neutrinos passing through the Earth as a function of zenith angle, for an observer at (or just below) the surface. At \sim TeV energies neutrinos are barely attenuated even when they pass through the dense core, but at 100 TeV to PeV energies the effect becomes quite strong even for directions passing through only the mantle and crust.

2.2 Astrophysical Production of Neutrinos


While cosmic rays produce extended showers in dense media such as Earth's atmosphere (Section 1.4), they manifestly succeed in escaping from the regions in which they are produced, so these regions must be optically thin for cosmic rays compared to Earth's atmosphere. Therefore lengthy showering processes with many interactions are often less important than single interactions whose final state products are free to decay without losing energy. At high energies, protons may interact with photons, and the cross section for such interactions is particularly large near the Δ resonance, whose dominant decay modes contain pions:




Protons can also interact with other hadrons, and again the dominant final states tend to include pions:



When the pions produced in such interactions decay, they produce either photons, if neutral, or charged leptons (mainly muons) and neutrinos, if charged:



Finally, any muons which are produced will also decay to electrons, producing still more neutrinos:



These processes tend to produce charged and neutral pions in roughly equal amounts, which leads to the interesting possibility of linking photon production fairly directly to neutrino production. Furthermore, the

details of muon production and decay lead to a characteristic ratio of neutrino flavors of $\nu_e : \nu_\mu : \nu_\tau$ of $1 : 2 : 0$, however it should be noted that neutrino mixing implies that any combination of flavors at production will become more similar to 1:1:1 after propagation. A useful rule of thumb can also be derived for the energies involved; for the $p\gamma$ interaction, which will be most of interest in the following discussion, it can be shown [39] that the mean pion energy is $\langle E_\pi \rangle \sim E_p/5$, while the mean neutrino energy is $\langle E_\nu \rangle \sim E_\pi/4 \sim E_p/20$.

2.2.1 Bound on the Total Diffuse Neutrino Flux

Since production of cosmic rays is coupled to production of neutrinos, it should be possible to estimate a limit on the total flux of neutrinos simply from the amount of energy available in the cosmic rays themselves. Such a calculation is performed in [40] and updated with newer inputs in [41]. It is assumed that since the cosmic rays escape from their production sites, these regions must be optically thin for cosmic rays, and the production spectrum of the cosmic rays is E^{-2} . If the cosmic rays lose a fraction of their energy independent of total energy which goes into the production of pions (and thus neutrinos) the spectrum of neutrinos should also be proportional to E^{-2} . The sources themselves are assumed to be distributed through the universe (or equivalently through time) similarly to the formation rate of stars in the nearby portion of the universe (redshift < 2) and constant at greater distances. This bound is shown in Figure 2.6 for the sum of all neutrino flavors, with a numerical value of $\sim 4.4 \times 10^{-8} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$ for an E^{-2} for high energy charged particle interactions in which the production of charged pions is somewhat greater than that of neutral pions (the latter producing no neutrinos on decay). This model is compared to several models of neutrino production in AGN Jets, which the authors of the bound argue should be disregarded as they violate the bound, requiring more energy being put into neutrinos than is available from the cosmic rays at the corresponding energies.

2.2.2 Neutrinos from Active Galactic Nuclei

While the generic bound on neutrino production is in conflict with most models of neutrino production in AGN jets, from which cosmic rays can be expected to escape, there is another plausible mechanism for neutrino production in AGN which is not subject to this constraint: If the neutrinos are produced deep in the core of the object (close to the central black hole) the high matter and radiation densities will be opaque to charged particles (cosmic rays) while the neutrinos can still escape. Such production cannot be directly related to cosmic ray observations, but it can be estimated from X-ray observations of AGN. A calculation




Figure 2.6: The Waxman Bahcall diffuse neutrino flux bound from [41]. Several predictions of AGN jet production of neutrinos are also shown.

of this type is presented in [42], where it is assumed that protons are accelerated in a small volume around the black hole, and then lose energy via interaction with photons (rather than matter since the x rays are observed to escape without absorption features) to produce the observed power law spectra of x rays and emission peak of ~ 10 eV ultraviolet photons. The protons which will interact most readily with the UV photons will have energies of order 10^7 GeV, and the maximum proton energy, limited by the mean free path time for interactions with photons, can be as high as $\sim 10^{11}$ GeV, so the resulting neutrinos will be cut off above $\sim 10^9$ GeV, follow the proton spectrum of E^{-2} down to energies of $\sim 5 \times 10^5$ GeV, and have an approximately uniform spectrum at lower energies. As shown in Figure 2.7 this spectrum (when multiplied by E^2) has a characteristic peak at \sim PeV neutrino energies.

2.2.3 Neutrinos from Starburst Galaxies

A different class of sources from AGN which also have substantial energy output which might in part be channeled into neutrinos are starburst galaxies, or events. These refer to cases of galaxies undergoing elevated rates of star formation, typically for relatively short periods of time, triggered by disturbances like mergers with other galaxies. During such events stellar deaths through supernovae can be expected to rise along with the birth of massive, short-lived stars, and activity from central black holes may increase as well due to additional material becoming available for accretion. These effects can reasonably be expected to accelerate charged particles, and in [43] it is argued that the radio emission of these galaxies can be

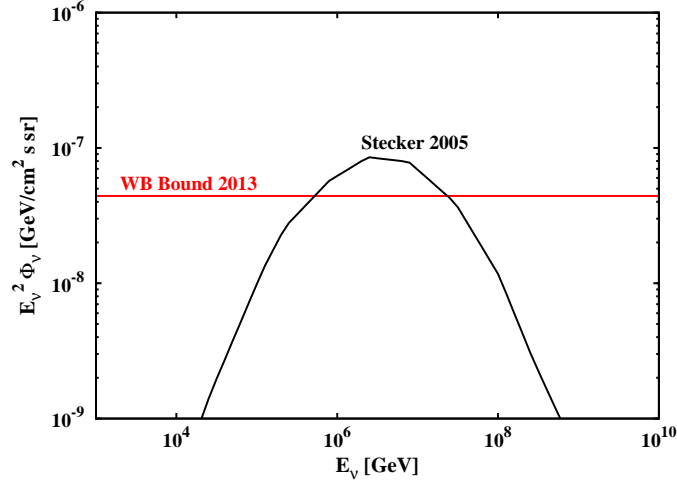


Figure 2.7: The AGN core model of [42] compared to the Waxman Bahcall diffuse bound. Note that this model is not constrained by the bound, so its prediction of larger fluxes at a range of energies is not forbidden.

connected to synchrotron emission from electrons produced when accelerated protons lose their energy to produce pions. Based on measurements of starburst galaxies and estimates of cosmic ray diffusion in our own galaxy the authors conclude that all cosmic ray protons with energies below $\sim 3 \times 10^6$ GeV will lose all of their energy in the denser gas and stronger magnetic fields thought to be present. The proposed synchrotron radiating electrons are produced in the same pion decays as the neutrinos, and so have similar energies, allowing the rate of neutrino production to be determined through measurements of the galaxies radio output. The spectrum of neutrinos will be expected to have the same shape as the proton spectrum, which could theoretically be a power law as hard as E^{-2} , but is likely to be somewhat softer, and the author's tentatively choose $E^{-2.15}$ as the most likely index, although the uncertainty is large. The proton spectrum, and thus the neutrino spectrum, may also contain a 'knee' like the cosmic ray spectrum in our own galaxy, but the shape of this transition and its energy are also unknown, further increasing the flux uncertainty, although it is suggested that this knee may be at higher energies than observed locally. The result of this calculation is the large, shaded band depicted in 2.8. Since this model allows so much latitude, in the rest of this work we will deal only with the $E^{-2.15}$ hypothesis.




Figure 2.8: The starburst galaxy model of [43] is shown as the shaded band, within which the $E^{-2.15}$ hypothesis is shown as a darker green line. A generic GZK neutrino flux (§2.2.5) is also shown.

2.2.4 Neutrinos from Gamma Ray Bursts

A final class of possible sources which should be considered are gamma ray bursts. These events intrinsically involve large energy outputs and the production of high energy particles (certainly including gamma rays), so they are natural candidates for production of charged particles and neutrinos. Gamma ray bursts are generally now assumed to consist of narrow beams of matter moving at relativistic speeds, and these events occur on short timescales, so the sources are likely to vary sufficiently to produce shocks within the beam, and this would allow the shock acceleration of protons, which would then undergo photo-meson production in conjunction with the gamma rays known to be present. In [44] it is argued that the cosmic rays above $\sim 10^{10}$ GeV may be entirely produced by GRBs if equal amounts of energy go into accelerating electrons, which radiate gamma rays, and protons, which will give rise to neutrinos. The observed photon spectra of GRBs from keV to MeV energies can be described by a broken power law $\propto E^{-1}$ at low energies and E^{-2} at higher energies. The proton spectrum is expected to have the same form, and the corresponding break energy for protons is calculated to be of order 10^7 GeV, so the same feature should appear at roughly 5×10^5 GeV in the neutrino spectrum. Furthermore, a second break is expected at higher energies, where the acceleration mechanism loses efficiency, implying a final segment of the spectrum something like E^{-3} . The spectrum of neutrinos from GRBs should thus have a characteristic trapezoidal shape, which is shown in Figure 2.9.

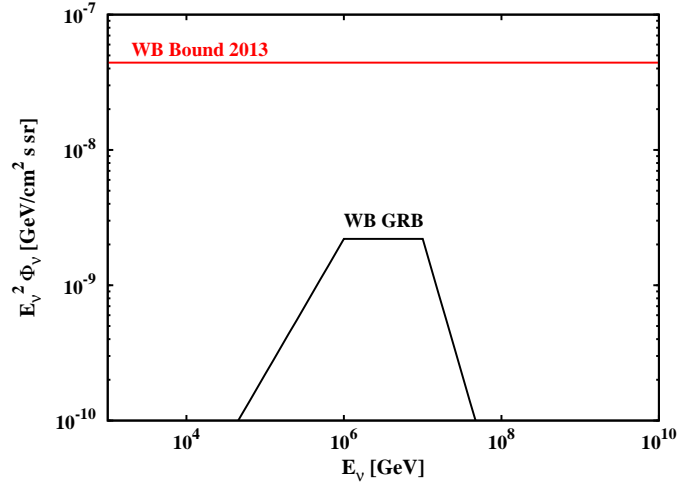


Figure 2.9: The gamma ray burst model of [44] compared to the Waxman Bahcall diffuse bound.

While this line of reasoning is not obviously flawed, it should be noted that dedicated searches, which use the sharp definition of the source in position and time to observe with very little background, have already been performed for neutrinos produced in GRBs [45]. These studies have observed no neutrinos, which severely constrains the possibilities for hadronic acceleration in these sources. It is still possible for cosmic rays and neutrinos to be produced by GRBs, but only if there exists a class of ‘hidden’ GRBs which accelerate cosmic rays but produce no detectable gamma rays. This class of model therefore seems disfavored as a source of diffuse neutrinos.

2.2.5 Cosmogenic Neutrinos

Neutrinos may also be produced by cosmic rays in regions well removed from the cosmic ray accelerators. As noted in Section 1.1, the interaction of cosmic rays with photons via the Δ resonance can occur whenever the total energy in the center of mass frame of the interaction is greater than the Δ mass, so even very low energy photons can become an effective interaction target for sufficiently high energy protons. The original calculations of this effect [2, 46] suggested that the cosmic microwave background photons which fill the universe will therefore attenuate all cosmic rays with energies greater than $\sim 3 \times 10^{10}$ GeV, and later calculations [47, 48] showed that through the usual photomeson production this leads to the production of neutrinos. These neutrinos are formally known as ‘cosmogenic’ neutrinos, but are often referred to colloquially as ‘GZK’ neutrinos. A number of detailed models for this type of neutrino spectrum exist

[49, 50], however they generally share the feature shown by the cartoon included in Figure 2.8 that the neutrino flux peaks at very high energies ($\sim 10^8$ GeV) and is generally quite small, due to the rarity of cosmic rays at such high energies, so GZK neutrinos are mostly outside the scope of this work.

2.3 Atmospheric Production of Neutrinos

Since protons interacting with matter tend to produce neutrinos they can certainly be expected to do so when reaching the Earth's atmosphere just as when encountering matter or radiation closer to their sources. There are, however, some key differences. First, while neutrinos produced at the sources of cosmic rays would be expected to follow the production spectrum of the cosmic rays, the atmospheric neutrinos will derive from the softer spectrum of the cosmic rays after diffusion. Second, most models of astrophysical neutrino production assume production in regions with low enough densities that the products of proton interactions can be expected to decay without further interaction; in the much more dense atmosphere of the Earth, however, energy loss to interaction plays an important role.

2.3.1 Conventional Atmospheric Neutrinos

The 'conventional' atmospheric neutrinos are so named because they have been reliably observed and are reasonably well understood. For reasons which should become clear shortly, the conventional neutrinos are defined as those which derive from pions and kaons, which are in turn produced from the interaction of cosmic ray primary particles, or as the final stages of the complex hadronic products of cosmic rays. All of the particles participating in an air shower are coupled together by differential equations relating their production (from other particles in the shower), decay (to other particles in the shower), and removal from the shower by catastrophic energy loss to the surrounding air. In general, for a particle type labeled by the index i the form of the equation is

$$\frac{dN_i(E, X)}{dX} = -\left(\frac{1}{\lambda_i} + \frac{1}{d_i}\right) N_i(E_i, X) + \sum_j \int \frac{F_{ij}(E_i, E_j)}{E_i} \frac{N_j(E_j)}{\lambda_j} dE_j \quad (2.2)$$

N_i is the number density of particle type i , the first term describes the loss of particles either to interaction (controlled by λ_i) or decay (controlled by d_i), and the second term describes the injection of particles from the interaction of other particle types j of all energies E_j , controlled by the density of j particles, N_j , and the probability for a particle of type j to produce a particle of type i in terms of the energies of both, F_{ij} .

The position coordinate, X , is a slant depth, which accounts for the amount of atmosphere traversed, rather than physical position.

For the purpose of producing a daughter particle type in a decay (for instance if i is π^\pm and the daughters are ν_μ) it is important whether loss of energy through interaction (which will reduce the energy which can go to the daughter) or direct decay is more likely to occur first. If we approximate the atmosphere's density profile as exponential with scale height h_0 , these two processes will be balanced when $h_0 = \gamma c\tau$, where τ is the lifetime of the (parent) particle in its rest frame. One can then solve for the critical particle energy ϵ_i at which this occurs, finding it to be

$$\epsilon_i = \frac{m_i c^2 h_0}{c\tau_i} \sqrt{1 - \left(\frac{c\tau_i}{h_0}\right)^2} \approx \frac{m_i c^2 h_0}{c\tau_i} \quad (2.3)$$

For the upper atmosphere $h_0 \approx 6.4$ km, so for pions ($c\tau_\pi = 7.8$ m, $m_\pi c^2 = 139.6$ MeV) $\epsilon_\pi = 114.5$ GeV, and for kaons ($c\tau_K = 3.7$ m, $m_K c^2 = 493.7$ MeV) $\epsilon_K = 851$ GeV. At lower energies than the critical energy, the particle is more likely to decay than to interact, as its mean lifetime multiplied by its speed (here taken to always be c) is smaller than its mean free path. At higher energies, however, it will be sufficiently boosted that its mean survival distance from decay will be longer than its mean free path, so it is more likely to interact. So,

$$\frac{1}{d_i} = \frac{\epsilon_i}{E_i} \quad (2.4)$$

This derivation made two assumptions: that the particle is at unit slant depth, and travels through the atmosphere vertically. Neglecting the curvature of the atmosphere, the slant depth, X , for a zenith angle θ is related to the equivalent vertical slant depth, X_v , by $X_v = X \cos \theta$, and the probability of interaction increases proportionally with density, which is $\rho = X_v/h_0$, so more generally

$$\frac{1}{d_i} = \frac{\epsilon_i}{E_i X \cos \theta} \quad (2.5)$$

The dependence of this quantity on the parent particle energy has important implications for the energy spectrum of daughter particles. At energies below the critical energy, $1/d_i$ is large compared to $1/\lambda_i$, so the decaying parent particles are able to give their full energy to their daughters, and the daughter energy spectrum will follow the same form as the parent spectrum. At higher energy, the probability of a parent




Figure 2.10: Left panel: The conventional atmospheric neutrino flux predicted by the HKKM calculation [51] for ν_e and ν_μ and for two different zenith angles. The flux is enhanced for both flavors at the horizon because progenitor mesons have more length in which to decay, and at energies greater than ~ 100 GeV the ν_μ flux is always larger than the ν_e flux. Right panel: The zenith angle distributions of the ν_μ and ν_e fluxes at a neutrino energy of 1 TeV.

to confer its full energy on its daughters is proportional to $1/E$, so the daughter particle spectrum will be steepened by one power of E .

It can then be shown by integrating the cascade equations (Equation 2.2) for a power law primary spectrum with index γ that the spectrum of muon neutrinos produced by pion and kaon decays is [9]

$$\frac{dN_\nu}{dE_\nu} \approx \frac{N_0(E_\nu)}{1-Z_{NN}} \left(\frac{A_{\pi\nu}}{1+B_{\pi\nu} \cos \theta E_\nu/\epsilon_\pi} + 0.635 \frac{A_{K\nu}}{1+B_{K\nu} \cos \theta E_\nu/\epsilon_K} \right) \quad (2.6)$$

$$A_{i\nu} \equiv Z_{Ni} \frac{(1-r_i)^\gamma}{\gamma+1} \quad (2.7)$$

$$B_{i\nu} \equiv \left(\frac{\gamma+2}{\gamma+1} \right) \left(\frac{1}{1-r_i} \right) \left(\frac{\Lambda_i - \Lambda_N}{\Lambda_i \ln(\Lambda_i/\Lambda_N)} \right) \quad (2.8)$$

$$r_i \equiv \frac{m_\mu^2}{m_i^2} \quad (2.9)$$

where the Z_{ij} are the spectrum weighted moments $Z_{ij} \equiv \int_0^1 x_L^{\gamma-1} F_{ij}(x_L) dx_L$, Λ_i is an attenuation length $\frac{1}{\Lambda_i} \equiv \frac{1}{\lambda_i} (1 - Z_{ii})$, and the weighting factor of 0.635 is the branching ratio for kaons to decay to a μ/ν_μ pair. $A_{i\nu}$ has the same energy dependence as the primary flux, so the neutrino spectrum matches the primary spectrum at low energies (where $E_i < \epsilon_i$) and is one power steeper at high energies. The transition between these two regions depends on zenith angle, since the effective critical energy increases with θ , becoming maximal (although not infinite, as suggested by this treatment, which has ignored the curvature of the atmosphere) at the horizon.




Figure 2.11: The separation of the conventional ν_μ flux into pionic and kaonic components as a function of neutrino energy at the horizon (left) and zenith angle at an energy of 1 TeV (right). The shorter lifetime of the kaons makes their contribution to the flux less dependent on angle, while the pion flux increases sharply at the horizon where pions are more likely to have enough distance to decay.

The spectrum of electron neutrinos produced from pions and kaons is not as simple, because the branching ratios from either of these mesons directly to states with electrons and electron neutrinos are small, of order 10^{-4} for pions and 10^{-5} for kaons. Instead, electron neutrinos are mostly produced by more complex decays of kaons with three particles in the final state, which are less suppressed.

The results of a detailed numerical calculation of conventional atmospheric neutrino fluxes [51] are shown in Figure 2.10 and Figure 2.11. In the first, the suppression of the ν_e flux with respect to the ν_μ flux can be seen, as well as the enhancement of both fluxes at the horizon, compared to the vertical. In the second the different behaviors of the portions of the ν_μ flux produced by pions and by kaons can be seen: The higher critical energy of the kaons at all zenith angles mean that they are the dominant source of neutrinos at high energies despite being less numerous overall than pions. The enhancement effect at the horizon is clearly visible in the flux as a function of zenith angle, and it can be seen that this effect is more pronounced for pions than for kaons, due to their lower critical energy.

As a technical detail, it is important to note that the HKKM flux calculation [51] was intended for use at energies below 10 TeV, where the knee structure of the cosmic ray spectrum is mostly irrelevant. Unfortunately, this work is interested in neutrino energies for which this behavior is very important, so it has been necessary to apply after-the-fact corrections to adapt the original calculation to newer models of the cosmic ray spectrum [52]. This was done in [53] by recalculating the spectrum weighted moments for both the original and new input cosmic ray fluxes to create a correction factor. The resulting factor is plotted in




Figure 2.12: The correction to the HKKM flux calculation to include the effects of the CR knee. The main change is a further suppression of the flux above neutrino energies of ~ 100 TeV.

Figure 2.12, where it can be seen that this change introduces a further suppression of the conventional ν_μ flux at energies above ~ 100 TeV.

2.3.2 Prompt Atmospheric Neutrinos

The previous section discussed production of neutrinos in the atmosphere only from two types of mesons: pions and kaons. Many other types of mesons can produce neutrinos, but they differ from the first two in that they are substantially heavier and shorter lived (they decay ‘promptly’, thus the naming of the neutrino flux they produce the ‘prompt’ neutrino flux). Their greater masses mean that these heavy mesons are less frequently produced, and so they should create much lower fluxes of neutrinos (indeed the flux from heavier mesons has not yet been distinguished from the conventional flux from pions and kaons). Their shorter lifetimes mean that these mesons can also be expected to produce qualitatively different fluxes from the lighter mesons, since they have much higher critical energies (4.3×10^7 GeV for the D^\pm , which would produce the next largest flux). These high critical energies mean that the spectrum of the prompt neutrino flux can be expected to follow the cosmic ray spectrum up to much higher energies, and will not depend on zenith angle.

Calculating this flux is also quite difficult, because it requires calculating the production rates of heavy mesons at high energies and in momentum regions where accelerator data is not available. One such calculation [54] is shown in Figure 2.13, where it can be seen that the prompt flux is harder than the conventional

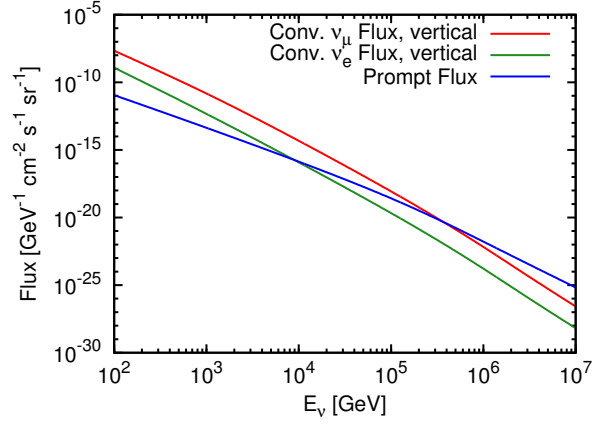


Figure 2.13: Comparison of the prompt atmospheric neutrino flux from the ERS calculation [54] to the conventional flux from the HKKM model. Note that the prompt flux is essentially the same for all zenith angles and for both ν_e and ν_μ , so only one of these is shown. All fluxes are shown with the correction for the cosmic ray knee [53] included.

flux (as expected) and should therefore become the dominant source of atmospheric neutrinos at high energies. The decay of heavier mesons to electrons (with electron neutrinos) is not suppressed as for the lighter mesons, so the prompt flux of ν_e should be essentially the same as the prompt flux of ν_μ . It should be noted that this calculation was also originally ignorant of the cosmic ray knee, so we again apply a correction calculated in [53].

Chapter 3

Detector

3.1 Detection Principle

In order to detect neutrinos, it is generally necessary to wait for them to interact with charged particles, whose disturbance or excitation can be observed via electromagnetic means. Either charged-current or neutral-current interactions will serve this purpose; in the former case, a charged lepton is directly produced, and in either case, the particle in the medium with which the neutrino interacts will recoil. Once a charged particle has been excited, if the medium is transparent and the particle has sufficient energy for its velocity to be greater than that of light in the medium, Čerenkov light is emitted. For neutrinos of TeV energies the velocity requirement on the secondaries is easily met.

The Čerenkov emission is emitted at an angle from the direction of the particle determined by the index of refraction of the medium, n and the velocity of the emitting particle, β , which in this context can usually be approximated as 1 [55]:

$$\theta = \arccos\left(\frac{1}{n\beta}\right) \quad (3.1)$$

The spectrum of radiation is given by the Frank-Tamm formula [56]:

$$dE = \frac{e^2}{4\pi} \mu(\omega) \omega \left(1 - \frac{1}{n^2(\omega)\beta^2}\right) dx d\omega \quad (3.2)$$

where E is the total radiated energy, x is distance traveled by the emitting particle and ω is the frequency of radiation. Neglecting the dependence of the permeability, μ , and the index of refraction on the frequency, the factor of the frequency indicates that the most energy is radiated at high frequencies (short wavelengths). In practice, an ultraviolet catastrophe is prevented by the index of refraction of any real material going to

one for sufficiently high frequencies, so that at some frequency the angular criterion can no longer be satisfied by a particle of finite energy.

A light detector located in a large volume of transparent medium may then detect the arrival of neutrinos by measuring the Čerenkov light produced by secondary particles. This process is shown schematically for ice as a medium in Figure 3.1. Assuming that the medium is dense, there are two major topologies of events which may be observed. The first may be grouped under the label of ‘cascades’: If the charged secondaries only travel relatively short distances (where shortness must usually be judged in comparison to the position resolution of the detector) while producing further secondaries, the dense shower of particles will produce a somewhat diffuse aggregate light emission, and whether its length can be directly observed at all depends on the detector. The events of the second topology are labeled ‘tracks’, corresponding to particles which interact relatively little in the medium, and so may travel long distances emitting light as well as stochastically scattering and producing occasional daughter particles (which may then develop into cascades). ‘Tracks’ are so named because they extend far enough that their extent is obvious even in sparsely instrumented detectors, and so the direction of the particle’s travel is readily apparent.

Of the realizable types of neutrino interactions, all neutral-current scatterings tend to be observed as cascades, since recoiling quarks rapidly produce showers of hadrons and eventually the usual electromagnetic component. A charged-current interaction by an electron neutrino likewise produces as a cascade, which is purely electromagnetic from the start. The range of a muon in relevant media (water, ice, and even rock) can be quite large, however, so charged-current muon neutrino interactions produce tracks. It should be noted that there is still a recoil at the neutrino interaction vertex, so a cascade is actually produced as well, but in many cases an arbitrary section of the muon path is observed while the initial cascade is far away, and can be neglected. Finally, charged-current interactions of tau neutrinos provide a more complex version of the muon neutrino case; the tau path length is generally limited by the particle’s short lifetime (a highly relativistic tau can be expected to travel only about 10^{-8} times as far as a muon of the same energy before decaying, neglecting interaction with the medium) and the majority of the decay modes of the tau involve hadrons, so these interactions tend to produce short tracks with initial and final cascades. Depending on the energy of the tau it is possible for both cascades to take place inside a detector, and at lower energies they may not even be well separated.

This work will focus on the observation of muon ‘track’ events, so the details of muon energy loss, which regulate both the total range the particle may travel and the total emitted Čerenkov light, are of

concern. The major mechanisms by which muons lose energy in a material are shown in Figure 3.2, and include ionization of atoms in the medium, photonuclear interaction, bremsstrahlung, and pair production of electrons and positrons. At low energies the loss due to ionization is dominant, and the energy lost in this manner is nearly constant with muon energy. The other forms of energy loss, however, are approximately proportional to muon energy, so they dominate for high energies. The expected range of a muon then depends approximately linearly on the muon energy for low energies, but grows only logarithmically at high energies. The impact of the energy loss mechanism on light emission derives from the fact that the energy lost by the muon goes into accelerating secondary particles, producing additional Čerenkov radiation on top of that produced by the muon alone. As a result, the total energy loss rate of high energy muons being proportional to their energy means that they also trigger light emission proportionally to their energy. This can be exploited to measure the energy of a muon even when its full path cannot be observed, but only when the muon energy is significantly greater than 100 GeV.

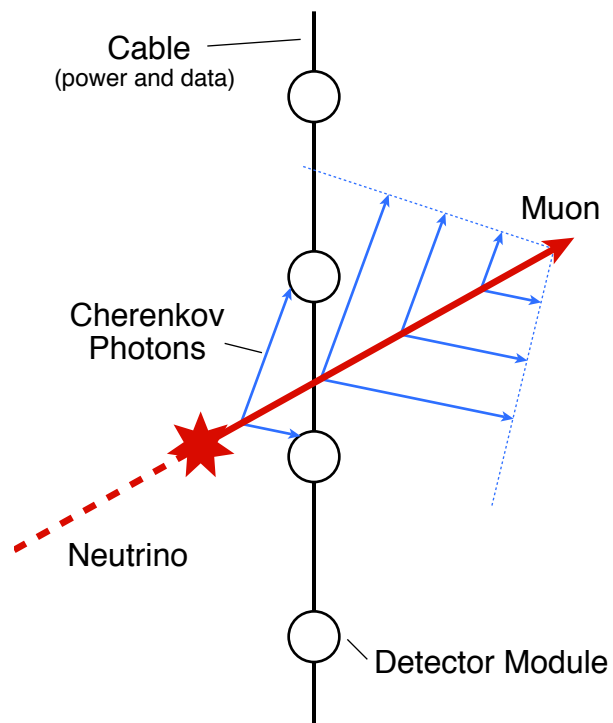


Figure 3.1: Schematic of the detection process. A muon neutrino (dashed red line), entering from the lower left, undergoes a charged-current interaction to produce a muon (solid red line). The muon is of sufficiently high energy to exceed the phase velocity of light in the ice, so Čerenkov photons (solid blue lines) are produced, at an angle 41° (for 400 nm photons) from the muon's direction. Some of these photons strike detector modules and may be recorded.




Figure 3.2: Cross sections for different mechanisms of energy loss by muons in ice as a function of muon energy from [57]. At high energies, stochastic losses (particularly Bremsstrahlung and pair production of secondary particles) dominate over continuous losses from ionizing the medium. Decay is not actually an energy-loss process, but has been included for comparison by multiplying the decay probability by the particle energy.

3.2 The IceCube Detector

The IceCube detector is a large volume neutrino detector utilizing the Čerenkov effect in the South Polar ice [58]. The detector consists of 5160 partially independent light-sensitive modules (known as ‘Digital Optical Modules’ or ‘DOMs’) lowered into boreholes drilled through the ice sheet using hot water. Within each borehole, the modules are connected by a cable which provides mechanical support during the lowering process (and until the ice refreezes), electrical power, and communications. The assembly comprising 60 modules and their support cable within one hole is referred to as a ‘string’, and 86 of these assemblies have been installed. 78 of the strings are arranged approximately on a hexagonal grid with a spacing between strings of 125 m, and the modules on these strings are separated from each other by 17 m vertically. The modules are located on the portions of the cable starting at a depth of about 1.45 km down to a depth of nearly 2.45 km; the full volume so instrumented is approximately 1 km^3 . The final 8 strings make up a subdetector known as DeepCore, which is located at the bottom center of the main array. These strings are more closely spaced, about 60 m apart horizontally, and the modules are also more closely packed vertically, about 7 m apart. The high density of DeepCore is intended to enable the detection of neutrino interactions down to $\sim 10 \text{ GeV}$ [59], but these modules also participate in high energy detections spanning larger portions of the full detector. Finally, the detector includes 324 modules placed inside sealed tanks on the surface of the ice sheet (or rather shallowly buried by drifting snow) which form the IceTop subdetector for cosmic ray air-showers. That portion of the detector will not be used in this work.

Each detector module contains a single large area (10 inch diameter) Hamamatsu R7081-02 photomultiplier tube protected inside a 13-inch glass pressure sphere. This photomultiplier is usually tuned to a gain of 10^7 to allow the detection of individual photons, and has a quantum efficiency of around 0.25 for wavelengths near 400 nm [60] (the majority of the modules in the DeepCore subdetector have an enhanced quantum efficiency to further aid detecting low energy events). Each module also contains its own digitization electronics, allowing the output of the photomultiplier to be precisely recorded without distortion by transmission over a long cable, and a collection of LEDs which can be used for calibration purposes [61]. The details of the digitization are discussed in Section 3.3.

Because IceCube is located in a remote location where conditions are harsh for most of the year, construction required several years to complete (from 2003 to 2010). While construction was not complete, the detector was operated in partial configurations, which have been designated by the number of strings included. The pattern of installation can be seen by the colored markers in Figure 3.3. For this work it is




Figure 3.3: Cut-away view of the IceCube detector, showing the strings (vertical lines) and the detector modules at depths between 1.5 and 2.5 km (small spheres). The volumes approximately corresponding to the DeepCore subdetector and the older AMANDA-II detector are marked as cylinders. The colored dots at the ice surface indicate the year of deployment of the string beneath; the IC79 partial detector included all strings except those with orange markers, while the IC86 detector included (and continues to include) all strings.




Figure 3.4: Schematic view of the IceCube Digital Optical Module. The photomultiplier tube is oriented downward, with its upper end surrounded by the readout, communications, and calibration electronics.

relevant that the strings indicated with orange markers were not installed until the end of 2010, so of the two years of data used in this analysis, 2010 and 2011, the former used only the first 79 strings to be installed.

3.2.1 Ice

The IceCube detector exploits the naturally occurring ice of the south polar ice sheet. This was chosen not only because the quantity of ice far exceeds what could be produced artificially, but also because the quality of the ice is extremely high. Nonetheless, its properties do vary due to the combination of slow deposition over time by snowfall, which contains varying amounts of dust and volcanic ash, and pressure and temperature variations within the ice sheet, and these properties must be measured so they can be included accurately in both reconstruction and simulation of data observed in the detector. The main properties of interest are the scattering and absorption of photons, since these affect the relationship of particle paths and energies to the light recorded in the detector modules. The primary method for measuring these effects is to use the calibration equipment of the modules themselves: one or more modules emit light with their built-in LEDs while the surrounding modules record the results. Once a sufficiently large corpus of data is collected

spanning a variety of regions in the ice, the optical parameters can be fit to it [62]. A few general features (shown in Figure 3.5) arising from the fit are worth noting: Foremost, at depths near 2000 m there is a large concentration of dust, which causes far greater absorption (as well as scattering) than in the rest of the ice volume. Additionally, the deepest ice, below the dust layer is extremely pure, with absorption lengths in excess of 100 m, allowing light to be observed from quite distant emitters, while above the dust layer the ice is not quite as clean, and so both scattering and absorption are greater. More recent studies of the ice have revealed additional complexity, such as anisotropic scattering in the horizontal plane [63]. Unfortunately, it was not possible to include these newer results in this work.




Figure 3.5: Results of two different measurements of the optical properties of the South Pole ice for light with a wavelength of 400 nm. This work uses the newer ‘SPICE MIE’ model of [62], from which this figure is also taken.

3.3 Data Acquisition and Processing

In order to reliably measure single photons with high precision, the output of the photomultiplier in an IceCube detector module is digitized locally before transmission to computers on the surface of the ice sheet [64]. Several digitizers are used to obtain good results over a range of signal types: For good time resolution of detections a high-frequency digitizer (a custom ASIC, known as the ‘Analog Transient Waveform Digitizer’ or ‘ATWD’), calibrated to operate at 300 MHz (one sample per 3.33 ns) is used. To support such rapid sampling this digitizer does not fully operate at all times, but instead stores analog levels in an input buffer (128 samples in length) until instructed to digitize. However, this limits the time span which it can record to ~ 427 ns, so a second, lower frequency digitizer, operating at 40 Mhz (one sample per 25 ns) is used as well. The low-frequency digitizer can readout continuously, but readouts are deliberately limited to 256 samples (spanning 6.4μ s) to control data volume. Each digitizer has 10 bits of output bandwidth per sample, which does not provide enough dynamic range to include both single photon signals and the bright, many photon signals (whose amplitude is limited in practice by the saturation of the photomultiplier, see Figure 3.12), so multiple digitization channels are employed with different gains. Only the high-frequency digitizer includes this redundancy, as complex, high brightness signals are expected to have only short durations. The highest gain channel is used for typical single photon detections, and the channels with progressively lower gains are read out as well when the preceding channel has saturated.

In order to determine when digitization should occur, each module has a discriminator attached to the output of the photomultiplier which is calibrated to fire if the voltage exceeds one quarter of that expected for an ideally amplified single photoelectron. When the onboard computer observes that the discriminator has fired and if the high-frequency digitizer is not busy from a previous readout request, readouts are triggered from both the high- and low-frequency digitizers. This process, or its output, is referred to as a ‘launch’, and may contain one to four digitized waveforms. The high-frequency digitizer requires 29μ s per channel digitized, plus 950ns to clear each undigitized channel, plus a constant 225ns to complete one readout and be ready again, yielding a variable dead time from $32 - 88\mu$ s. In order to combat this, a second, identical digitizer is included in each module, so that when one digitizer is busy, the other may still be used. The worst-case dead time is still 82μ s, but repeated signals bright enough to require digitizing all gain channels are sufficiently rare that under normal conditions the module experiences negligible dead time [65].

While each module is capable of detecting and recording signals independently, it is connected to its immediate neighbors on the same string by short cables which allow messages to be exchanged about when

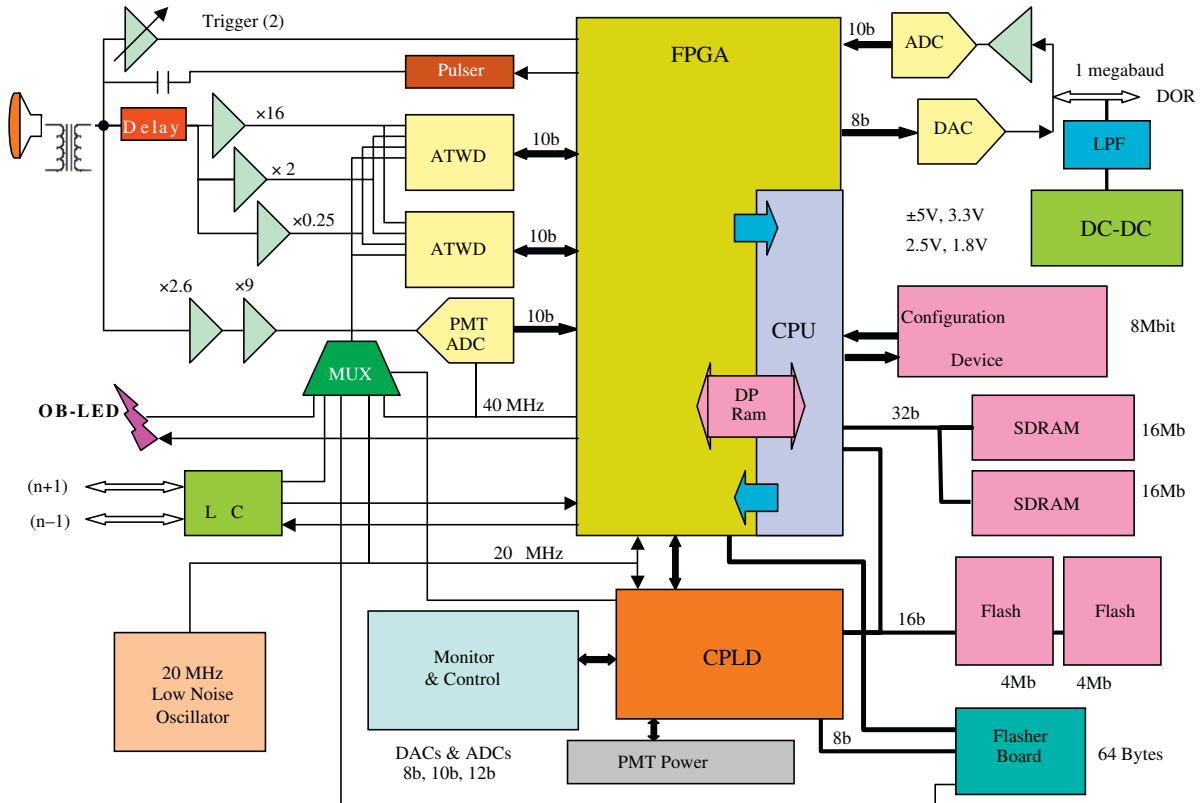


Figure 3.6: Schematic diagram of the detector module readout electronics [64]. Note that following the delay line connected to the the photomultiplier output (red box labeled ‘Delay’) are three parallel amplifiers with different gains, each of whose outputs are connected to channels of both high-frequency digitizer chips (labeled ‘ATWD’), while a pair of series amplifiers whose combined gain is also relatively high connects the photomultiplier to the low-frequency digitizer (labeled ‘PMT ADC’). Each ATWD has a fourth digitizer channel which can be used to record other signals from within the electronics (via the multiplexer labeled ‘MUX’), but is not connected to the photomultiplier for recording physics signals.

a possible signal has been detected (specifically, when the triggering discriminator has fired). This system, referred to as ‘local coincidence’ is used to throttle the readout of data for isolated signals, as these are less likely to be reconstructable particle events. Under their current configuration, IceCube modules produce two types of readouts (‘launches’). One case is that discriminator fires, digitization begins on both the high- and low-frequency digitizers, and within a $1\mu\text{s}$ window a message is received from a neighboring module (or the message may be received up to $1\mu\text{s}$ before the firing of the discriminator, or may originate from a module one neighbor removed). In this case, the readout proceeds, and 128 samples are recorded and transmitted from one to three high-frequency channels along with 256 samples from the low-frequency digitizer. This is known as a ‘Hard Local Coincidence’ or ‘HLC’ launch. If, however, no local coincidence signal is received within the required window, digitization of the high-frequency channels is aborted, and only 16 samples are recorded from the low-frequency digitizer. Of these 16 samples, only three are transmitted: The sample with the largest value and the immediately adjacent samples. This is known as a ‘Soft Local Coincidence’ or ‘SLC’ launch. The inclusion of SLC launches in the final data stream gives extra information about low energy events and the fringes of high energy events with minimal increase in the total data rate or the dead time of the detector modules (as restarting the high-frequency digitizer after an abort can be done much faster than a full digitization and normal restart).

Although each detector module makes decisions about when to transmit data, and there is even limited coordination among modules in the form of the local coincidence, it is desirable to add more sophisticated logic to the system to determine which periods of data are likely to contain usable information for physics purposes. To do this, a system of ‘triggers’ is implemented in software on the computers which aggregate the detector’s output on the surface. A number of triggers are available and in use, but only the most basic of these are actually used in this work. When a trigger determines that its criteria are met, it defines a ‘readout window’ of time within which data should be further considered by the system. Data falling within the union of all overlapping readout windows is forwarded to the next stage of processing as a unit (an ‘event’ for the purposes of the software), while launches outside of any readout window are discarded.

The trigger most commonly used to select potentially interesting physics events is the ‘Simple Multiplicity Trigger’ (this is also sometimes incorrectly rendered as ‘Simple Majority Trigger’, although no majority is involved), which simply watches for a threshold number of detector modules launching within a time window. To prevent the trigger from reading out uncorrelated noise in widely separated parts of the detector, it is configured to consider only HLC launches (those with local coincidence) as counting toward its threshold.

Standard settings for this trigger, called ‘SMT-8’ are a threshold of at least 8 HLC launches within a time window of $5\mu\text{s}$ and a readout window spanning from $4\mu\text{s}$ before the time at which the condition was met to, $6\mu\text{s}$ afterward.

Finally, the detector records many legitimate particle events which are nonetheless uninteresting; the majority of these are low energy muon bundles from cosmic ray air-showers which occur above the detector. The trigger has neither the complexity to understand what distinguishes these from other events, nor the computational capacity to add such complexity since it must operate on the entire stream of data from all detector modules. After the reduction in data rate provided by the trigger system, however, it becomes feasible to employ reconstruction techniques to classify events. This task is performed by a system named ‘PnF’ (Processing and Filtering), which runs a standardized set of reconstructions on all events produced by the trigger readout and applies selection criteria known as ‘filters’. PnF is able to process events concurrently, and can tolerate a latency of order tens of seconds, which gives room for substantially more detailed analysis. Again, a number of filters are in use, each typically with a different physics goal. Each filter has an associated ‘prescale’ value which may be used to restrict its output; for a prescale of N all but 1 (selected randomly) in N of the events which the filter marks as interesting will be ignored. After all filters have considered an event, if any filter passes it and that filter’s output is not ignored due to its prescale, the event is queued for transmission from the South Pole Station to data storage centers in the northern hemisphere via satellite. The total bandwidth allocated for this transmission is approximately 100GB/day.

3.4 Event Reconstruction

Interpretation of IceCube data requires reconstruction of the properties of the particles which produced the recorded signals. What properties are of interest, and what techniques are used to estimate them, vary according to the type of data under consideration and the physics goal of the analysis. For this work, only track-like muon events are considered, and the only necessary properties (or at least the intersection of necessary and calculable properties) will be the directions and energies of the muons.

3.4.1 Waveform Feature Extraction

The raw digitized readouts produced by the IceCube detector modules are not a particularly suitable starting point for high-level event reconstruction: They include substantial effects from the electronics themselves, and may contain information in a rather sparse form, since the record of the arrival of a photon




Figure 3.7: Example of waveform unfolding. The calibrated waveforms are shown as the stepped blue and black solid lines and correspond to the left vertical axis. The unfolded basis functions are depicted as red lines, with heights proportional to the number of photoelectrons they represent, corresponding to the right vertical axis, and the waveform predicted by the superposition of basis functions is superimposed over the actual waveforms as the dashed blue and black curves (where each curves corresponds to the digitizer shown in the same color).

may be smeared across many samples. For this reason, the waveforms are first calibrated, using the known properties of each module’s electronics to recover as closely as possible the analog voltages which were digitized, to counteract the complicating effects of the inductive coupling between the photomultiplier output and the digitizers, and to remove the time delays introduced by the hardware between when a photon is absorbed by the photomultiplier and the resulting signal is recorded in the digitizer. After the waveforms are calibrated, they are unfolded in terms of basis functions designed from the average response of the combined photomultiplier-amplifier system to an ideally amplified single photoelectron. This process is called ‘feature extraction’. A separate template function is used for each amplification channel (since each has a measurably different frequency response) and the data from all overlapping digitizer outputs is unfolded jointly, using a non-negative least-squares algorithm [66]. This provides the best possible estimate of the times at which incoming photons were collected by the module, and largely, if not entirely, removes detector hardware effects from requiring consideration. Since the output is not strictly photon times and amplitudes, but only a statistical approximation, they are typically referred to as ‘reconstructed pulses’ or simply ‘pulses’.

3.4.2 Directional Reconstruction of Track-like Events

Reconstruction of the direction of a track is generally intertwined with reconstruction of its position and time as well. This yields six variables: two directional coordinates, and three positional coordinates, and a time coordinate; however, the time is degenerate with the positions, so only 5 variables must be solved for in total. The main information which is useful in constructing this solution is about timing: the amount of light produced by a particle typically depends somehow on its energy, but the shape of the distribution of light from it mostly depends only on where it went and when, since it will always emit light according to the same geometric pattern.

The most basic reconstruction algorithm typically used for muons in IceCube is known as ‘LineFit’. This is a near-trivial least squares linear regression to the times of observed pulses, and is usually further restricted to consider only the time of the first observed photon at each detector module. This last simplification deserves some explanation, after all, all emitted photons carry information about their emitter and it seems undesirable to lose this. However, it is important that the the time distribution of photons arriving at a given point in the detector from an emitter at some other point is governed in large part by the scattering properties of the ice between and around the two points. If the scattering length is short, photons will undergo many scatterings, and may take substantially longer paths than they would have otherwise, causing them to arrive later, producing a broadening in the time distribution. An algorithm like LineFit is not prepared to deal appropriately with this complexity, so it is better to remove it, if possible. However strong the scattering it is still possible, if increasingly unlikely, for a photon to travel from the emitter to the receiver without scattering, and choosing between any two photons, the one which arrives earlier will generally be closer to this ideal. It is therefore not completely incorrect to conduct the LineFit algorithm on the set of first pulses, and it proves to be a useful ‘first guess’ for a directional reconstruction, which provides a starting point for other, more detailed algorithms.

Because LineFit uses a physical model which is not strictly correct, it is somewhat vulnerable to data points which are outliers. While the assumptions made by this algorithm are not badly wrong for light which travels only short distances from the emitting muon, even the earliest pulse recorded by a module may still be from a scattered photon and will not agree well with its simplistic picture for any possible set of muon parameters. Likewise, the algorithm assumes that all pulses must originate from the muon, even though independent noise is always present in the detector. A way that these problems can be treated is to remove or reduce the influence of these points on the fit. This can be accomplished by applying a cleaning

operation to remove pulses which look likely to have been significantly scattered, due to arriving later than neighboring pulses despite the necessary distance of travel assumedly being similar, and by replacing the pure least-squares weighting in the fit with a weighting which assigns lesser penalties to distant pulses which never fit well (and so are likely to be unrelated noise) [67]. This combination of improvements is referred to as ‘Improved LineFit’, and has largely superseded the plain LineFit in IceCube usage.

A still better result may be achieved by explicitly accounting for the timing effects due to the ice, rather than trying to avoid them entirely. One means of doing this is to use the so-called Pandel Function to describe the expected distribution of photon arrival times at every module depending on the parameters of the muon track, and perform a likelihood fit to best match the expected distributions to the observed data [68]. The Pandel Function is a simple analytic model of the time distributions which arise from diffusive scattering over varying total distances [69]. It does not, however, incorporate any treatment of the layered structure of the ice, the fact that the scattering is variable as a function of depth. In spite of this limitation, using the Pandel Function accounts for the distribution of photon arrival times widening at greater distances from the muon, unlike LineFit’s uniform treatment, which is a substantial improvement to the model. Attempting to use this more precise model directly tends to reveal an imprecision in the data itself, namely that the flight time of a photoelectron through the photomultiplier varies by around 2 ns [70], the times reported by different detector modules are only synchronized to about 1-2 ns [64], and the feature extraction unfolding has a precision of around 1 ns for pulses which are isolated in waveforms when data from the high frequency digitizer is available and about 8 ns when only low frequency digitizer data is available [66]. These uncertainties yield a combined precision on observed photon times of $\sim 2.5 - 8.5$ ns, which is not always small compared to the timing variability predicted by the Pandel Function, and may cause pulses to arrive at times which appear earlier than the earliest possible time from simple consideration of geometries and velocities, before which the Pandel Function is defined to be zero, causing the model to break down when applied to such data. This can be addressed by accounting for the timing uncertainty in the model by convolving the Pandel Function with an additional Gaussian term for the uncertainties [71]. Finally, a constant term can be added to the likelihood function to account for the possibility of noise photons appearing at any point in the detector at any time, without regard for the time or location of the muon.

This more detailed likelihood fit for muon geometries is usually applied to IceCube data in two variants. The first is known as ‘SPEFit’ because like the usual case of LineFit it considers only the first detected photon on each detector module (so it operates only on single photons or photoelectrons). For this case a

Gaussian term for timing uncertainty with a standard deviation of 15 ns is convolved with the Pandel time PDF. It should be fairly obvious that this is not strictly correct on multiple levels: If the Pandel Function is defined to give the PDF for the arrival of any photon, it cannot be the correct PDF for the time distribution of the first photon potentially selected out of many (which will be biased toward earlier times). Furthermore, the applied time jitter is substantially larger than can actually be justified by what is believed to be known about the detector’s timing properties. These details mean that the SPEFit compares photon arrival times with PDFs which are considerably widened, imposing a limit on the possible precision in the fit, but leaving it fairly robust against other systematic errors in the model and poor seed values for the muon properties to be fitted. SPEFit is typically used following LineFit, using the output of the simpler reconstruction as a seed, but is not usually taken as the final reconstruction. For a high precision, final reconstruction, the other common variant of the Pandel-based likelihood, known as ‘MPEFit’ is frequently used. In this model, only the time of the first detected photon on each module is used, as in SPEFit, but the timing PDFs against which that photon’s time is compared are corrected to account for the selection bias of always choosing the first photon. In addition, a more physical timing jitter width of 4 ns is applied in the Gaussian convolution step.

3.4.3 Event Splitting

The preceding section depends on the assumption that the observed event in fact contains a single muon which has a well-defined position and direction to be reconstructed. Unfortunately, there is no strict enforcement of this assumption in real data. In fact, given a detector trigger rate of ~ 3 kHz, and a readout time of $10 \mu\text{s}$, per-trigger simple statistics imply that around 3% of triggered events will in fact, contain within their readout window a second event which could have triggered the detector on its own. 3% is perhaps not a large fraction, but this simple analysis neglects the influence of the selection by the filters: In fact the output from the filter used in this analysis contains $\sim 10\%$ events with a second particle in the detector (of which 10% or about 1% of the total contain a third particle, and so on). The reason for this bias is that the filter attempts to select for neutrino events which go up, rather than atmospheric events which go down, and selects more energetic events over less energetic ones (see Section 4.1). Coincident events are mostly the coincidences of down-going muons, but they confuse reconstructions designed for single muon tracks and tend to be assigned reconstructed directions entirely randomly on the sky, making large numbers of them look up-going even though they are not. Furthermore, if an energy cut is designed with single muons

in mind, these events will be far more likely to pass it, since their average energy is twice that of single muon events. As a result, it is necessary to somehow protect the reconstructions which assume a single track from these inputs if precise reconstructed outputs are desired.

One useful approach is to attempt to identify individual particle tracks within a coincident event, separate them, and proceed treating each independently. Within IceCube, this process is termed ‘Event Splitting’. Although multiple algorithms have been developed, one of the oldest (which forms the basis of several newer variations) is known as ‘TopologicalTrigger’ [72]. It is mostly heuristic, without strong theoretical justification, but has proven effective in practice. The basis is the observation that the light deposited by a moving muon is distributed by an emitter moving at the speed of light, so each detected photon should have an approximately light-like spacetime separation from (many of) the other photons produced by the same muon. The heuristic aspect is that this is not a strict requirement; spacetime points through which the muon passes will have this exact separation, as will points along the same Čerenkov ray radiating from the muon path, but points on different Čerenkov rays do not necessarily have any particular relationship. Nonetheless, photon arrivals originating from unrelated muons tend to be so far off from this condition that applying it with a generous margin of error tends to behave as desired. The algorithm then associates observed pulses (representing photons) into clusters within which each member is approximately at a light-like separation from some minimum number of others. All clusters which share members (overlap) are then taken to belong to the same ‘subevent’ consisting of the union of their pulses. Ideally, the clusters associated with distinct muons will never share any pulses, and so all pulses generated by each muon will end up in separate subevents. Noise pulses can destroy the effectiveness of the algorithm by bridging from one subevent to another, so it is typically used in conjunction with another algorithm to reduce noise.

While such splitting algorithms can be quite helpful, they still do not successfully treat all cases. It is then sometimes necessary to fall back on an older, more brute-force approach to event splitting which has historically been called ‘split fits’. This technique depends largely on applying a small number of arbitrary splitting algorithms, reconstructing all of the resulting subevents, and attempting to assess whether any of the resulting partitions has succeeded in improving the description of the data. The two standard splitting algorithms used in this approach are the ‘Time Split’ and ‘Geometry Split’. The Time Split is based on the observation that, in many cases, one muon passes through the detector and leaves it entirely before the second enters. There is then some time which separates the pulses associated with each track, and since on average the tracks have the same brightness a simple way to guess this time is to simply choose the

mean time of all pulses (rather than performing some sort of analysis for a gap in the time distribution). The Geometry Split is likewise arrived at by observing that the two muons are usually physically separated by some significant distance, and one arrives earlier than the other, so a single track reconstruction will attempt to fit both by passing approximately through the pulse cloud of one actual track, then through the other, and so some plane perpendicular to the single reconstruction will probably separate the two clusters of pulses. Again, the two sets of pulses will, on average, be of equal size, so simply using the center of gravity (average position) of all pulses will frequently be successful. Some noticeable features of these techniques are that, like the more sophisticated algorithms, they make little effort to be correct in all cases in favor of quickly getting right answers in common cases; and that they are only defined for the specific case of two muons. This inflexibility has led to their falling out of favor, but they have (somewhat) different strengths and weaknesses from newer algorithms, so they can still be used profitably.

3.4.4 Muon Energy Reconstruction

While reconstruction of the geometric properties of muons can be done mostly using timing information, the reconstruction of muon energies must use light amplitude information. As shown previously, at high energies (Figure 3.2), muons emit light (largely through the production of secondary particles) proportionally to their energy. This means that once the rate of light emission has been computed the calculation of the energy can be seen as relatively trivial (some caveats will appear later, however). Determining the light emission rate of a given muon is equivalent to the problem of relating apparent magnitude of an astronomical object to its absolute magnitude; the required information is just the distance between the point of emission and the point of observation. This problem, however, is already solved if a geometric reconstruction of the muon's path has already been performed, since this is a determination of the emission point(s), and the locations of the observation points (the detector modules) are already known to acceptable precision. It then remains to write down the relationship between apparent and absolute brightness, since the emitters in question are not observed in vacuum but in a non-uniformly absorptive medium. The typical form of the reduction of received light with distance from the source is shown in Figure 3.8, and can be understood to effectively be comprised of a term inversely proportional to distance (purely from the geometry of a one dimensional emitter in three dimensional space) and a falling exponential term (due to the absorption). This becomes more complex when the positional variation of the absorption is considered, and in general this can only be treated completely via Monte Carlo simulation of the diffusion process. In practice, either analytical

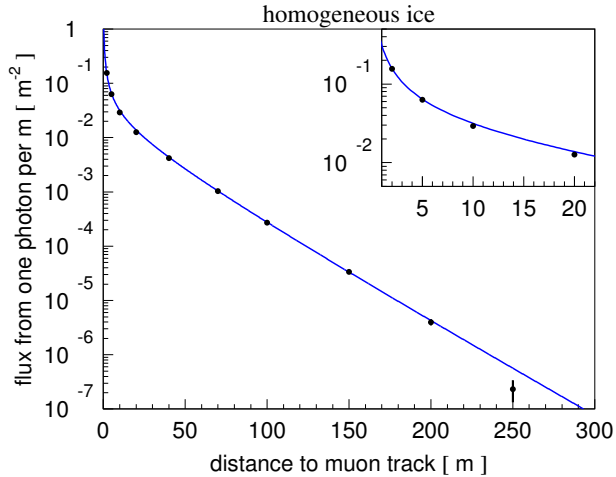


Figure 3.8: The expected distribution of received light amplitudes as a function of distance from a muon track of unit brightness. Points are the results of simulations, while the continuous line is a fit. (Figure from D. Chirkin)

approximations or tabulations of simulation output can be used. A final detail, which cannot be ignored if high precision is sought, is that while the proportionality of light output to energy holds on average, it does not hold exactly at all times. Specifically, the forms of energy loss whose influence increase proportionally with muon energy are also those which are composed of discrete events which occur stochastically, many segments of the path of a high energy muon will be relatively dim, with only Čerenkov light being emitted by the muon itself along with a near-constant baseline of ionization loss, while there will be short periods of intense emission, for example when a bremsstrahlung emission event occurs. Passable results from energy reconstruction can be achieved by glossing over the variation in brightness by simply averaging the energy as reconstructed from each detector module along the length of an observed muon track, but in principle, treating this more carefully should be advantageous. Possible approaches include dividing the track into segments to be reconstructed separately, or attempting a joint unfolding of all segment energies from all observations [66]. Finally, as in geometric reconstruction, noise photons unrelated to the muon should be somehow treated, as they can appear at arbitrary distances from the muon, and will cause energies to be systematically overestimated when the algorithm attempts to explain their presence.

For this work, the MuEx energy reconstruction algorithm [73] is used. It uses an analytical approximation for observed light distributions which includes the layered structure of the ice, and it includes a widening of the PDFs used for energy-loss expectations as a function of muon energy which mitigates the effects of

stochastic energy losses. The treatment of noise by this algorithm is minimal, so it performs best when a noise removing algorithm has been used to filter the pulses.

3.5 Production of Simulated Data

As will be discussed in Chapter 5, the analysis presented in this work requires accurate simulation of the response of the IceCube detector to fluxes of neutrinos. Such simulation is produced using a chain of simulation tools developed by a number of authors over the lifetime of the project which encompass knowledge of the physics processes, detector hardware, and calibration to produce Monte Carlo-style simulations of possible neutrino (and cosmic ray air-shower) events observed by the detector.

3.5.1 Neutrino Interaction Generation

In order to simulate neutrino interactions, a generation program, NeutrinoGenerator [74], similar to ANIS [75] is used. ‘Primary’ neutrinos (which may later be treated as either astrophysical or atmospheric in origin) are sampled on the surface of the Earth from an arbitrary (but known) energy spectrum, and directed towards a volume surrounding the IceCube detector. While the Earth is essentially transparent to neutrinos at low energies (below 1-10 TeV), at higher energies the rising cross-section means that neutrinos may lose energy by neutral current scattering or be transformed into charged leptons far from the detector. In most of the latter cases, the neutrino can be considered entirely lost, as electrons and muons tend to lose their entire energies via electromagnetic interactions, but taus are likely to decay while they still carry a substantial fraction of their original energies, effectively regenerating tau neutrinos at lower energies. To account for these effects, the individual neutrinos are tracked through the volume of the Earth and selected to undergo interactions randomly according to cross-sections and local matter densities (Figure 3.9). For greater efficiency, neutrinos which would be totally lost in the simulation to charged-current interactions are instead assigned a reduced weight (probability of occurrence), rewritten as if a neutral-current interaction had occurred instead, and propagated onward. During the entire neutrino-propagation phase, neutrino oscillation effects are ignored, under the assumption that they are negligible for the neutrino energies at issue, which is true for this work as long as unproven exotic effects such as sterile neutrino oscillations are not allowed.

While the simulation could allow the simulated neutrinos to interact at their natural rates in and around the detector, as it (approximately) does in other regions, this would be extremely inefficient in some energy ranges since most neutrinos would not interact at all. For this reason, each simulated event is forced to

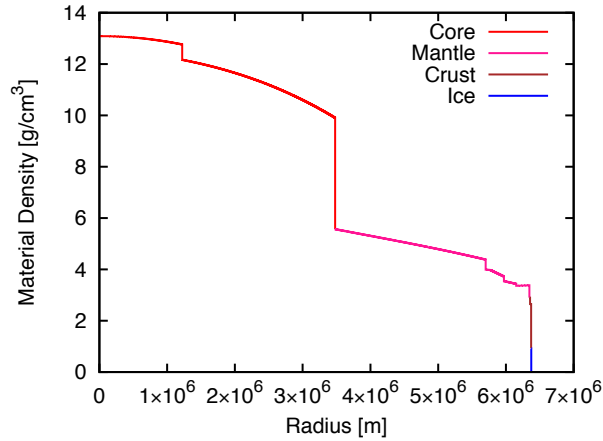


Figure 3.9: The version of the Preliminary Reference Earth Model [76] used when generating neutrino interaction events for IceCube simulation. The model has been augmented by including the ice sheet where IceCube is simulated. This figure strictly represents only a slice through the Earth from the core to the South Pole, as the ice sheet is tapered off (according to a simplistic parameterization) until it vanishes at latitude 70°S; the rock is treated as being spherically symmetric.

interact in the region of the detector, and a weight penalty is recorded depending on how unlikely such an interaction was. For neutrinos interacting via the neutral-current process or electron neutrinos interacting via charged-current, the daughter particles do not travel long distances, so it is sufficient that the region of forced interactions be only slightly larger than the instrumented volume of the detector. For charged-current muon neutrino interactions which produce high energy muons with ranges of kilometers, it is necessary to expand the potential interaction region along the direction of the particle’s travel to be long enough that all possible muon propagation distances can be sampled. The same must be done for tau neutrino simulations, since the decay of a tau has a probability of $\sim 17\%$ to decay to a muon (and muon neutrino) with a substantial fraction of its energy. This unfortunately reduces the efficiency of the simulation and poses a considerable barrier to importance sampling.

3.5.2 Air-shower Simulation with CORSIKA

In addition to simulating neutrino events which can be used to fit possible signals, it is also necessary to simulate cosmic ray air-showers in order to learn how to remove them as a source of background. The generation of these events and their simulation until they reach the ice sheet (after which all remaining particles

other than muons and neutrinos can usually be neglected), is performed by the CORSIKA software package [77], which is an efficient and widely-tested implementation of air-shower physics. Some of the physics effects necessary in these calculations are quite well understood (decays of light mesons, electromagnetic interactions of gamma-rays and electrons), but some, primarily the interactions and production of hadrons at high energies, have never been measured and cannot be exactly calculated. For these components of the calculation it is necessary to rely on extrapolations from existing measurements using theory or phenomenological models. For the simulations used by this work the Sibyll 2.1 [78] hadronic model has been used. It is extremely difficult to judge whether any well-made hadronic model is strictly superior to another even when they differ, and while Sibyll is not assumed to be particularly more correct than other models of the same generation (a number of newer models now exist, and Sibyll itself will soon be updated based on measurements made since the release of the previous version). It is chosen largely for the pragmatic reason that it can be calculated far more rapidly than most other models, allowing a greater volume of simulated data to be produced. Sibyll 2.1 does not include charm quark physics, so it generally cannot produce the prompt decay air-shower component, however this is expected to be of greater importance for neutrinos than air-shower muons. So, both for reasons of simplicity (as well as bugs in the available versions of CORSIKA in handling air-shower neutrinos) atmospheric neutrinos for this study are treated using direct neutrino simulation (the same as astrophysical signal simulations), rather than using the output of CORSIKA.

3.5.3 Muon Propagation and Cascade Generation

For both charged-current muon neutrino events and cosmic ray air-showers, produced muons must be propagated from the points where they enter or are produced in dense media (ice, rock) to the detector. Similar to neutrino propagation, this is a largely-straightforward application of the interaction (energy loss) cross-sections discussed in Section 3.1. Interaction with energy losses smaller than some threshold are treated as continuous, while larger losses are recorded as discrete events so that their individual Čerenkov light can be simulated. Two software implementations are used for this work [57, 79].

3.5.4 Photon Propagation

Because in a literal sense IceCube detects only photons, the production and propagation of Čerenkov light from the particles passing through the detector must be simulated. In the past this was performed using lookups in precalculated tables of photon arrival times as functions of relative source and receiver positions,

but this method has become disfavored as it has proven difficult to capture important, fine detail of the light propagation process in histogram tables which are small enough to be efficiently usable on current computers [80, 81] (a new method [82] for constructing correct but compact representations of high detail tables became available somewhat too late for use in this work, and is not fully general). Since the tables were produced by generating large numbers of photons, tracking their random walks of scattering and absorption, then tabulating the results, a viable technique for simulating events is to omit the tabulation step and instead propagate photons on demand for each event. This approach is sometimes referred to as ‘direct propagation’ since no intermediate tables are used, and for low energy events (which produce few photons) it can become more efficient than the tabulation method. Direct propagation for simulation purposes is implemented by two software packages [83, 84], both of which were used to produce the data used in this analysis. In order to process photons rapidly enough for practical generation of the necessary quantities of simulated data, both implementations depend on propagating numerous photons in parallel using graphics acceleration hardware (GPUs). Photons are generated randomly according to parameterizations of the Čerenkov output of different particles (muons, taus) or groups of particles (whole electromagnetic or hadronic showers [85]). Each photon is then tracked through a series of random scatterings according to the ice model 3.2.1 until it is either absorbed or strikes the surface of a detector module. As an approximation the modules are typically treated as being larger than their physical size for the purposes of this collision detection. This technique, known as ‘DOM oversizing’ is useful because it allows fewer photons to be propagated (a factor of 25 for a typical oversizing of a factor of 5 in module radius), and can be justified because, even when oversized, the simulated module is still small compared to the structures in the ice itself, which are treated as being 10 m at the smallest. An important detail, however, is that this scaling is applied only in directions perpendicular to the path of the photon under consideration, so that the path length is not changed; otherwise a detectable shift in timing results. Once a photon has been determined to reach a detector module a correction is applied for the wavelength and position of incidence to determine whether a photoelectron is ejected from the photomultiplier cathode.

3.5.5 Hardware Simulation

After photoelectrons are generated at the detector modules, the remaining hardware response of those modules must be simulated. The first component of this is the behavior of the photomultiplier tube itself: The quantum efficiency of the photocathode is actually treated as the final step of the photon propagation




Figure 3.10: The probability density function from which amplified photomultiplier pulse charges are sampled, as a function of charge relative to an ideal, single photoelectron pulse.

software, but the amplification process must be handled. Due to fluctuations in the electron cascade within the photomultiplier, the output charge at the anode must be sampled from a distribution of charges, based on a determination from hardware measurements. Although each photomultiplier may have its own unique distribution of charges, a single distribution (Figure 3.10) is simulated for all modules [86, 87].

However, not all photoelectrons undergo the amplification process as intended [70]. Some fraction of the photons bypass the photocathode and directly strike the first dynode, reducing their charge amplification by a factor but decreasing their time of flight. These produce what are termed ‘prepulses’, as the resulting pulse has abnormally small amplitude (simulated as $1/15$ of the charge of a standard pulse) and occurs early (~ 30 ns), since the tube is traversed by the photon at the speed of light, rather than by an electron which must accelerate from zero velocity. This outcome is simulated as accounting for 0.7% of all photoelectrons. Another possibility are ‘late’ pulses, in which the photoelectron scatters elastically from the first dynode, and requires some time to reverse direction again and cause the standard amplification cascade. The generated charge pulse has the standard amplitude, but is delayed in time by anywhere from about 25-160 ns. This outcome is simulated for 3.5% of pulses. Finally, a photoelectron may scatter off of a residual gas atom inside the photomultiplier tube, ionizing it in the process. This does not substantially reduce the velocity or energy of the original photoelectron, so it proceeds to be amplified normally, but the now-positively charged ion drifts (slowly compared to the electrons due to its far greater mass) back toward the photocathode. When it strikes the photocathode several (around ten) electrons are typically ejected, which then accelerate to the dynodes causing a larger than usual charge pulse, referred to as an ‘afterpulse’ to be output at a time




Figure 3.11: The simulated time distributions of different photomultiplier output pulse types, relative to the nominal time for a signal to traverse the photomultiplier tube. The afterpulse component is shown separately due to its much longer time scale.

substantially later than the original single photoelectron pulse (ranging from ~ 250 ns to $\sim 10\mu$ s). Each photoelectron is simulated as having a 5.9% chance to also create an afterpulse, and each afterpulse then has the same chance to create further afterpulses (although this geometric progression dies off quickly) [87].

In addition to systematic variations in the timing and amplitude of the output pulses, there is also a random contribution to the timing due to nonuniformities in the photomultiplier itself. These include the fact that its surface is not exactly a sphere centered on the first dynode, and that the electric field within the tube is likely not entirely uniform. Based on lab measurements, this can be approximated using an asymmetrical jitter in time, for which current simulation uses a Gumbel distribution with mode 0.15 ns and variance 6 ns^2 [87].

Finally, while the photomultipliers used by IceCube have excellent linearity for low-brightness inputs (up to ~ 50 photons arriving simultaneously, and more if their arrivals are distributed over longer times), at high brightnesses their output lags behind the input until it saturates entirely at some maximum current. This effect is simulated by estimating the approximate ideal current at the photomultiplier anode and then reducing the amplitudes of the produced pulses by the ratio of the expected actual current to the ideal current using a parameterization by T. Feusels of lab measurements [88] (Figure 3.12). This effect has been found to operate on very short time scales (of order 2 ns), so even when some portions of the photomultiplier’s output are highly saturated, both earlier and later remains faithfully linear in the input.

Once the output of the photomultiplier has been generated, it must be digitized in the same manner that the actual detector modules employ. This consists of three stages: determining whether the module’s




Figure 3.12: The simulated treatment of photomultiplier saturation for the IceCube nominal photomultiplier gain of 10^7 . At this gain the peak ideal output current for a single photoelectron is about 0.36 mA.

discriminator is triggered, leading to a readout being generated, the simulation of the intermodule communication of such triggers (‘local coincidence’) which determines whether the readout is full or abbreviated, and finally the recording of the readout as if by 1-4 of the digitizers [89]. The simulation of the discriminator behavior is accomplished by convolving the stream of pulse times and amplitudes produced by the photomultiplier simulation with a template function based on measurements of the pulses reaching the discriminator in the real electronics, and likewise the digitized output uses pulse templates corresponding to the different amplifier stages preceding each digitizer [90, 91]. The local coincidence and digitization are performed taking into account the known latencies and deadtimes of the hardware. In addition, the voltage ‘droop’ caused by the inductive coupling of the photomultiplier output to the readout electronics is accounted for in the pulse templates and is treated over the full time span of each simulated event, as it not only alters the digitized waveforms, but can even suppress triggering and readout on small pulses which follow during the recover period after pulses with very high charge.

3.5.6 Triggering Simulation

Digitization of photomultiplier outputs is the last stage of IceCube data handling which deeply involves hardware, however data at this level is generally not stored, for simple reasons of volume, so to make simulated data fully equivalent to the output of the real detector the trigger criteria must be applied as well as the first level reconstruction and filtering. The latter steps can be performed with the exact same

software, but for technical and historical reasons the triggering is simulated using distinct software from that actually used for real observation [92]. (Reasons include that the actual data acquisition must operate on a continuous stream of input, while the simulation has already had its input reduced to separate time blocks, and simply that the actual data acquisition is implemented in a different computer language which is inconvenient to interface with the simulation and data processing software.) There are no known differences in the output of the two trigger implementations, but full and rigorous consistency checks have not yet been undertaken.

Chapter 4

Data Selection

As the purpose of this study is to identify astrophysical muon neutrinos, they must be distinguished both from other types of events in the IceCube detector and from other fluxes of neutrinos. The majority of the data recorded by IceCube is produced by cosmic ray air-showers, specifically the muons produced in such air-showers, as they are able to penetrate through the ice sheet to reach IceCube. Since this analysis seeks to take advantage of the long muon tracks and cannot depend on observing the neutrino interaction vertex inside the detector, some other means must be employed to ensure that the observed muons are likely produced by neutrinos. This can be accomplished expediently by selecting only muons whose directions imply that they have passed through substantially more material than the probable maximum range of the muon, indicating that some of the distance must have been traversed as a particle less prone to interaction, i.e. a neutrino. This approach cannot distinguish neutrinos by origin (atmospheric or astrophysical); that analysis is left to Chapter 5.

This chapter describes the particular details of the selection choices that were made to create a sample of neutrinos with minimal contamination from air-shower background. In practice, this means comparing the data rates predicted by simulations of signal and background fluxes, and defining a set of criteria which preserve the signal while removing the background (termed ‘cuts’ because each criterion removes from consideration the events which fail to meet it, thus reducing the dataset). For this sub-study, 805.5 hours of experimental data from the 2010 data-taking period are compared to a cosmic-ray simulation dataset equivalent to approximately 264 hours of air-shower background with a spectrum as given by [93] and a neutrino simulation dataset weighted (as described in Section 5.2.1) to both a conventional atmospheric neutrino spectrum and a hypothetical E^{-2} power law spectrum (representing a possible astrophysical flux). If the data selection is successful, the simulated air-shower events should be eliminated, the remaining

experimental data should generally match the atmospheric neutrino simulation, and as many as possible of the events from the test astrophysical flux should be retained.

4.1 Online Filter

The data available as input for this analysis are those selected by the ‘online’ filters (Section 3.3). The ‘Muon Filter’ is specifically designed to select events which are likely to be muons passing through the detector in any direction, although with a bias for events which travel down through the ice to select only those with high energies in order to avoid wasting bandwidth on the many low energy air-shower muons. It should be noted that the Muon Filter was designed by other researchers, and rather than as a part of this work.

4.1.1 2010 Muon Filter

During the 2010 data period the Muon Filter [94] was defined in terms of two angular regions: an ‘up-going’ region of events reconstructed with large zenith (which, if correct, means that they have traversed a substantial distance through ice or rock) and a ‘down-going’ region of events reconstructed to arrive at the detector more vertically. In order to control the overall data rate requiring more advanced reconstruction, all events in which fewer than 8 detector modules launched with local coincidence are cut (the number of launched modules is frequently called ‘NChannel’, although the requirements for launches to be considered vary by context), as are all events with fewer than 10 launched modules ($N_{\text{Channel}} < 10$) and zenith angles reconstructed by LineFit which are less than 70° . The events passing these cuts are then reconstructed using the slightly more detailed SPEFit.

Since coincident air-shower events may be mis-reconstructed as up-going (see Section 3.4.3), it is useful to test whether the events reconstructed as up-going are actually successful reconstructions or not. SPEFit is designed around a likelihood maximization, and so, if the final likelihood it reports is larger, the event is more likely well described by the fit. For computational reasons, negative logarithms of likelihoods are reported, so a cut is applied that $\frac{-\text{Log}L_{\text{SPE}}}{N_{\text{Channel}}-2} \leq 8.1$ for events reconstructed with zenith angles greater than 78.5° . (The denominator is chosen to make the cut behave more consistently for events detected by substantially different numbers of detector modules.)

Many of the events with down-going reconstructions are in fact down-going, but there are so many air-showers with low energies that it is not practical to transmit the data for all of them, well reconstructed or




Figure 4.1: Illustration of the cut used by the Muon Filter for down-going events during the 2010 data taking period. The colored histogram shows the expected rate of atmospheric muons from cosmic ray air-showers as a function of direction (cosine of the zenith angle from an SPEFit reconstruction 3.4.2) and total recorded charge in the detector (Q_{Tot}). The white line represents the cut which was used, where data below the line is eliminated, while the black lines show a similar cut which was considered but not used. Note that all data the cosine of whose zenith angle is < 0.2 is left untouched by his criterion.

not. A cut is therefore placed on events with zenith angles reconstructed by the SPEFit as more down-going than 78.5° to remove those which are lower energy and therefore less likely to be interesting air-showers or neutrinos. This cut is defined as a reconstructed zenith angle-dependent criterion on the total number of photon-electrons recorded in the event (referred to as ‘total charge’ or Q_{tot}), as a slightly better proxy for event energy than NChannel, and is shown in Figure 4.1.

The estimates for the results of this filter are that in the up-going region conventional atmospheric neutrino events are expected to pass at a rate of 1.04×10^{-2} Hz, while for an E^{-2} flux 92.7% of events are expected to pass in the up-going region (used for this analysis), while cosmic-ray air-shower events are expected to pass the filter with a rate of 24.2 Hz. This means that for every neutrino event available in the data, there are approximately 2300 air-shower events remaining which must be rejected by the data selection for this analysis to obtain a pure sample of neutrinos. The total data output rate of this filter is expected to be 33.3 Hz of events over the whole sky.

4.1.2 2011 Muon Filter

For the 2011 data-taking period the Muon Filter underwent a minor redesign for the up-going region [95]. The reconstruction quality cut was rewritten to be $\frac{-\text{Log}L_{\text{SPE}}}{N_{\text{Chan}}-3} \leq 8.7$, which has the effect of both increasing efficiency for retaining low-energy neutrino events while simultaneously rejecting a greater proportion of mis-reconstructed air-shower events. The cut in down-going region was also altered so that for a given reconstructed zenith angle, an event is required to have 26% more deposited charge to be retained, but this change has no impact on this analysis.

The result is that the expected rate of atmospheric neutrinos is increased to 1.29×10^{-2} Hz, 93.7% of events from E^{-2} flux are expected to pass, and 19.9 Hz of air-shower events are expected to remain. Since this provides a more pure input to the analysis selection there is no particular need for the selection to be altered from one which is applicable to the 2010 data period, so the selection was optimized with respect to 2010 data only and applied unchanged to 2011 data. Neglecting the potential for bias as a function of energy in the higher-level selection, approximately 24% more atmospheric neutrino events can be expected in the 2011 data than in 2010, but since the effect is confined to low energies it has little significance for the behavior of the actual analysis.

4.2 Data Volume Reduction

After data has been retained by the Muon Filter and transmitted to the northern hemisphere for permanent storage, more reconstructions are performed (since computing resources are less constrained away from the experiment site). In particular, the MPEFit directional reconstruction algorithm is applied to the data, in addition to the SPEFit calculated during the filter decision process. In many cases this more-precise reconstruction reveals that an event would not have passed the online filter had this reconstruction been available at that stage, so this suggests that a good way to quickly remove uninteresting events is to reapply cuts similar to those of the filter. In addition to replacing the filter's SPEFit with MPEFit, the total collected charge in the event (used as a proxy for event energy) is replaced. The IceCube detector is not entirely homogeneous, particularly because of the presence of the DeepCore infill which gives both higher instrumentation density and has more efficient collection of light per unit area of photocathode. As a result, events whose particles have the same true energy may appear much brighter depending on whether they pass through the infill volume. For sophisticated reconstructions, this additional information should be purely beneficial, but total charge is overly simplistic as a proxy for energy for this to be true. A slightly better energy proxy is

then the total charge observed in detector modules which are part of the larger grid spacing only. A cut was then applied to the data such that all events whose zenith angles, as reconstructed by the MPEFit, were smaller than 78.5° were required to meet the criterion $\log_{10}(Q_{\text{tot}}) \geq 0.6 * (\cos(\text{zenith}_{\text{MPE}}) - 0.5) + 2.5$ to be retained for further consideration.

A reduction in the rate of data with up-going reconstructions was accomplished by selecting based on a simple, geometrical quantity which functions as an indicator of correct reconstruction of track-like events. Since the Čerenkov light both spreads out and is attenuated at larger distances from the emitting particle track, detector modules near the particle should generally collect more charge. This means that the average distance away from the particle track of a detected photoelectron should be small, and this property can be checked for reconstructed tracks. In cases where the reconstruction correctly passes through the cloud of observations the mean distance of modules which launch, weighted by charge observed in the module (the ‘charge-weighted distance’), will be smaller than if the reconstruction is incorrect and does not pass near some or many of the observations. Events with up-going reconstructions whose charge-weighted distance was greater than 200 meters were therefore rejected, unless they were reasonably bright (more than 100 photoelectrons observed), as they were likely to be mis-reconstructed, and not to contain an interesting, high energy neutrino-induced track which might be separated in the next phase.

Testing this data-reduction phase on the 805 hours of experiment data showed that the data was reduced from a rate of 32.6 Hz (the actual output of the Muon Filter over the entire sky) to 10.6 Hz, or a passing efficiency of 0.325. This is similar to the result found when applying the same criteria to cosmic ray background data simulated by CORSIKA, for which the efficiency was 0.362. Testing on a hypothetical E^{-2} spectrum of muon neutrinos gave an efficiency of 0.921 averaged over all zenith angles, and 0.962 for events whose true direction was up-going. Combined with the efficiency of the muon filter, this means that 89.2% of neutrinos from a signal flux would survive at this stage of processing for the 2010 data period, and 90.1% for the 2011 data-taking period.

4.3 High Level Reconstruction

With approximately 10 Hz of data remaining, it is reasonably practical to attempt further reconstruction which may be able to separate coincident events into their constituent components which can then be better separated into promising candidate neutrinos and background cosmic rays. The core of this technique was the use of the TopologicalTrigger algorithm, but this algorithm performs poorly when confronted with large

amounts of random noise unconnected to particle tracks. Therefore, a ‘hit cleaning’ algorithm known as ‘SeededRT Cleaning’ [96] was first applied to the data. This algorithm is based on two assumptions: first, the local coincidence criterion on detector module readouts strongly suppresses random noise, so any module launches which have coincidence with neighbors are unlikely to be noise; second, (like the assumption central to TopologicalTrigger) that light signals propagating outward from the same light source will generally have approximately light-like spacetime separation. SeededRT Cleaning therefore starts with the collection of all reconstructed pulses which were extracted from readouts with local coincidence, and iteratively adds to the set all other pulses which are ‘close’ to at least one pulse already in the set, by being less than 150 meters away and occurring within one microsecond of the expected arrival time of a signal traveling at the speed of light from the position of the included pulse to the position of the pulse under consideration.

After cleaning noise pulses, the events were processed using an implementation of the TopologicalTrigger algorithm [97]. This algorithm has three main, user-adjustable parameters, which are the maximum horizontal distance within which two pulses are considered causally connected (measured in meters), an equivalent maximum vertical distance (measured in detector module spacings to allow for the different vertical densities of the normal IceCube and DeepCore strings), and the error tolerated relative to ideal propagation from one point to a possibly causally connected point (measured in nanoseconds). A small optimization study was performed over this parameter space using simulated coincident events, the results of which are shown in Table 4.1. The criterion for an event to be considered ‘correctly split’ is difficult to define rigorously, since any given particle which was actually simulated in the event may produce little or no detected light. For this study, a particle was considered non-trivial if it produced at least 4 photons which reached detector modules which then launched and read out data. The result of the study was that for the settings considered (chosen to be similar to those arrived at by other similar studies for other selections) most choices have similar success rates, so the particular choice is not critically important (combinations which performed substantially worse have been suppressed from Table 4.1, however). The 300 meter, 20 spacing, 800 ns combination was chosen, as its performance was second best, but was not quite as close the edge of the parameter space which had been used previously as the (almost identical) highest ranked combination.

A detail which proved to be important was that for events which produce large amounts of photomultiplier afterpulsing the TopologicalTrigger algorithm tends to group some or all of the afterpulses together, placing them in their own subevent in the output. There is nothing wrong with this behavior (in fact, it may be viewed as desirable since no current IceCube reconstruction algorithms are aware of afterpulses, and so

they should perform better with afterpulses removed), but it complicates counting whether the number of subevents output by the splitting algorithm matches the number of significant particles actually present in the event. It was found to be fairly simple to devise a conservative heuristic algorithm to identify subevents consisting of afterpulses as follows: Since afterpulses are triggered by actual photoelectrons, the set of detector modules contributing to an afterpulse subevent should be approximately a subset of the detector modules comprising some earlier subevent. At least 75% of the pulses in the afterpulse subevent being present in the earlier subevent seems to be a practical threshold. Next, the afterpulse subevent should have its pulses at a substantially later time than the subevent which produced it. A simple condition which seems to work well is to require that the mean time of the potential afterpulses be at least 3 microseconds later than the mean time of the pulses in the potential progenitor subevent. Finally, while each afterpulse tends to have an output pulse charge higher than that of a single photoelectron, afterpulses are relatively rare. Empirically, the afterpulses typically contain 5-10% as much charge as the collection of progenitor pulses. Applying this set of criteria appears to be conservative in the sense that it essentially never flags a subevent as afterpulses when it is not, and at the same time the rate of failing to flag subevents which appear to a human observer to be afterpulses is less than $\sim 1\%$. In addition to being useful for excluding afterpulse subevents from the counting in this splitting study this identification algorithm was used as a cut in the main analysis to avoid running expensive reconstruction algorithms on collections of afterpulses.

After the splitting algorithm, each resulting subevent was reconstructed again (unless there is only one subevent reported and it was identical to the original entire event). As usual for analysis of track-like events, directional reconstructions were run in sequence using the previous results as seeds, beginning with LineFit, SPEFit iterated several times (5 in total), and finally MPEFit. In addition, other reconstructions and parameters which were later used as cut variables were computed. These are discussed in greater detail below, but included a profile likelihood estimate of the reconstruction error of the final MPEFit, a directional reconstruction with a Bayesian prior that events should be distributed like background cosmic rays in zenith angle, and ‘split fits’ which try to treat the event as a pair of muons, in case the TopologicalTrigger algorithm has failed.

Horizontal distance (m)	Vertical distance (module spacings)	Time Error (ns)	Fraction of events correctly split
300	15	800	0.932
300	20	800	0.932
300	30	800	0.932
300	25	800	0.932
300	20	500	0.932
300	15	500	0.932
300	30	500	0.932
300	25	500	0.932
300	15	1000	0.931
300	20	1000	0.930
300	30	1000	0.930
300	25	1000	0.930
300	15	1200	0.929
300	20	1200	0.928
300	30	1200	0.928
300	25	1200	0.928
400	20	500	0.926
400	25	500	0.926
400	15	500	0.926
400	30	500	0.926
400	15	800	0.924
400	20	800	0.924
400	30	800	0.924
400	25	800	0.924
400	15	1000	0.921
400	20	1000	0.920
400	30	1000	0.920
400	25	1000	0.920
400	15	1200	0.918
400	20	1200	0.918
400	30	1200	0.918
400	25	1200	0.918
500	20	500	0.916
500	15	500	0.916
500	25	500	0.916
500	30	500	0.916
500	15	800	0.914
500	20	800	0.914
500	30	800	0.914
500	25	800	0.914
500	15	1000	0.911
500	20	1000	0.910
500	30	1000	0.910
500	25	1000	0.910

Table 4.1: Fraction of coincident track events which were correctly separated by the TopologicalTrigger algorithm for various choices of its parameters.

4.4 Final Selection Criteria

The final data selection criteria used in this analysis are presented here in a particular order which is intended to be reasonably intuitive, but it should be noted that all criteria were (approximately) optimized simultaneously. Two types of plots are shown to illustrate the reasoning behind these criteria: Distributions of simulated and real data after the applications of the cuts which have already been described at a given point are used to illustrate the logic of choosing the given variable(s) to separate signal data from background. Additionally, cumulative plots are shown of the rate of simulated and experimentally observed events retained in the selection as a function of single cut parameters with all other cut parameters held fixed at their final values to illustrate why the selected value for that parameter was deemed suitable.

One known weakness of the procedure used to derive these selection criteria is that there was no single, quantitative metric which was optimized. In principle, the correct method would be to construct the full analysis method of Chapter 5, and to maximize the analysis sensitivity (Section 5.3) with respect to the selection parameters. This was not possible for a combination of reasons: Firstly, these portions of the analysis were prepared in the order shown in this document, which is arguably the wrong one. Secondly, estimating the analysis sensitivity is a rather computationally-intensive activity and it is not clear that it would have been feasible to do this iteratively for even the final stages of the selection optimization. Finally, there is the theoretical difficulty that the analysis method developed in Chapter 5 is predicated on the assumption that the contamination of the data by cosmic ray air showers is quite small, so the sensitivity could not be meaningfully estimated until the data selection was sufficiently mature as to give nearly the final purity, and attempting to use it as a final ‘polishing’ step would have to be done cautiously to avoid leaving the portion of the phase space in which the analysis assumptions hold (i.e. cosmic ray contamination remains low and never dominates in any part of the observable space).

Instead, the selection was subjectively optimized by the human experimenter according to a set of simple guidelines: the analysis assumptions must be met, so air shower contamination must be minimized, and only high energy neutrinos can be expected to have a high probability of being astrophysical in origin, so events with higher observable energies should be preserved rather than the larger number of expected neutrinos with low energies.

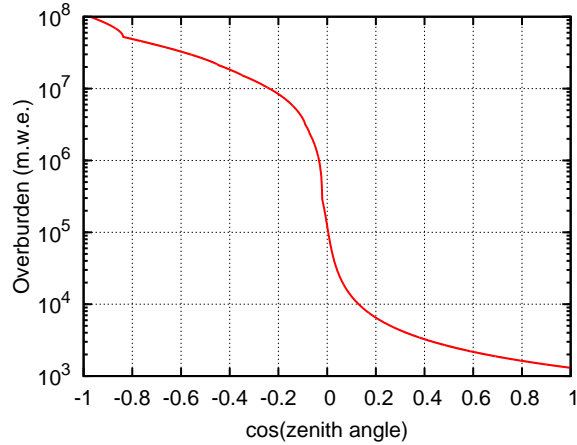


Figure 4.2: The overburden in meters of water equivalent above the top of the IceCube detector (1450 meters below the ice surface) as a function of zenith angle. The change in slope just below the horizon is due to the transition from ice to rock, and the similar change at $\cos(\text{zenith angle})$ somewhat smaller than -0.8 is due to the Earth's core.

4.4.1 Zenith Angle Restriction

The main power of this selection to reject air shower events is expected to derive from selecting events which pass through considerable overburden to reach the detector, but at some zenith angle the overburden is no longer sufficient and air shower events will reach the detector at a non-negligible rate. Figure 4.2 shows the overburden as a function of zenith angle, which begins to rise rapidly for $\cos(\text{zenith angle})$ smaller than 0.1 , where it is equivalent to more than 12 kilometers of water. A cut is therefore placed on the cosine of the reconstructed track zenith angle, rejecting all events which are too down-going, as shown in Figure 4.3.

4.4.2 Track Reconstruction Quality Selection

In most of the angular (zenith angle) space used by this analysis, the material of the Earth blocks cosmic ray muons entirely from reaching the detector. Such events are present in this part of the observable space in the reconstructed data only because the reconstructions are wrong, and since neutrino-induced events whose directions are severely misreconstructed are similarly undesirable it is natural to attempt to detect and remove poorly reconstructed events from the data sample to be analyzed. In earlier stages (Section 4.2) simple, largely geometrical variables (the charge-weighted distance, in particular) were used for this purpose, driven by a need to process large data volumes quickly, but for higher selection precision it is necessary to turn to more advanced techniques. Two methods are used here, both of which derive from the fact that the




Figure 4.3: Top: The angular distribution of all events before the zenith angle cut. The sudden change at $\cos(\text{zenith angle})$ of 0.2 is caused by the data reduction cuts from Section 4.2. Bottom: Due to other cuts the rate for all types of events is essentially constant for all choices of the $\cos(\text{zenith angle})$ cut above 0.1.

final MPEFit used to reconstruct these data is a maximum likelihood calculation, and additional information can be extracted from the the likelihood description besides the parameter values of the best fit. The first variable used, known as the ‘paraboloid sigma’ uses the shape of the likelihood space around the maximum point to estimate the statistical uncertainty on the location of that point; the second, the ‘reduced log likelihood’ attempts to use the best obtained likelihood value as a global measure of the success of the fit.

The purpose of the paraboloid sigma is to estimate the probable error of the reconstruction in the zenith and azimuth angle track parameters. It does this using a profile likelihood construction, as described in greater detail in Section 5.1.1, in order to treat the likelihood only as a function of these two parameters, essentially rendering the other fit parameters (the three spatial coordinates of a point through which the track passes) unimportant. Interpreting the likelihood in terms of probability suggests that parameter values within some difference from the maximum are allowed within a corresponding confidence level, and so the statistical uncertainty on the best fit is described by the region of the likelihood space where the likelihood is close to the maximum, within a suitably computed tolerance. For a well-behaved fit, it should be possible to Taylor expand the likelihood around the maximum point, and by the definition of a maximum the gradient term will be zero, but the second order term corresponding to curvature will not. The region of the likelihood within a chosen confidence level can then be approximated by using the Taylor expansion to estimate where the likelihood value is within the tolerance of the maximum value, and if the Taylor expansion around the maximum is truncated at second order its functional form will be a paraboloid. So, to estimate the angular error of the directional reconstruction the profile likelihood is scanned (with a small number of points) in the zenith and azimuth angle space and fit with a parabolic function. The elliptical region of the paraboloid which surrounds the best fit at the chosen confidence (usually 68%) is computed, and its parameters (axis lengths and orientation) are reported. For using this information solely as a quality measure for the reconstruction it is simpler to further reduce it to a single variable, discarding the unnecessary orientation parameter, and combining the two axis lengths by interpreting their geometric mean as a single radius, since this gives a circle whose area is the same as that of the ellipse. This final radius is the paraboloid sigma (called sigma because it can also be interpreted as the width of a Gaussian probability density for the true direction of the event).

The paraboloid sigma is in theory a correct description of the statistical angular uncertainty of each event’s directional reconstruction. However, in practice, the likelihood function used by the reconstruction is not actually correct, and this inexactness becomes noticeable in that the actual reconstruction error does not




Figure 4.4: The ad-hoc correction function applied to the paraboloid sigma as a function of the number of detector channels which contribute to the likelihood.

fall within the computed radius the correct fraction of the time. The paraboloid sigma does still correlate with the actual error, so this situation can be patched up substantially by multiplying with an ad-hoc correction function derived by inspection of the ratio of the mean paraboloid sigma to the mean true error, usually as a function of an observable correlated with energy. The correction function applied in this work is shown in Fig. 4.4. This ‘corrected paraboloid sigma’ then has a roughly consistent meaning for events of all energies, making it easier to use for the purpose of data selection.

Unlike the paraboloid sigma, the reduced log likelihood (or ‘RLogL’) is defined less in the sense of comparing the direction of a reconstruction with other possible directions and more with whether the reconstruction of that event has been as successful as the reconstructions of other events. Since the best fit is defined by maximizing the likelihood function, or minimizing the negative logarithm of the likelihood, a fit should generally describe the data better if it has a smaller negative log likelihood. The comparison between events is hampered, however, by the fact that events with different brightness have different likelihood distributions, since by Wilks Theorem (see Section 5.1) the distribution depends on the number of degrees of freedom of the fit, which in this case is the number of observations (light detections by distinct detector modules) less the number of parameters being fit (two angular and three positional, see Section 3.4.2). Ideally, the distribution should be a χ^2 with this number of degrees of freedom, whose mean will then be equal to the number of degrees of freedom, so by dividing this quantity out (as long as the number of degrees of freedom is not too small), all events should belong to a shared distribution with the same mean. Therefore, the reduced

log likelihood is computed as the event's best fit likelihood, divided by the number of detector modules contributing to the fit minus the five fit parameters.

As shown in Figure 4.5 the distribution of simulated air shower background events is shifted to larger (worse) values of both paraboloid sigma and RLogL compared to the simulated distribution of neutrino-induced events from a hard, signal-like spectrum. A cut (drawn over the distributions) has been designed to select for small values of these variables simultaneously, eliminating the vast majority of the background events surviving the previous selection criteria. The cut is defined in terms of two linear rays in the space of the two variables simply as an evolution of the simple case of cutting on each variable separately (in which case the two rays would be perpendicular and aligned with the variable axes).

4.4.3 Bayesian Background Rejection

After the reconstruction quality criteria of the previous section have been applied, most air shower events misreconstructed as up-going are eliminated, but some remain. Furthermore, since this analysis seeks to use a portion of the down-going zenith angle range it is also important to treat the background which remains there. Another tool for eliminating misreconstructions is to compare the unconstrained track reconstruction with one which has incorporated a Bayesian prior that the majority of observed events are truly down-going, and should be reconstructed as such. An approximate prior, given by a simple analytic form which has been used in past analyses is shown in Figure 4.6. It is based on past observations of the zenith angle distribution of air shower events, and can be substantially understood in terms of the overburden distribution of Figure 4.2: Where the overburden is smallest, at angles near vertically down-going, many air shower events are able to penetrate. Approaching the horizon, the overburden increases rapidly, decreasing the rate of air showers, until it becomes large compared to the range of even the most energetic muons so no air showers should reach the detector at all and the prior becomes a flat penalty floor.

In the case that an event is correctly reconstructed, the evidence due to the data should overwhelm this prior, and both reconstructions will have the same parameters. Since the data and the best fit parameters would then be the same, the only remaining difference between the two would be the contribution of the prior, which will simply modify the final fit likelihood. In the up-going region, one can make the simplifying assumption that if the unconstrained reconstruction performed better by some constant factor (corresponding to the floor in the prior) than the constrained reconstruction, it was probably trustworthy. For the down-going region it is important to be slightly more careful, as events of either type may be correctly reconstructed

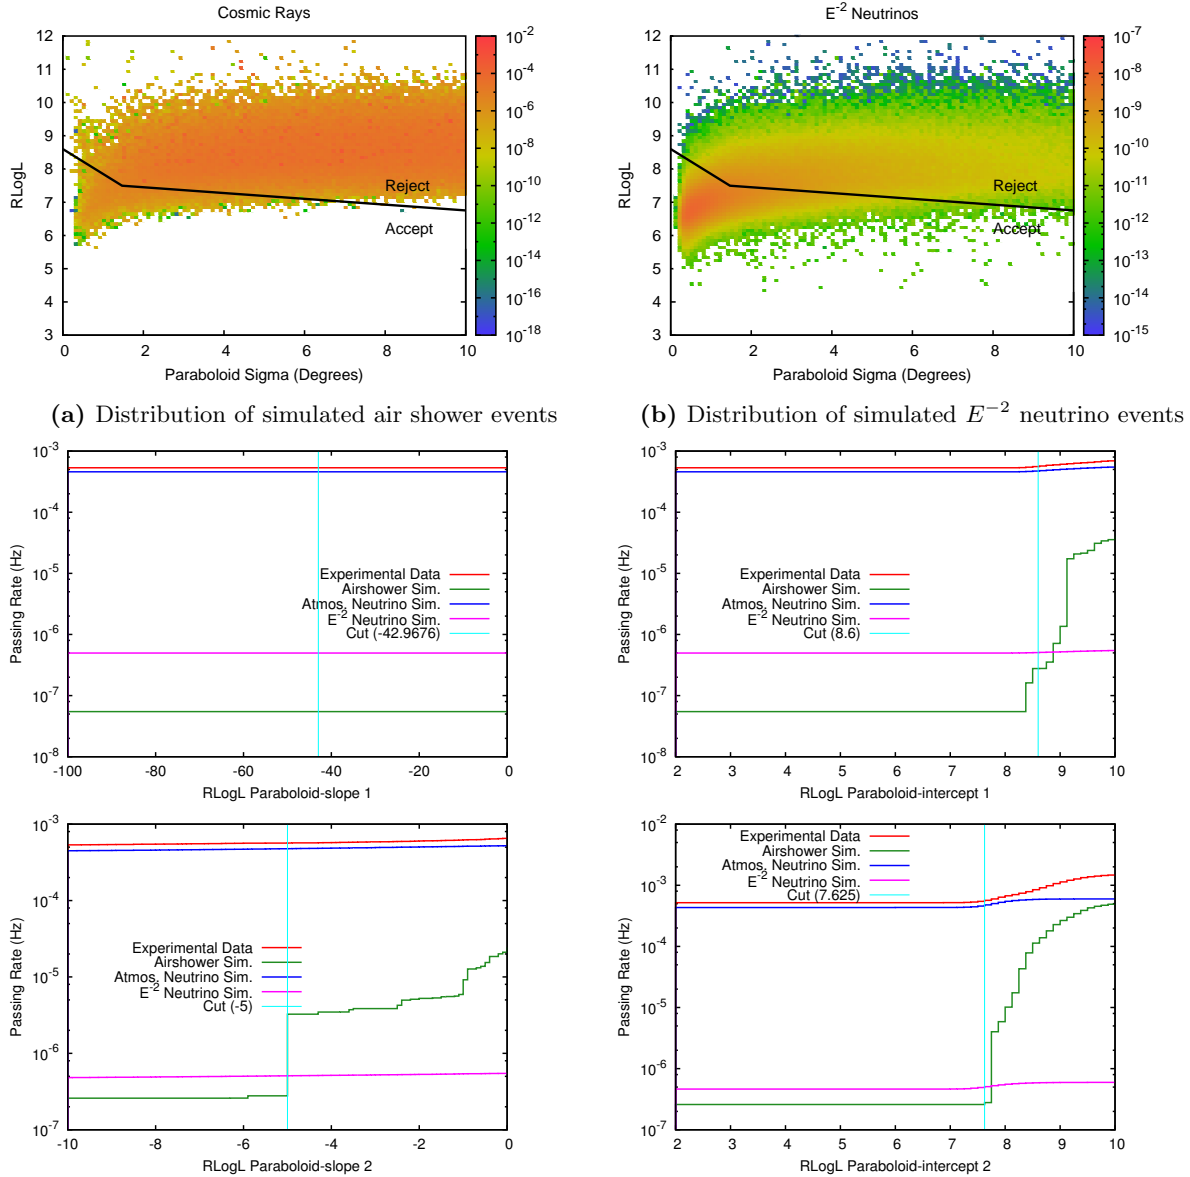


Figure 4.5: Upper two panels: the distributions of simulated air shower background events and E^{-2} signal neutrino events in two measures of directional reconstruction quality. Many air shower events remain in the sample at this stage only because their reconstructed directions are incorrect, so these variables allow separating these types of data. Lower four panels: optimization of the four parameters used to describe the two line segments used to form the selection criteria in the space of the two reconstruction quality measures.




Figure 4.6: The prior used on the zenith angle of the reconstructed track for the Bayesian reconstruction. Note that all up-going directions are heavily penalized, and down-going directions are less penalized the closer they are to vertical.

in this region. However, if they are correctly reconstructed, the reconstructions with and without the prior should be the same, and thus their likelihoods will differ only by the value of the prior, and so the distribution of well-reconstructed events should trace out the prior as a function of zenith angle. This will not be exact, for a variety of practical reasons (mainly, the fits will essentially never actually be identical), but it makes sense to place a cut which surrounds the prior.

Figure 4.7 shows the cut which was defined for the difference of reconstruction likelihoods as a function of zenith angle. In Figure 4.7b it can be seen that the neutrino-induced events from a signal-like spectrum, which are generally relatively high energy and well reconstructed follow a narrow band in the region above the horizon, around which the cut has been placed. The air shower background events in Figure 4.7a are concentrated at the same likelihood difference values, but have much larger spread due to low energy events which are not well reconstructed, and which are eliminated by the cut.

4.4.4 Direction Dependent Brightness Cut

After the Bayesian reconstruction criterion, the remaining air shower background events should generally be well described by down-going reconstructions. As shown in Figure 4.8 this means that neutrino-induced events dominate the up-going region of the data sample, but air showers are dominant above the horizon. The air showers are typically dim, however, since they are forced to penetrate through the overburden equivalent to more than 12 kilometers of water, which the muons produced by neutrinos do not and so can arrive at the

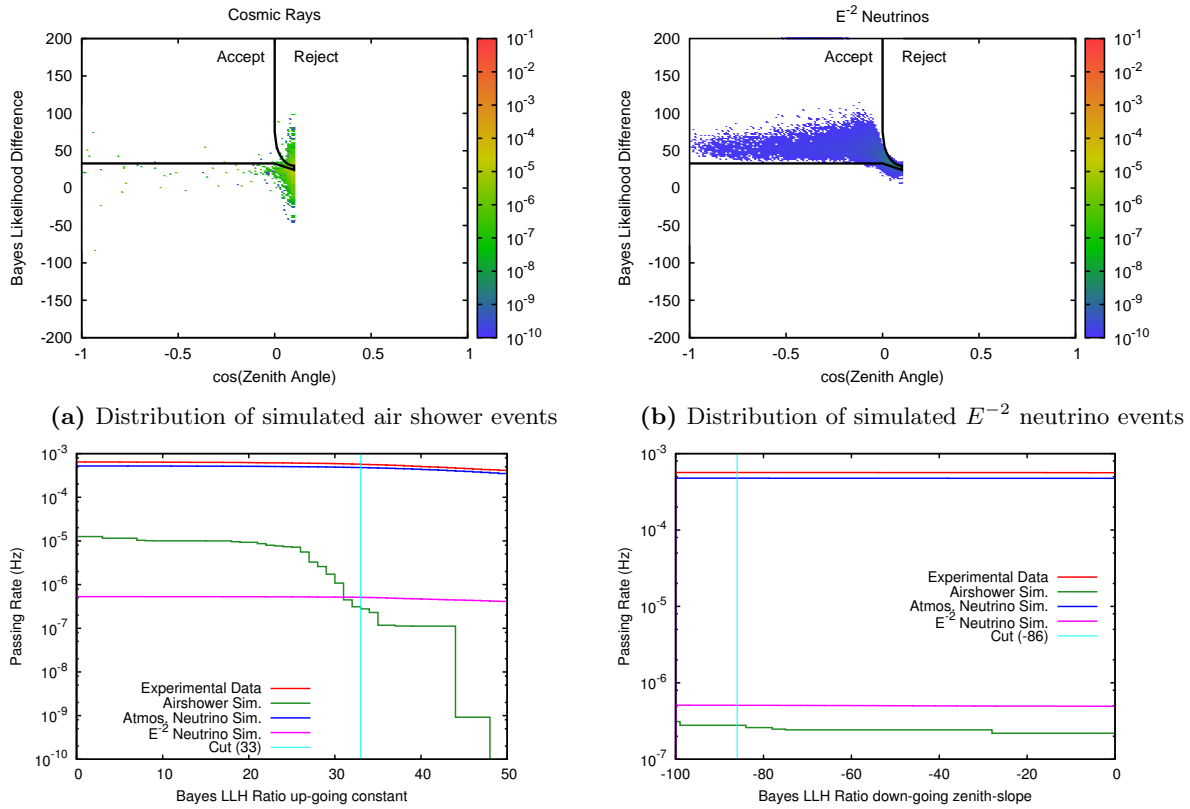


Figure 4.7: Upper two panels: the distribution of simulated air shower background events and E^{-2} signal neutrino events as a function of reconstructed zenith angle and difference of log likelihoods between reconstructions which do and don't contain a Bayesian prior that the events should be reconstructed down-going according to an approximate, known distribution of background event angles. Lower two panels: optimization of the difference in log likelihoods required for up-going events, and the rate (slope) with which this restriction is loosened above the horizon.




Figure 4.8: The reconstructed zenith angle distributions of events after the zenith angle, track quality, and Bayesian reconstruction cuts. Air shower background is low and the data sample matches well to the expectation of conventional atmospheric neutrinos except in the most down-going region, where the experimental data matches air shower simulations.

detector with much higher energies. By sacrificing all events with relatively low energies, as simplistically measured by the number of detector modules which receive light (NChannel), it is possible to eliminate virtually all of the expected air showers at a given zenith angle, while leaving a window for higher energy neutrino events, as shown in Figure 4.9. At the same time, very dim events with up-going reconstructions are also eliminated, as this accounts for few neutrinos which are actually valuable to this analysis, but also a few stubborn air shower events with poor reconstructions.

4.4.5 Overlapping Muon Rejection

Very few simulated air shower events remain after eliminating low-energy events as a function of zenith angle. Of those which do, some are actually coincident muon bundles, which have survived not only the attempt to split up such events but also all cuts so far intended to eliminate poor directional reconstructions. A reasonable approach is then to return to a very simple technique, namely attacking the events by brute force with split fits (Section 3.4.3). Each event is split into two halves twice: once by dividing the observed pulses at the median time, and once by dividing geometrically with a plane through the center of gravity of the pulses and perpendicular to the MPEFit direction. Each half event is then reconstructed using LineFit and iterated SPEFit, although MPEFit is not run for reasons of computational expense. A simple method for testing whether the separate reconstructions of the two halves of each event has actually been more successful than the single reconstruction is to test whether the number of pulses which are ‘on-time’

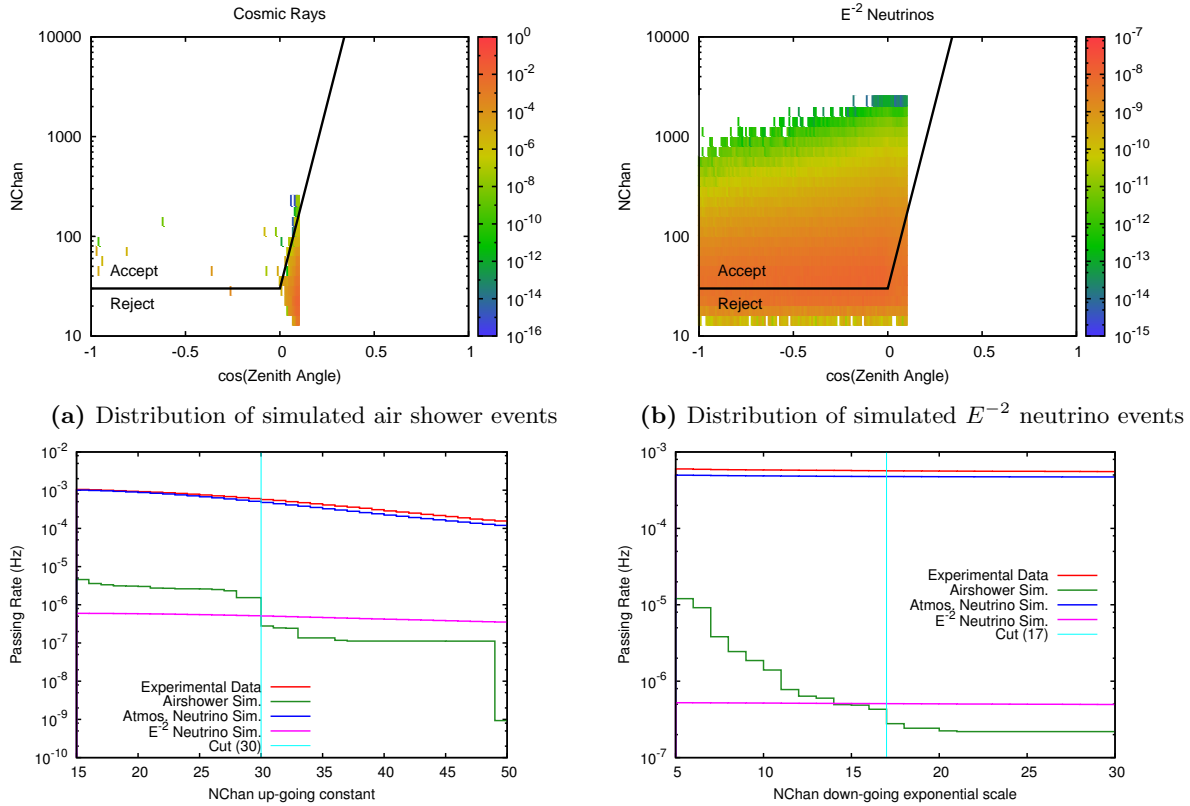


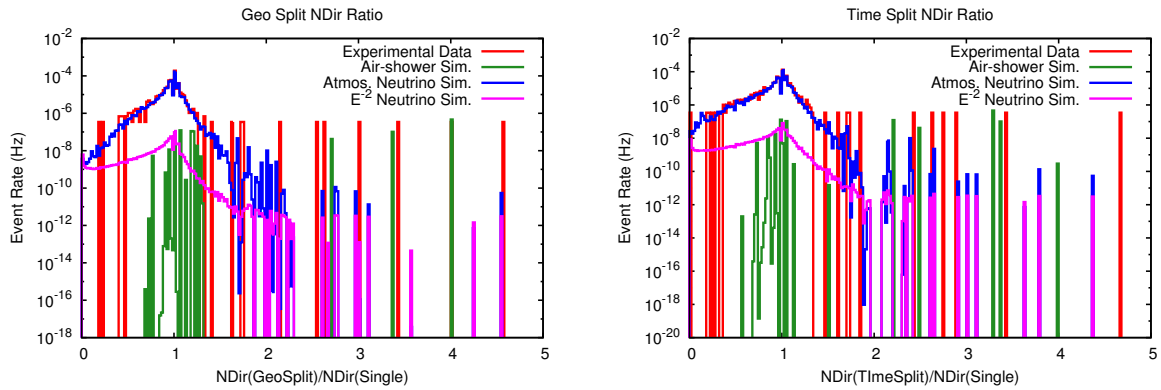
Figure 4.9: Upper two panels: the distribution of simulated air shower background events and E^{-2} signal neutrino events as a function of reconstructed zenith angle and number of detector modules which collected light and were not part of the DeepCore subarray. The neutrino events are often bright (which correlates with having high energy), while the majority of the remaining air shower events are dim. For more vertical angles, where the overburden protecting the detector decreases more air shower events survive with higher energies, so a cut is constructed requiring that events must be brighter (be observed by more modules) the more vertical their angle. A brightness floor is also placed on up-going events, as this eliminates a few more mis-reconstructed air shower events, and while it eliminates many neutrinos these are also dim events which have little usefulness for this work. Lower two panels: optimization of the level of the floor on allowed brightness and the rate of increase (slope) of required brightness above the horizon.

Cut level	Air shower background		Conv. atmos. neutrinos		E^{-2} signal neutrinos
	Rate (Hz)	Number in full sample	Rate (Hz)	Number in full sample	Fraction surviving
MuonFilter §4.1	32.6	1.858×10^9	1.170×10^{-2}	666672	0.932
Data Reduction §4.2	10.6	6.040×10^8	9.127×10^{-3}	520061	0.897
Zenith Angle §4.4.1	3.88	2.211×10^8	8.080×10^{-3}	460403	0.688
Track Quality §4.4.2	4.586×10^{-4}	26130	1.356×10^{-3}	77248	0.313
Bayes. Reco. §4.4.3	1.661×10^{-4}	9463	1.120×10^{-3}	63797	0.286
Zenith Brightness §4.4.4	2.014×10^{-6}	115	5.753×10^{-4}	32784	0.238
Split Event §4.4.5	5.562×10^{-7}	31.7	5.749×10^{-4}	32761	0.238

Table 4.2: Rates and fractions of data surviving by type as a function of the level of selection applied. Efficiencies are with respect to the detector trigger.

compared to the expected time residual from the fit(s) has substantially increased. This parameter will generally always increase, since even in the single-particle case the two fits will have more freedom to match fluctuations of the data, but as shown in Figure 4.10 it rarely improves by a factor of two, while this is a likely occurrence when the event being reconstructed actually consists of two separate tracks. This final cut eliminates a portion of the remaining air shower background events, but is included mostly for safety against contamination unforeseen due to the limited simulation statistics used in developing the selection.

Table 4.2 shows the overall results of the data selection process, including the Muon Filter and the filtering designed specifically for this study. Background from cosmic ray air showers is reduced by a factor of approximately 5.8×10^7 so that it makes up only about 0.1% of the final data, while 23.8% of the neutrinos from a hypothetical E^{-2} flux which trigger the detector are expected to be retained.



(a) Distribution of events by geometric splitting

(b) Distribution of events by time splitting




Figure 4.10: Upper panels: Distributions of the ratio of on-time pulses with two reconstructions to a single track reconstruction for two different possible subdivisions of each event. Lower panel: Cutting on both ratios eliminates some of the remaining air shower background, and provides some protection in case this background is not well described by the limited amount of simulated data.

4.5 Muon and Neutrino Effective Areas

With the full definition of the data selection it is useful to further characterize whether it is efficient. One means of doing so is calculating the effective area the detector presents to a test flux of neutrinos, that is, the area of a hypothetical detector with perfect efficiency which would be required to collect neutrino interactions at the same rate. In general the effective area depends on almost any observable, since the detector and data selection are likely to vary in efficiency for events with different properties, but the behavior as a function of energy and direction are generally most interesting. The upper panel of Figure 4.11 shows the effective area of this data selection for the IC79 data-taking period as a function of true neutrino energy for several different ranges of zenith angles, and averaged over all other event properties. The trend in energy is dominated by the growth of the neutrino interaction cross-section with neutrino energy, which for most energies and angles enhances the rate of neutrinos interacting near the detector and being included in the sample. At large zenith angles, however, the amount of material encountered in passing through the core of the Earth combined with the larger cross-section at high energies to give a substantial attenuation of the neutrinos before they can reach the detector. The lower panel of the same figure shows a comparison between the effective area of this data selection for the IC79 detector configuration to the previous analysis selection, used for the IC59 configuration [98]. At low energies, the effective area of this data selection is smaller, simply because no particular attempt was made to preserve these events, in anticipation of their being quite numerous but also incapable of containing a discernible signal. The two selections are equally effective at around 10 TeV, and at higher energies the larger size of the IC79 configuration, combined with improvements such as including data above the geometric horizon, cause this selection to have a larger area, by about 33% at 100 TeV and more than 60% at 1 PeV.

While the neutrino effective area is a useful diagnostic, it has one major weakness for this type of selection. Since the selection includes (and favors) events with interaction vertices outside of and potentially far from the instrumented volume, interesting signal neutrinos with very high energies will be counted toward the neutrino effective area even if the muons they produce have very low energies by the time they reach the detector, and so are not distinguishable from the more numerous low-energy atmospheric neutrinos. A good sanity check is to also examine the effective area for muons, to ensure that in the cases where high energy muons do reach the detector they have suitably high probability of being detected and retained in the data sample. This is shown in Figure 4.12, where it is compared with the instrumented volume of the detector. From this plot, it can be seen that this selection retains about 60% of muons at 10 TeV, 75% at 100 TeV,




Figure 4.11: Upper panel: The effective area of the detector for a flux of neutrinos after all data selection criteria have been applied. The area is shown for five zenith angle bands, showing the impact of absorption in the Earth for high energy neutrinos at high zenith angles, and the average over all zenith angles for which there is any acceptance is overlaid in green. Lower panel: Comparison of the effective area for this analysis (labeled ‘IC79’) to the previous analysis (labeled ‘IC59’). This analysis has greater acceptance for neutrinos with energies greater than about 10 TeV, but makes no particular effort to optimize for lower energies and so has lower acceptance for such.




Figure 4.12: The effective area of the detector after all data selection cuts for a flux of muons, compared to the true geometric area of the detector, averaged over all included zenith angles.

and just under 90% at 1 PeV. So, as desired, this selection is efficient for collecting the neutrino events in which a high muon energy indicates a necessarily high-energy neutrino energy.

4.6 Final Data Properties

In addition to the total amount of data selected and its purity, it is important for a successful analysis that the data be sufficiently well reconstructed that the observables contain useful information about the physics. The main observables of interest in this work are the direction and energy of the muons, and the resolution with which these are reconstructed can be determined by examining simulated data, for which the true values are known. The results for the angular resolution are shown in Figure 4.13, which shows the median difference in the angle between the true and reconstructed directions of muons produced by two different neutrino spectra as a function of reconstructed muon energy. Over the whole energy range relevant to this analysis the angular resolution is better than 1° , and at the high energies where signal events may be detected it is 0.5° or better. It should be noted that the two spectra have slightly different angular resolution behaviors because the mapping of true energies to reconstructed energies is not one-to-one (as will be discussed shortly), so depending on the true neutrino spectrum, an observed event with a given reconstructed energy will have




Figure 4.13: Median angular resolution for neutrino events in the final data sample for two representative spectra. At all energies the resolution is better than 1° , and above a few tens of TeV it is generally about 0.5° .

a different most probable true energy. (Plotting in terms of true energy eliminates this difference, but is somewhat less useful since we would like to understand the expected properties of actually observed events, for which this quantity will not be known.) In any case, the angular resolution is easily better than what is particularly needed for this analysis.

More central to this analysis will be the energy resolution. Unfortunately, energy resolution for long ranged muons is intrinsically poor in a detector of the style of IceCube, since it is based on measuring energy loss rates, rather than calorimetric measurement of the total particle energy (see Section 3.4.4). Figure 4.14 shows two different ways of visualizing this aspect of the data. First, the distribution of reconstructed energy values as a function of true muon energies (at the point of closest approach to the detector center) is shown, with each (vertical) band in true energy individually normalized, so that the spectrum of true muon energies is irrelevant. As expected, there is substantial similarity to the theory of muon energy loss from Figure 3.2, namely that at high energies muons lose energy proportionally to their total energy, allowing the total energy to be distinguished, while at low energies the loss rate becomes nearly constant, so different total energies become indistinguishable. The vertical spread of the distribution arises from the statistical variation of events with the same true energies, and is the factor which limits the energy resolution. There is also a bias in the distribution, in that the most probable reconstructed energy for a given true energy has systematic differences from the true energy, but this is merely an annoyance which does not affect the

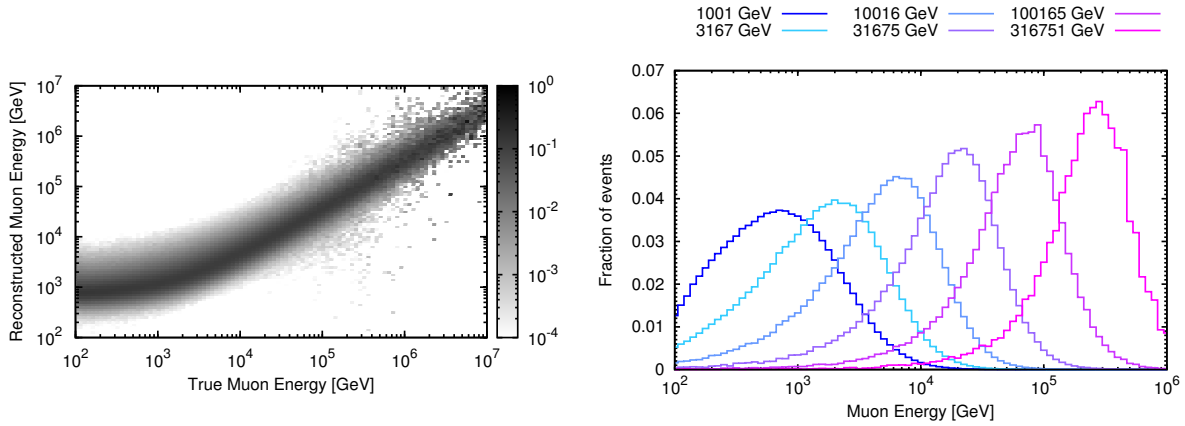


Figure 4.14: Left panel: The relationship between true and reconstructed muon energy in the final event sample, which is mostly linear at energies above 1 TeV, but has not been calibrated to minimize bias. As such, the most probable reconstructed energy for a given true energy is not necessarily the same, e.g. for a true energy of 10 TeV the most probable reconstructed value is nearer 5 TeV. Right panel: A visualization of the uncertainty of the energy reconstruction. Each distribution is the distribution of inferred true energies for the single true energy labeled in the legend, and it can be seen that at high energies the uncertainty in inferred energy approaches a factor of two.

forward-folding analysis used in this work (see Section 5.1.2). The energy resolution itself is represented in the second panel of Figure 4.14 by showing the distribution of inferred most-probable muon energies for several different true muon energies. These distributions do depend on the true neutrino energy spectrum, so for convenience a measure of foresight has been applied by using the spectrum found as the result of the analysis in Section 6.2. The widths of these distributions are essentially the uncertainty in the reconstructed energy, and it can be seen that they steadily decrease with increasing energy to an uncertainty of about a factor of 2. This is not as good as one might desire, but it is important to remember that this is not the only effect which smears the true distribution of neutrino energies; there is also the energy lost to the hadronic recoil at the interaction vertex and the energy lost by the muon while traveling from the vertex to the detector. These are an inescapable limitation on the precision of analyses which use events which interact outside the instrumented volume, and is the price which must be paid for the corresponding increase in effective collecting area.

Chapter 5

Analysis Method

The general approach taken in this work to study possible astrophysical neutrino fluxes is a ‘forward-folding’ analysis, which consists of combining a hypothesis of an input neutrino flux with a description of the detector response, derived from simulation, to predict the events observed by the detector. In this chapter, both the theory of such tests and the particular model constructed for this analysis are discussed.

5.1 Likelihood Fit

Because this analysis depends on Monte-Carlo simulations to estimate the response of the detector to signals, continuous PDFs for event observables are not directly available. While such could be made, they would be approximations to the discrete, simulated events, so it is simpler to classify the simulated events by their observable properties into histogram bins and use the result as a binned pdf. This pdf, multiplied by the time for which the detector operates to record data (the ‘lifetime’) is then a binned prediction of the detector’s observations (the expectation), and the experimental results may be likewise binned in the same observables to create a corresponding distribution of what is actually observed (the observation). It is then necessary to define a measure of the agreement between an expectation and an observation. A suitable construction is a Poisson likelihood: In some bins the expected number of events may be quite small so that statistical fluctuations are not Gaussian, while in bins where large amounts of data reside Poisson errors asymptotically become Gaussian. The total Poisson likelihood is simply the product of Poisson probabilities for all bins, where e_i is the expected number of events in bin i and x_i is the observed number:

$$\mathcal{L}(\{\theta_i\}|\{x_i\}) = \prod_i \frac{e_i(\{\theta_i\})^{x_i}}{x_i!} e^{-e_i(\{\theta_i\})} \quad (5.1)$$

Here we have also allowed the pdf, and thus the expectations and the likelihood, to be a function of some model parameters θ_i .

A hypothesis which matches the data more closely will have a larger likelihood than another hypothesis which matches less well, so we can find the best hypothesis out of a parameterized family by maximizing the likelihood with respect to all parameters [99]. For computational purposes it is more convenient to minimize the negative logarithm of the likelihood (converting the product in Equation 5.1 into a sum), which yields the same parameter estimate, due to the monotonicity of the logarithm.

In addition to determining the best parameter values for the hypothesis, it is valuable to be able to quantify the degree to which that choice of parameters is better than other choices. A case of particular interest is comparing the best fit obtained when maximizing over all model parameters (the global best fit) to the conditional best fit obtained maximizing over some subset of the parameters while some number, n , of the parameters are held fixed to specific values of interest (the constrained or conditional best fit). Letting $\{\hat{\theta}_i\}$ be the resulting parameters of the global best fit and $\{\hat{\theta}'_i\}$ be the parameters of the constrained best fit, we can define a test-statistic:

$$\lambda = 2 \ln \left(\frac{\mathcal{L}(\{\hat{\theta}_i\}|\{x_i\})}{\mathcal{L}(\{\hat{\theta}'_i\}|\{x_i\})} \right) \quad (5.2)$$

Large values of the test statistic indicate that the observed data are more likely sampled from the distribution given by the model with parameters $\{\hat{\theta}_i\}$ than by $\{\hat{\theta}'_i\}$. Specifically, as stated by Wilks' Theorem [100], in the case that the total amount of observed data, N , is large, λ is distributed as a χ^2 distribution with n degrees of freedom, up to errors of the order $1/\sqrt{N}$. Specifically, this means that if data is drawn repeatedly from a distribution described by the n fixed parameters of the conditional best fit (treating this as the null hypothesis), and both the global and constrained maximum likelihood fits are performed, the distribution of test-statistics should be a χ^2 distribution. This result, known as Wilks' Theorem, provides a convenient means of estimating the p-value or significance for distinguishing the global best fit obtained via likelihood maximization from a null hypothesis of related form which restricts some parameters. The theorem does not always hold, however, so it is sometimes necessary to obtain the test-statistic distribution empirically using Monte-Carlo simulation in which a possible data realization is sampled from the null hypothesis distribution, both maximum likelihood fits are performed, and the resulting test-statistic is added to the accumulated distribution.

Finally, we may wish to include constraints on model parameters based on existing knowledge. This is most easily done using Bayesian prior probability distributions, each of which is simply a probability

distribution which is a function of one of the model parameters, which is multiplied onto the likelihood (or added to the log likelihood.) Priors penalize the likelihood for values of the parameters which have small prior probability, so that when the data does not constrain the parameter strongly its value will be determined by the prior. If, however, a parameter disfavored by the prior is required to explain the data, sufficient data will overwhelm the prior (that is, if the penalty to the likelihood from failing to match the model to the data is larger than the penalty imposed by the prior).

5.1.1 Error Estimation

Given a maximum likelihood set of model parameters it is useful to also be able to compute confidence intervals for these estimates. For a model with a single parameter the obvious approach is to form the confidence interval using the likelihood ratio test to identify all models which are not excluded compared to the best-fit model at the chosen confidence level. This entails simply computing the likelihood at each point in the space for comparison to the best-fit likelihood, which can be viewed as a conditional best-fit with no free parameters (and thus if Wilks' Theorem is used to approximate the test-statistic distribution the necessary number of degrees of freedom is just the same as the total number of model parameters). As the number of model parameters increases, however, this becomes unfeasible, as the parameter space volume to scan increases exponentially, and reporting the shape of the multidimensional surface is awkward at best. In general, it is not possible to capture all possible correlation among parameters with a lower-dimensional approximation, but if this information can be reasonably foregone a profile likelihood may be used. In this technique, the dimensionality is reduced by considering only one parameter or a few parameters at a time (termed 'structural' parameters in the traditional nomenclature), while all others are treated as secondary 'nuisance' parameters (or 'incidental' parameters, traditionally). The parameters or parameters of interest are then scanned over, being held fixed while the likelihood is maximized at each point with respect to the nuisance parameters. This is a conditional likelihood with a test-statistic whose distribution can be approximated under Wilks' Theorem as a χ^2 distribution with a number of degrees of freedom equal to the number of parameters of interest. Constructing a confidence interval in this way for each parameter in turn (while treating all others as nuisance parameters) gives a reasonable approximation of separate estimates of the uncertainty on each parameter.

5.1.2 Note on Biased Observables

It is generally desirable that a statistical estimator have both low bias (correlated, systematic difference from the true parameter) and low variance (random statistical difference from the true parameter). It is then a common cause for concern that a likelihood fit such as has been discussed here will give incorrect results if one or more of the event observables (e.g. reconstructed event energy) used to form the observable distribution is biased, so that it systematically misestimates the true physics parameter (true event energy). In fact, it is a strength of the forward-folding approach that this typically causes no difficulty. The reason for this lies in the way that only observable distributions are compared: the extracted fit parameter values are those which maximize the likelihood by best matching the experimental observable distribution. This hinges critically upon the observables having the same behavior for both the experimental and simulated data, but as long as this is true, any monotonic transformation of the observables is theoretically irrelevant. The requirement of monotonicity arises because the transformation must not lose information by mapping two input values to the same output value, as this might mean that two different combinations of fit parameters could produce the same observable distributions. Likewise, in practical terms the transformation function should not be too slowly varying, as it would then map many inputs to outputs which while technically different might be too similar to be distinguished. This still leaves substantial room for commonly occurring biases, such as constant shifts, or observables which differ from the true value by a factor.

It should be noted that, in the context of this analysis, it is not necessarily possible to remove the bias from relevant estimators using only a priori knowledge. In the case of the muon energy reconstruction, the relationship between the true parameter and the observable depends on the distribution of true parameter values, that is, the true neutrino energy spectrum. Without performing the analysis fit, the spectrum is assumed to be unknown (at least in the high energy region of interest) due to the possible inclusion of an astrophysical component whose strength and exact form are not known, so the observable cannot be calibrated for use as an input to the fit. Furthermore, after the fit has been performed, the knowledge gained about the spectrum is limited, so it must be understood that any a posteriori calibration is conditional upon the fit result, and may still not be a very good approximation to the truth.

5.2 Model Construction

5.2.1 Note on Weighting

In order to perform the forward-folding analysis it is necessary to have a description of the detector response, but, as has been mentioned previously, for IceCube this response is too complex to be written in an analytic form, and must be approximated using simulated data. In principle the calculation of the likelihood for a given model is conceptually trivial: simply perform a complete simulation with behavior governed by the model in question, and use its output as the expectation. Simple though this idea is, it is also completely impractical because of the vast amount of computation required to create near duplicate simulations for different model parameter values. A far better approach is that whenever possible the simulation should be designed so that the dependence of its output on the model parameters can be calculated after the simulation has been run, so that only one simulation dataset must be produced, which can then be adapted to form all necessary models.

We therefore introduce the concept of a simulation weight, which is a function of both model parameters and event properties so that when all simulated events are entered into a histogram with contributions proportional to their weights, the complete histogram approximates the ideal pdf for the chosen model. Simple inspection indicates that if creating a histogram of all events with unit contribution produces arbitrary, simulated distribution, then the per event weight must be defined in terms of the ratio of the desired distribution to the simulated distribution. Specifically, for each event the weight should be the ratio of the probability for that event to occur in the desired (model) distribution to the probability for that event to be produced by the simulation. This frees us not only from having to run the simulation to produce any particular physical distribution but also from having to produce any physical distribution at all: As long as both the target and simulated distributions are fully understood, one may be transformed freely into the other. This procedure does not eliminate all possible problems, however. In particular, as defined, it will produce a distribution whose mean value is that of the target distribution in any region of phase space, but it will have fluctuations which depend on the amount of simulated data being weighted to form that mean. If in some region the simulation has a high generation probability compared to the observation probability in the target model, many events will be produced there and each assigned small weights, and the statistical error on the resulting mean can be far smaller than the counting error on the same mean produced by unit events. This situation is useful but the inverse may also occur: if in a region of interest the simulation produces

very few events where the model has high probability, each event will be assigned a large weight to represent many more events, and the mean will have far larger errors than, might be expected from normal counting. This problem can be solved by brute force (simply producing a larger corpus of simulated data), but we see then that it is advantageous for the simulated distribution to be sufficiently similar to the distributions we wish to study, as this will also help to reduce occurrence and extremity of the second case (although it will at the same time reduce the first).

5.2.2 Atmospheric Parameters

The model to be fit must account for all components of the flux which are expected to be present in the dataset, so it must treat both the conventional and prompt atmospheric neutrinos. The forms assumed for these fluxes are the ones discussed in Sections 2.3.1 and 2.3.2: The HKKM conventional flux [51] and the ERS prompt flux [54], each corrected for the cosmic ray knee according to [53]. For each of these there are a few model parameters of interest. First, since both fluxes depend on the production rates of mesons in air showers, their total normalizations are somewhat uncertain. So, we choose to make both normalizations free parameters in the likelihood fit, constrained only to be non-negative: We could place a prior on the conventional normalization based on past measurements, but since this flux can be expected to thoroughly dominate the data sample we should have easily enough new data to constrain it. The prompt flux, however, has never been conclusively observed, so it makes sense it insert no assumption about its strength.

In this same vein, the relative contribution of pions and kaons to the conventional flux is not exactly known. This is somewhat constrained by measurements of the charge ratio of atmospheric muons, and its uncertainty is estimated to be 10-15% at 100 GeV-1 TeV, so we additionally allow the normalization of the kaon contribution to float within the overall conventional flux term, with a Gaussian prior of width 10%, centered on the kaon contribution predicted by the HKKM model.

Finally, both of these fluxes depend directly on the flux of cosmic rays. Since we have already allowed their normalizations to float we do not need to treat the uncertainty in the overall normalization of the cosmic ray flux, however we should also treat the uncertainty in its shape. At \sim TeV-100 TeV energies, measurements of the spectral index of the cosmic rays still show disagreement, such as 2.66 ± 0.02 observed by the CREAM experiment to 2.78 ± 0.009 observed by the AMS experiment [3]. We therefore include a parameter in the likelihood to express the change in the cosmic ray spectral index from what is assumed in our models [52]. This is implemented in the form of a term which multiplies the atmospheric neutrino

spectra by a power law with variable index, in each case pivoted about the approximate median energy of the spectrum to minimize degeneracy with the normalization parameters.

The atmospheric portion of the model is then

$$\begin{aligned} \Phi_{\text{atm}}(E_\nu, \theta) = & N_c (\Phi_{\text{HKKM},\pi}(E_\nu, \theta) + N_K \Phi_{\text{HKKM},K}(E_\nu, \theta)) \left(\frac{E_\nu}{E_{\text{conv,median}}} \right)^{-\Delta\gamma} \\ & + N_p \Phi_{\text{ERS},\pi}(E_\nu, \theta) \left(\frac{E_\nu}{E_{\text{prompt,median}}} \right)^{-\Delta\gamma} \end{aligned}$$

where the model parameters are N_c , the conventional normalization, N_K , the kaon normalization factor, N_p , the prompt normalization, and $\Delta\gamma$, the change to the cosmic ray spectral index.

5.2.3 Astrophysical Parameters

For the signal portion of the fit we must assume the form of the astrophysical neutrino flux. Our baseline choice, for consistency with past analyses is a pure E^{-2} power law flux, whose only free parameter is its normalization, on which we place no prior. For greater flexibility, we choose to also consider general power law fluxes with arbitrary spectral indices as well. These two cases of the model can be described by

$$\Phi_{\text{astro}} = N_a (E_\nu/100 \text{ TeV})^{-(2+\gamma)}$$

with model parameters N_a , the signal flux normalization, and γ , the change from an E^{-2} power law, which we fix to zero when we wish to test the fixed power law hypothesis only.

Finally, we may wish to explicitly test some of the theoretical models discussed in Section 2.2, and to do so we simply substitute parameterizations of these models into the formula above in place of the pure power law; we make no effort to allow for variation in the shapes of these models, leaving only N_a as a free parameter.

5.2.4 Detector Parameters

Systematic uncertainties in the detector itself can also be incorporated into the likelihood as long as they can be parameterized. One such parameter is the absolute optical efficiency of the detector modules (the ‘DOM efficiency’) for collecting photons from Čerenkov spectra. In principle the behavior of this parameter is simple: If a unit light source emits N photons/GeV, and we believe that our detection efficiency is ϵ , when

we observe an event with M detected photons we will conclude that it had energy $E = M/\epsilon N$. If we have overestimated ϵ by some factor, we will then underestimate E by this same factor. This simple picture is broadly true at high energies where the numbers of photons involved are large and the assumption that all behavior scales linearly holds, but it is not correct at low energies where the difference between observing and not observing a single photon can mean the difference between the detector triggering or not triggering. So to parameterize this effect we must first simulate it in detail. To do this, five simulated datasets were created which were identical except for the simulated value of the DOM efficiency; the values used were 0.9 (the baseline value assumed for all simulation at the time), 0.95, 0.99, 1.089, and 1.1979. The fact that some of these values are larger than one is not a concern, since within the software they are multiplied by several other values which are considered constant which make the overall efficiency less than one in all cases. The goal of the parameterization is then to be able to reweight a given simulation to produce the flux which would have been observed, had the actual DOM efficiency been some value ϵ' instead of the value which was simulated for that dataset of ϵ . It is trivial to observe that this can be accomplished by multiplying each simulated event by the ratio of rates predicted for events of that type in simulations made with values ϵ' and ϵ ; the only great difficulty occurs if ϵ' is not a value which has ever been explicitly simulated. The only practical way to address this is to interpolate the results of simulations with values above and below ϵ' to estimate what this result would have been, and so to do this an interpolating spline was constructed among four of the simulated datasets (those with efficiencies 0.9, 0.99, 1.089, and 1.1979) as a function of both DOM efficiency and event energy proxy. The fifth dataset could not be used in the fit for technical reasons (DOM efficiency values which were evenly spaced, or evenly spaced in log space, were required) but this was perfectly acceptable as it could then be held in reserve to test to success of the spline at interpolating to a value to which it had not been fit. Since the distributions of events in the observable energy differ for each spectrum component (and thus their ratios differ as well), this procedure was actually carried out six times: once for each spectrum component (conventional atmospheric, prompt atmospheric, and astrophysical) and for each data taking period (2010 and 2011); part of one of these fits is shown in Figure 5.1. With this parameterization in hand it was possible to insert the DOM efficiency into the flux model calculation as a free parameter, and based on previous measurements [66] it was given a prior with width 3%, centered on a value of 0.989.

Another systematic uncertainty which one would like to treat is the uncertainty on the optical properties of the ice from which the detector was constructed. In particular, in the calculation of the ice parameters

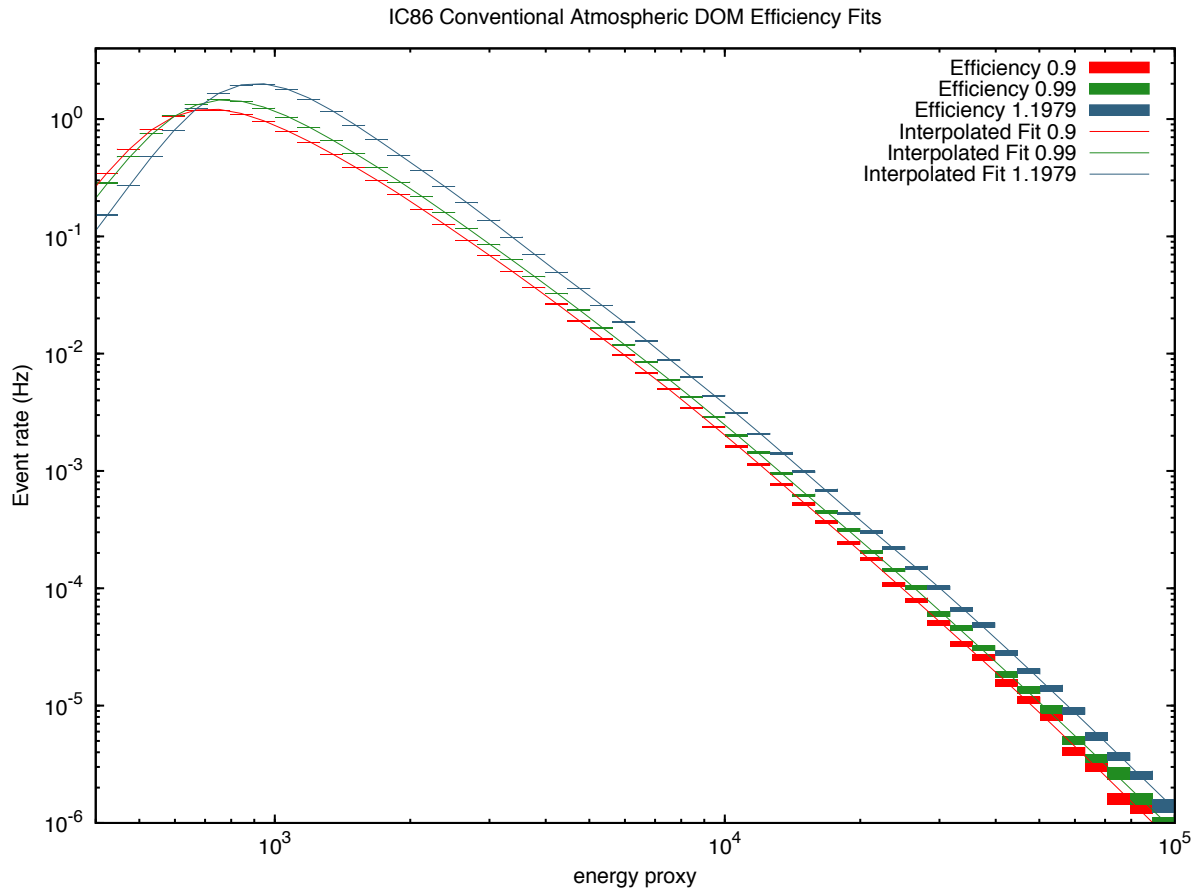


Figure 5.1: Partial representation of one of the six DOM efficiency parameterizations needed in this analysis. Each set of boxes represents the event rate predicted for this flux (the conventional atmospheric neutrinos) in this data taking period (2011) when the detector is simulated with a different DOM efficiency. Five datasets were actually used, although only three are shown here, for clarity; the efficiency 0.95 and 1.089 datasets have been hidden. The widths of the boxes simply show the histogram binning, while the heights show the statistical errors from the simulation. The smooth lines are the spline fits to the datasets with matching colors. The range of fitting also extends to higher energies which are cut off on this plot.

used by this analysis [62] there was found to be a $\sim 10\%$ uncertainty in the overall quantity of both scattering and absorption in the ice. Like the DOM efficiency, there is no obvious parameterization of the effect of changing these properties on the high level data, so direct simulation is necessary. Unfortunately, it was discovered that the quantities of simulation generated for this purpose were substantially too small, such that they contained statistical errors which washed out any systematic effect under study. It was therefore not possible to meaningfully include this uncertainty in this analysis.

5.3 Analysis Sensitivity

As noted in Section 4.4, one would in principle like to optimize the sensitivity of the analysis with respect to the parameters of the data selection, but this is generally too impractical to actually carry out. Now, however, with both the data selection and the analysis method fully defined we can compute the sensitivity and at least determine whether it is acceptable. By the sensitivity, we mean that we wish to know how small a signal this analysis will be able to reliably distinguish from the expected background, and we choose to define this formally as the median upper limit, at 90% confidence, which the analysis finds for the signal flux when run on data which contains no signal. This can be easily computed by weighting the simulated neutrino events to the conventional atmospheric spectrum, and then drawing from them possible realizations for the observed data by choosing a total number of events to draw from a Poisson distribution with mean equal to the mean expected number of events, and then sampling that number of events treating pool of simulated events as a multinomial distribution where each event has a probability to be selected equal to its weight (expected rate of observation) divided by the sum of all event weights (the expected total data rate). This calculation is shown in Figure 5.2, with the primary result that the analysis should have a sensitivity of $2.2 \times 10^{-19} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$ for an E^{-2} signal flux. This compares quite favorably with the sensitivity of the previous analysis [98], for which a sensitivity of 6.9 in the same units was expected when using one year of data (as opposed to the two used here) with a detector 70-75% as large as the one used here. One other advantage of this analysis is its use of data in the zenith angle range from $85 - 90^\circ$. To estimate the impact of this change the sensitivity was also calculated without this angular region, yielding a sensitivity of $2.7 \times 10^{-19} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$, which implies that this addition gives about a 19% gain in sensitivity (the number of events expected from an E^{-2} flux is increased by only 8.8%, but these are biased to higher energies which are more distinguishable from the atmospheric fluxes), but by no means dominates the data sample.




Figure 5.2: The expected sensitivity of this analysis estimated by Monte Carlo trials with and without the down-going portion of the data.

Chapter 6

Results

6.1 Observed Events

Before proceeding to the main analysis, in the form of the likelihood fit, it is useful to examine the experimentally observed events somewhat to ensure that there are no obvious problems. The full unblinded data sample turns out to contain 35322 events, which is too large a number to examine manually in detail or discuss here. Manual examination of a few thousand events found no major problems, although the rate of coincident background contamination appears to have been somewhat underestimated. The most important portion of the data sample is the tail of the energy proxy distribution containing the most energetic events, which are the most likely to be contributed by the target signal flux.

The twenty events with the largest energy proxy values are listed in Table 6.1 in order of time of observation, and further details are shown in Appendix A. As expected most of these events have rather horizontal directions, since if they truly have energies of 100 TeV or greater the Earth would have been likely to absorb them. An unexpected feature, however, is that several of these events are reconstructed as never actually entering the instrumented region of the detector. An example of this, which is actually the event with the highest reconstructed energy in the sample (116357,6324295 or ‘Dr. Heinrich Faust’), is shown in Figure 6.1. By eye, the reconstructions of some of these events are not entirely convincing; while it is possible that they are correct it also seems plausible that they could be completely wrong.

A small side study was therefore undertaken to attempt to determine whether this class of events does indeed represent well-reconstructed neutrino-induced muons or whether they are air shower background events which have survived the data-selection process (and would compromise the correctness of the subsequent spectral analysis). Two avenues of verification were explored: first to use more advanced reconstruction techniques to check the correctness of the reconstruction otherwise used in this analysis and for the event

Table 6.1: Basic properties of the 20 observed events with the highest energy proxy values. For more details including energy estimates based on the spectrum fit results see Section A.

Event ID	Data Period	R.A. (°)	Dec. (°)	Energy proxy (TeV)	Inferred Muon Energy (TeV)	Name
116026,44241207	2010	346.9	24.1	133.2	339	Dr. Henry Jekyll
116082,62251639	2010	139.0	47.6	51.1	116	Angus MacGyver
116147,14170716	2010	310.6	22.0	55.2	124	Dr. Moreau
116269,59516168	2010	323.5	2.8	66.7	156	Dr. Strangelove
116357,6324295	2010	254.1	16.3	290.1	755	Dr. Heinrich Faust
116574,20123342	2010	267.6	13.8	132.1	302	Dr. Emmett Brown
116701,6581938	2010	331.1	11.1	141.9	317	Dr. Richard Seaton
116807,9493609	2010	88.6	0.2	199.8	604	Dr. Hari Seldon
116876,63208734	2010	110.6	0.0	65.1	134	Dr. Strangepork
116883,17395151	2010	285.8	3.1	147.4	422	Captain Nemo
117639,30571557	2010	308.1	1.0	53.4	112	Buckaroo Banzai
117927,15766169	2010	207.4	6.7	60.9	139	Professor Joseph Cavor
117946,40122789	2010	181.8	38.6	52.0	116	Dr. David Bowman
118210,47538807	2011	235.4	19.3	107.4	252	Dr. Victor Frankenstein
118475,52691508	2011	152.4	6.8	60.2	125	Professor Abraham Van Helsing
118615,37865356	2011	31.4	11.9	50.3	109	Dr. Giacomo Rappaccini
118631,36844560	2011	9.6	7.9	64.0	147	Impey Barbicane
119037,60175569	2011	222.0	3.2	50.7	109	Professor James Moriarty
119136,66932419	2011	37.2	18.7	160.2	397	Dr. Susan Calvin
119739,41603205	2011	238.4	18.9	131.8	326	Dr. Henry Walton Jones, Jr.

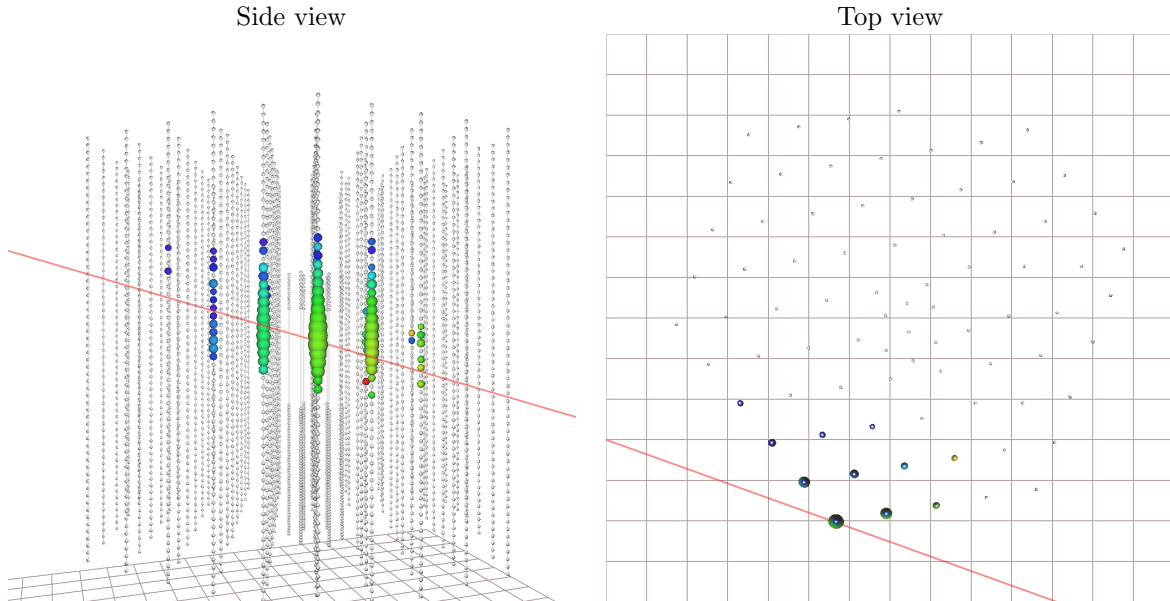


Figure 6.1: Visualization from side and top perspectives of the observed event with the largest energy proxy value. The event's reconstruction indicates that it is up-going, but that it never enters the instrumented volume of the detector, merely passing very close to one edge.

selection (MPEFit), and to attempt to generate simulated background events with properties similar to the observed event in large numbers to estimate more precisely the rate at which such events can survive the selection process. This full procedure (due in equal parts to its success and its time consuming nature) was carried out only for the highest reconstructed energy event, with results which will be described next. The reconstruction portion only was repeated for the remaining top events, without significantly different results.

The reconstruction portion of the study was conducted using the ‘Millipede’ reconstruction framework. Unlike the methods introduced in Section 3.4.2 and otherwise used in this work, this method is founded on basis of energy reconstruction (Section 3.4.4), in that it fundamentally fits predicted light output to observed charge, using the recorded amplitude information, although timing information is still used. In addition to using amplitude information, this framework is not restricted to fitting single uniform particles, but instead unfolds the necessary contributions from ensembles of particles, which are taken in this case to be the stochastic energy losses of the muon, represented by a collection of cascade-like energy depositions arranged collinearly, but with the entire group free to be moved as a unit by the reconstruction. The expected light yield from each energy loss at each observing detector module is estimated using multidimensional spline interpolations of photon propagation tables (Section 3.5.4, [82], which give reasonably good detail, although they are unable to treat some of features of the ice which break symmetries, such as tilt and azimuthal anisotropy). As a result, this reconstruction technique is both substantially different from the simple parameterization based method used otherwise both in its treatment of the structure of muon energy losses and ice variations (it is unfortunately far too computationally expensive to be used generally). Because generic optimization of the direction of the muon track for the Millipede likelihood has often been observed to perform poorly (possibly because of a significantly more complex likelihood space) instead a brute-force search over track direction parameters was performed, while leaving the vertex position parameters free to be optimized normally. The map of resulting likelihoods is shown in Figure 6.2, where lower values (blue colors) indicate better fits. The best direction found by this scan is not identical to the one found by the more simplistic reconstruction, but it is close and it is also up-going. Furthermore, the scan tested all down-going directions (with some granularity), so it can answer whether any of them provides a plausible alternative, the answer to which is unambiguously negative, as shown in the lower panel of the same figure: When picking the best azimuthal angle for each zenith angle, all down-going reconstructions are excluded by the best fit at over 13σ . There is, then, some reason to be confident that the reconstructed direction of this event, leading to its classification as a neutrino, is correct.




Figure 6.2: Results of the detailed reconstruction study for the observed event with the highest energy proxy. Upper panel: Likelihood skymap comparing possible reconstructed directions for the event. Blue colors indicate better likelihoods, red colors are worse, and gray pixels indicate missing data. The white \times indicates the best direction found by the scan, and can be compared to the red \times which was the standard reconstruction (MPEFit) for this event; although the two are not identical they are fairly similar, and in particular have nearly the same zenith angle. It is interesting to note that there is a second, less favorable minimum for a direction $\sim 150^\circ$ different in azimuth, which would correspond to a muon passing the detector edge in approximately the mirror image direction. Lower panel: Relative probabilities of the best reconstruction found by the scan at each zenith angle. All downward-going solutions are worse than the best (upward-going fit) by a factor of at least e^{90} , so all down-going (or horizontal) hypotheses are disfavored at more than 13σ .

The other approach to testing whether this event is a valid neutrino is to simulate a large volume of non-neutrino events to test whether they can ever survive the data selection process and yield similar observables. Two optimizations were applied to do this expediently: First, full air shower simulations were not run. Instead single, down-going muons were generated, and weighted according to the expected spectrum produced by air showers [101]. This is justified on the grounds that a blob-like event topology with little clearly defined direction can be best achieved by a single large, stochastic energy loss, which is far more likely to be produced by a single muon than a bundle of muons of the same energy, since in the bundle either one muon would have to undergo a proportionally larger (and thus less likely) energy loss, or multiple muons would have to undergo coincident energy losses. The second optimization in this simulation was based on the fact that the experimentally observed event is known to have passed near a particular point at one edge of the detector, therefore all simulated events were aimed at a volume surrounding this area by a few tens of meters. Since the actual rate of air shower muons reaching these depths is not particularly high, this made it possible to simulate a very large livetime. Since the energy reconstruction of the observed event might not, for the purposes of this study, be considered trustworthy, air shower muons were simulated over a quite wide energy range, from 10 TeV to 1 PeV, with an effective simulated livetime of about 100 years at the low end of this range and a far greater effective livetime for higher energies (due to the steeply falling cosmic ray spectrum). In total, 10^6 air shower muons were simulated, but none survived the data selection process. (In order to verify that this was not a software bug which simply eliminated all events, a sample of up-going muons was also simulated; a number passed the selection criteria.) This leads to the conclusion that an event with the geometric properties and brightness of the observed event cannot be expected to be produced by air shower contamination in a sample of this size (2 years, compared to the > 100 years simulated).

After these checks, the selected data seem to be sound, so the main analysis can be performed.

6.2 E^{-2} Fit

The primary fit originally planned for this analysis was to use the hypothesis of an astrophysical neutrino flux with a spectrum of exactly E^{-2} . One dimensional projections of this result for the energy and zenith observables are shown in Figure 6.3, and the fit parameters with their profile likelihood confidence intervals are given in Table 6.2. The agreement between the data and the sum of fitted fluxes is generally good, and, as expected, the conventional atmospheric component is generally dominant.

Parameter	Fit Result	Prior
Conventional flux normalization	$0.94^{+0.04}_{-0.04}$ times the HKKMS07 flux	Must be non-negative
Prompt flux normalization	$0.85^{+1.50}_{-0.85}$ times the ERS prompt flux	Must be non-negative
Astrophysical flux normalization	$9.8^{+4}_{-3} \times 10^{-19} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$	Must be non-negative
Cosmic ray spectral index change	$-0.027^{+0.011}_{-0.011}$	Gaussian: 0 ± 0.05
Detector optical efficiency	$+18.7^{+0.9}_{-0.5}\%$	Gaussian: $+9.9\% \pm 3\%$
Kaon production normalization	$1.15^{+0.08}_{-0.07}$ times the HKKMS07 flux	Gaussian: 1 ± 0.1

Table 6.2: Best fit parameters when an E^{-2} astrophysical flux is included. The listed error ranges are 68% confidence intervals.

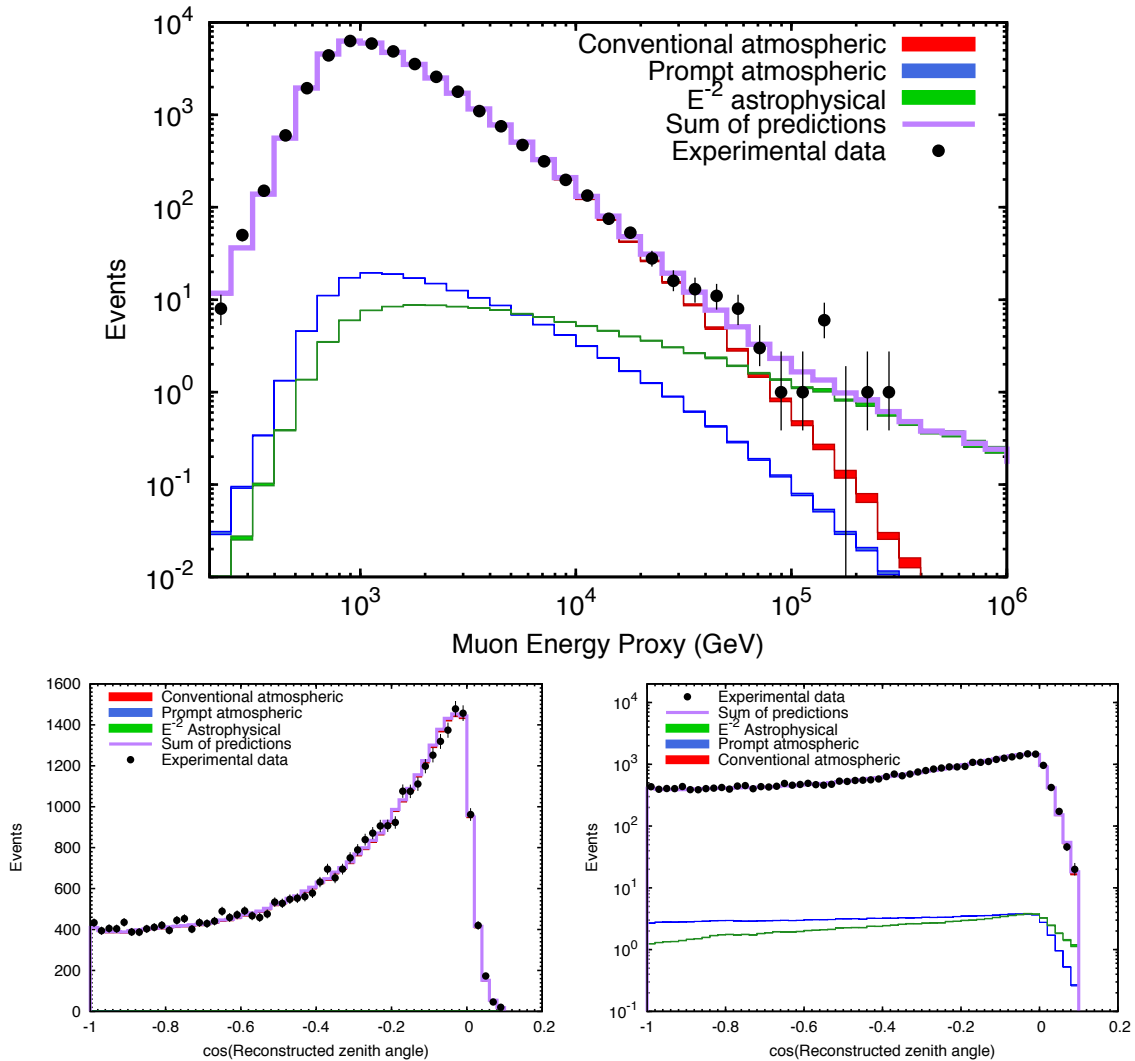


Figure 6.3: Distributions of reconstructed event energies and zenith angles compared to the best fit model for an E^{-2} power law astrophysical flux. The zenith angle projection is dominated by low energies, so it is shown in both linear and logarithmic scales, so that the subdominant flux components can be seen. Statistical 68% errors are shown on both the experimental and simulated data.

The most important result of this fit is that the astrophysical flux is fit to a non-zero normalization ($9.8_{-3}^{+4} \times 10^{-19} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$), and application of Wilks' Theorem suggests that the null hypothesis, the possibility of zero astrophysical contribution, is rejected at 3.6σ . There seems, then to be some sort of interesting signal in this flux, although an E^{-2} flux is not necessarily the best explanation. Next, the fitted contribution from astrophysical charm production of neutrinos is also non-zero, but this measurement is not distinct from zero, so we cannot conclude that this component has definitely been observed. The change to the cosmic ray spectral index is well within the prior, and unsurprisingly works to harden the atmospheric neutrino spectra somewhat, helping them to partially explain the excess of high energy events. The behavior of the DOM efficiency parameter is somewhat anomalous, as it is well outside the expected range given in the prior, and the error range computed from the data excludes the entire 1σ prior range. This portion of the fit is driven by the lowest energy events however, so it has more bearing on the measurement of the atmospheric fluxes than the astrophysical flux. Finally, the normalization for the conventional atmospheric flux is somewhat low, but this is essentially cancelled by the raised value fit for the contribution from Kaon decays to this flux, since over most of the energy range of this data sample these parameters are degenerate.

A one dimensional scan (in astrophysical flux normalization) over the likelihood space of this fit is shown in Figure 6.4. The substantial, although not total, degeneracy between the astrophysical and prompt atmospheric flux components can be seen clearly from the way that the prompt normalization falls rapidly as the astrophysical normalization is increased. The apparently similar though smaller anti-correlation of the conventional flux, however, is largely illusory, as it is more than cancelled over most of the energy range by the increasing fitted contribution of Kaon decays.

6.2.1 Separate Fits of Data-taking Periods

As a cross check, each of the two data-taking periods was also fit separately with this hypothesis. Given the substantial difference in the numbers of high energy events found in each period (Table 6.1) it is not surprising that the fits should return different strengths for the astrophysical flux, but it is important that these results should not be strongly different compared to what can be expected from statistics, as this would likely indicate either a problem with the data selection or some sort of time dependence in the flux. The results of these fits are shown in Figure 6.5 in the form of likelihood scans over the astrophysical flux normalization. Each result can be seen to be outside the other's 1σ error range (estimated using Wilks'

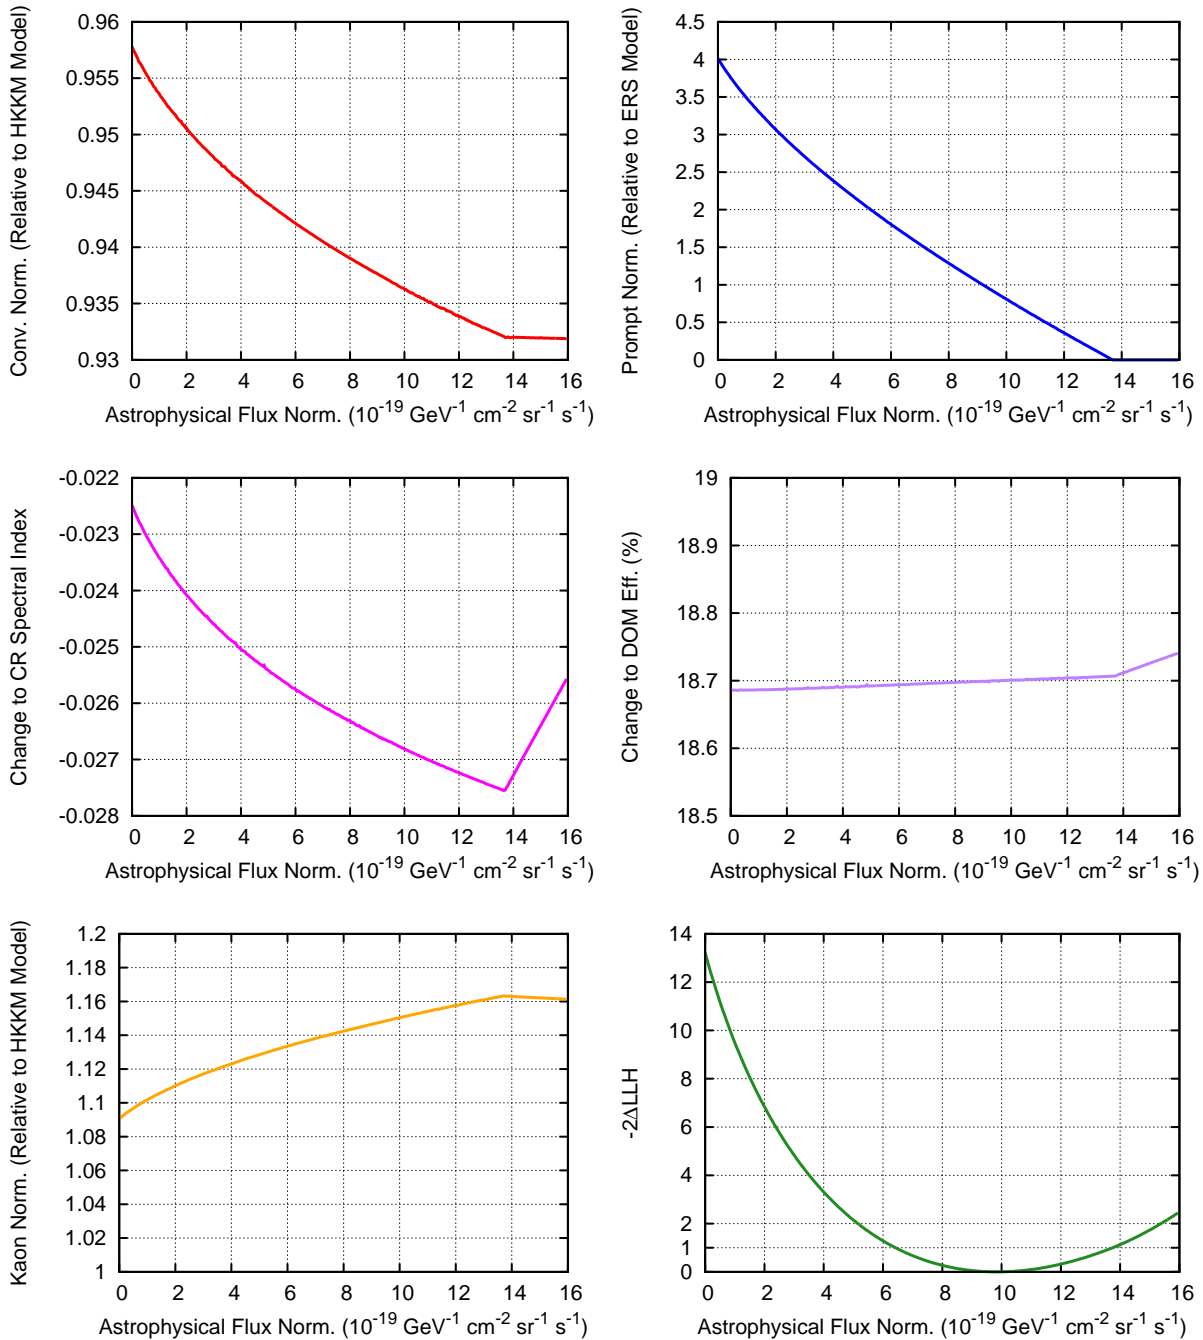


Figure 6.4: Behavior of all other fit parameters (and the fit likelihood) as a function of the normalization of an E^{-2} astrophysical flux. The parameter which experiences the most change is the prompt atmospheric flux normalization, indicating substantial degeneracy between this flux component and the astrophysical flux. A change occurs in the behavior of all parameters when the astrophysical flux level exceeds $1.37 \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$, as a prompt atmospheric flux is then unnecessary to explain the data, causing this term, which is coupled to the other nuisance parameters, to drop out of the likelihood.




Figure 6.5: Comparison of the fit likelihoods as functions of E^{-2} astrophysical flux normalization for each year of data. The best fit normalization for each case is found at the minimum point of the corresponding curve, and using the approximation of Wilks' Theorem the 68% confidence interval for each measurement is the domain in which the (doubled) likelihood difference from the minimum is less than or equal to one. While there is some tension between the separate results, it is not very strong, as their 1σ confidence intervals overlap.

Theorem), but assuming the two results are at least partially uncorrelated then the tension between the two is less than 2σ , and so is not particularly unlikely from random fluctuation.

6.2.2 Observation Domain

If a non-atmospheric signal has been observed, it is important to gain some understanding of the energy range in which it is observed. Clearly, this analysis is unable to extract meaningful information about the astrophysical flux at \sim TeV energies, where it is hidden beneath the atmospheric flux, and likewise at very high energies the flux becomes so small that this analysis will be unable to distinguish it from zero. We seek, then, a reasonable means of making this argument quantitative. A method which has been used previously for similar work is to compute the range of energies whose removal from the astrophysical flux template used to fit the data leads to a particular level of degradation of the estimated analysis sensitivity [102, 98]. This method is unsatisfactory for two reasons, however: Firstly, the sensitivity calculation is designed around the assumption of the lack of a signal (it is the median flux level computed for a dataset which actually contains no signal), which does not seem appropriate for making statements about the domain in which a non-negligible signal has been observed. Another consequence of this use of a zero-signal assumption is that since the backgrounds considered in this type of analysis (the atmospheric neutrino fluxes) decrease with

energy, if this method is applied to two analyses which have different sensitivities but observe the same flux, the analysis with the better sensitivity (i.e. the ability to distinguish smaller signals from the background) will general be found to have the larger minimum energy at which it can be claimed to observe the signal, in direct contradiction to what should be expected and desired. Secondly, the sensitivity calculation is by design independent of the observed data, instead being computed solely based on the simulation of models. This seems a wholly undesirable feature for this purpose, though, since one desires to know in what energy domain the observed data implies a given flux.

After considerable discussion with J. van Santen and L. Mohrmann, a new prescription has been devised which addresses these shortcomings. The key observation is that since likelihood differences are deemed suitable for distinguishing possible values of other parameters, it should be possible to use them to answer this question as well. We therefore make a more precise statement of the original question (‘In what energy domain is a signal observed?’): In what energy domain is the best-fit signal flux necessary to obtain a better fit to the observed data than any possible combination of the nuisance parameters? We then answer this question with the following calculation: Holding the parameters of the signal flux at their best-fit values, we progressively remove more of the lower (upper) energy tail, in true neutrino energy of the signal flux template while refitting all other parameters, until the fit likelihood has worsened by 0.5, thus obtaining the minimum (maximum) true neutrino energy to which the analysis is sensitive to the measured flux. Some of the details of this procedure such as the likelihood change threshold, are chosen arbitrarily, and other reasonable choices are possible, but there does not appear to be any argument that any of these other choices would be objectively better than the ones chosen here.

For this analysis, the procedure described above yields a sensitive energy domain for the $\Phi(E_\nu) = 9.8 \times 10^{-19} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1} \left(\frac{E_\nu}{100 \text{ TeV}}\right)^{-2}$ astrophysical signal flux of 500 TeV - 1.2 PeV. Illustrations of the fits resulting at these boundaries are shown in Figure 6.6. From the right plot is it clear that the maximum sensitive energy can be no lower, as lowering it further would substantially worsen the model’s ability to account for the highest energy observed events. The minimum energy is more subtle, as the left plot shows that introducing a hard cutoff in true neutrino energy has not led to a hard cutoff in the observable energy proxy. This is of course due to the fact that the energy proxy can only attempt to measure the energy of muons when they reach the detector, and each muon may lose any fraction of its energy while traveling to the detector. As a result, the minimum energy is fixed by the need to prevent a valley between where the atmospheric fluxes can be adjusted to fit the observed data and the highest energy events where they

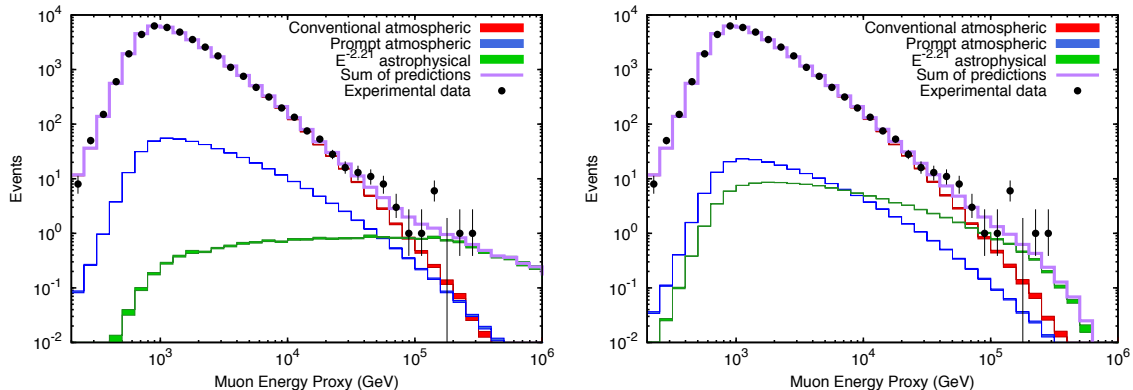


Figure 6.6: Best fits for an E^{-2} astrophysical flux with hard cutoffs below 500 TeV and above 1.2 PeV. Note that in the case of the low energy cutoff events are still present with arbitrarily low muon energies, since muons may lose an arbitrary fraction of their energy traveling from their point of production to the detector.

simply cannot. Since the normalization of the prompt atmospheric flux is considered entirely unknown in this analysis, the fit is able to use this component to compensate for a lack of astrophysical contribution up to fairly high energies, thus driving the minimum sensitive neutrino energy up to the fairly high value of 500 TeV.

6.2.3 Observation Significance

As has been previously mentioned, the significance with which the null hypothesis of a purely atmospheric origin for the observed data to be rejected in favor of the E^{-2} signal hypothesis can be estimated using Wilks' Theorem as about 3.6σ . However, it is not certain that Wilks' Theorem is a valid approximation for this problem, since relatively small portions of the otherwise large data sample are associated with the non-zero signal fit whose significance is of interest. Therefore, rather than assuming that the test statistic ($-2\Delta\text{LLH}$) is χ^2 distributed, the test statistic distribution was obtained empirically through a Monte-Carlo simulation study. The question to be answered is how often realizations of the null hypothesis fluctuate to yield test statistic values as large as or larger than the value obtained for the observed data (13.22). These fluctuations were simulated by brute force: Realizations of the experiment were drawn from the simulated events, weighted according to the atmospheric-only hypothesis, fit with both the atmospheric-only and E^{-2} astrophysical flux hypotheses, and the resulting test statistics tabulated. The results are shown in Figure 6.7: Out of 381388 trials performed 43 yielded more extreme test statistics than the experimental data, for a p-value of 1.1×10^{-4} , corresponding to a one-tailed significance 3.7σ . While it can be seen that the test




Figure 6.7: Distribution of test statistic values obtained from realizations of the atmospheric-only hypothesis, with the vertical marker placed at the test statistic for the observed data. Note that the distribution is not χ^2 with any number of degrees of freedom, thus necessitating sampling it directly.

statistic distribution so obtained has some structure not expected for any χ^2 distribution, the result is not substantially different from the one obtained using Wilks' Theorem.

6.3 Free Power law Fit

The E^{-2} astrophysical flux hypothesis fits the observed data reasonably well, but this was always a simplistic model. The easiest extension to this model is to allow the spectral index of the hypothetical signal flux to be a free parameter as well, and this additional freedom corresponds reasonably to known open questions, such as the fact that neither the spectral index of the cosmic ray flux at Earth at high energies nor the change in that spectral index during the diffusion of the cosmic rays is exactly known. Accordingly, the fit from the previous section was performed again, allowing the index of the astrophysical flux to vary freely, yielding the results in Table 6.3, which are also shown in the energy proxy projection in Figure 6.8. The first change noticeable from the previous result is that the fit has used the added freedom to make a somewhat softer astrophysical flux, with an index of -2.21, although with a one dimensional error range which is still compatible with -2. The normalization of the astrophysical flux has also changed somewhat, but this can be understood as a correction mostly tied to adjusting the spectral index; the energy about which the spectrum is allowed to pivot is 100 TeV, but as discussed in Section 6.2.2 the energies which contribute to the astrophysical component are probably somewhat higher, so softening the flux requires

Parameter	Fit Result	Prior
Conventional flux normalization	$0.94^{+0.04}_{-0.04}$ times the HKKMS07 flux	Must be non-negative
Prompt flux normalization	$0^{+1.05}$ times the ERS prompt flux	Must be non-negative
Astrophysical flux normalization	$1.6^{+0.63}_{-0.8} \times 10^{-18} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$	Must be non-negative
Astrophysical flux index	$2.21^{+0.21}_{-0.23}$	None
Cosmic ray spectral index change	$-0.026^{+0.01}_{-0.008}$	Gaussian: 0 ± 0.05
Detector optical efficiency	$+18.7^{+0.55}_{-0.5} \%$	Gaussian: $+9.9\% \pm 3\%$
Kaon production normalization	$1.15^{+0.08}_{-0.07}$ times the HKKMS07 flux	Gaussian: 1 ± 0.1

Table 6.3: Best fit parameters when an astrophysical flux with the form of an arbitrary power law is included. The listed error ranges are 68% confidence intervals.

raising its normalization to pass through the same data points. The only other fit parameter to undergo an interesting change is the normalization for the prompt atmospheric flux, which has dropped to zero. This is of course now possible because the softer astrophysical flux takes over more smoothly from the conventional atmospheric component, so the prompt component is not needed to fill in the transition.

A different method of visualizing this result is shown in Figure 6.9, which shows a scan over the likelihood space in both of the free parameters of the astrophysical flux. Both plots are overlaid with confidence contours constructed using Wilks' Theorem. The plot of likelihood values clearly shows the minor degeneracy between




Figure 6.8: Distribution of reconstructed event energies compared to the best fit model for an $E^{-\gamma}$ power law astrophysical flux. The fitted zenith angle distribution is not qualitatively different from that shown in Figure 6.3 since it is entirely dominated by the conventional atmospheric component.

the two astrophysical flux parameters, as well as the fact that while the fit with a free power law index is better than the fit with the index fixed to -2, the latter is not excluded with any great significance. The plot of fitted prompt normalization shows how the prompt flux component becomes unnecessary for any sufficiently strong astrophysical flux.

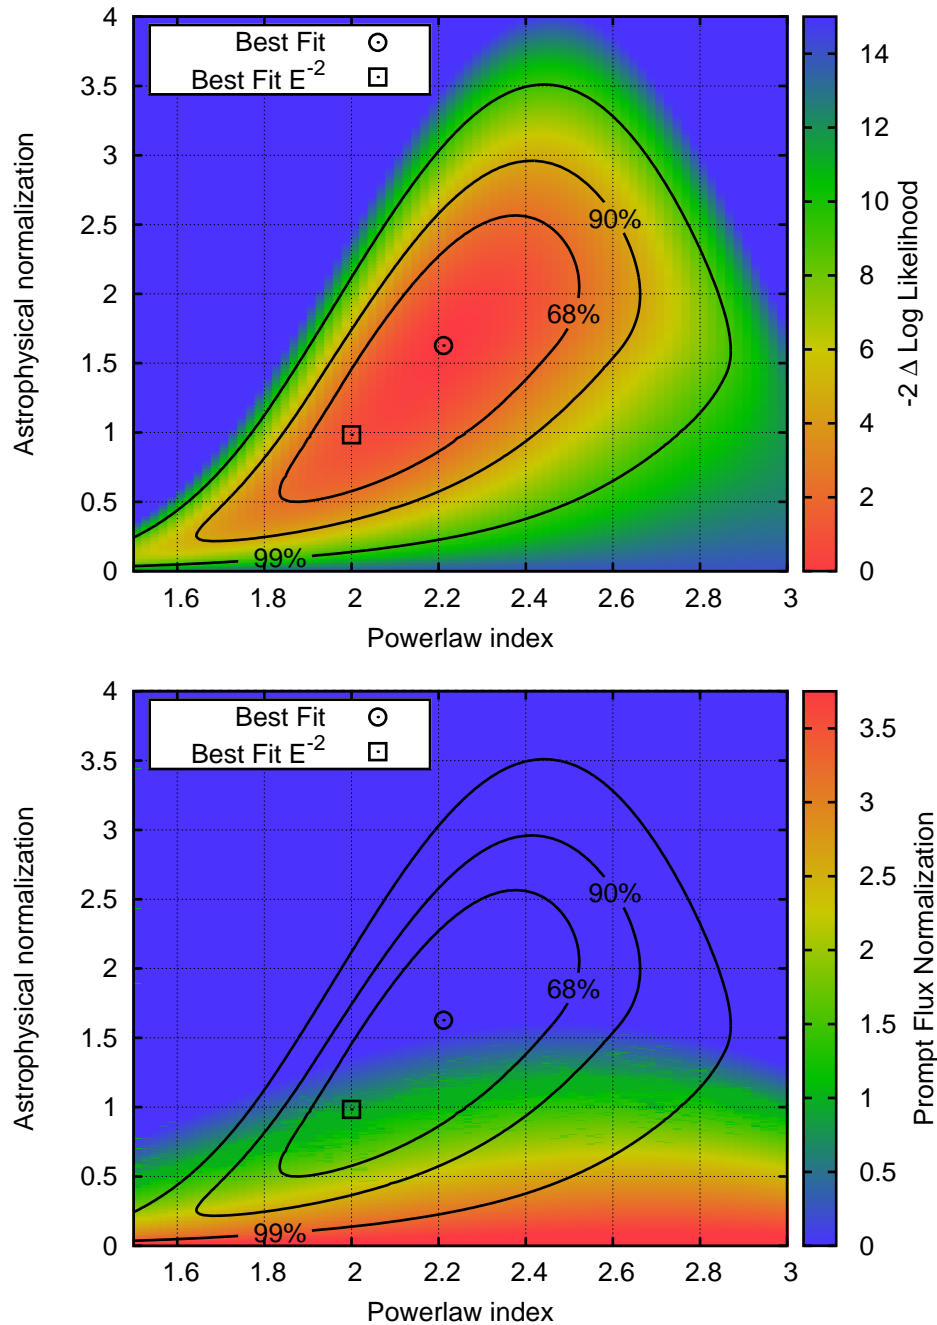


Figure 6.9: Comparisons of profile likelihood and prompt atmospheric normalization as functions of possible astrophysical power law fluxes (spectral indices and normalizations in units of $10^{-18} \text{GeV}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{s}^{-1}$). Confidence regions are computed using Wilk's Theorem and indicated by contours. The best fit E^{-2} flux is within the 68% region, but is clearly not the best possible fit. The approximate anti-correlation of the prompt normalization to the astrophysical normalization is shown in the second plot; for many allowed astrophysical flux models no contribution from atmospheric prompt flux is required at all.

6.4 Model Fits

Besides fitting power law hypotheses for the astrophysical flux, we can also directly use published models. We make the simplifying assumption that, like the prompt atmospheric flux the main uncertainty about these models is their normalization, treating the shapes as entirely fixed. As a survey of representative models we use the same ones discussed in Section 2.2 [42, 43, 44]. Table 6.4 summarizes these fits in comparison to the power law fits from the previous sections. For each model the best-fit normalization is listed, along with the largest normalization not excluded by the data at 90% confidence (using Wilks' Theorem), and finally the difference in log likelihood between the best fit for this model and the global best fit (the free-index, unbroken power law). From these fits it appears that none of these models is particularly close to having the correct normalization to match the observed data, although the Loeb & Waxman Starburst model comes the closest, and it is worth noting that this is well inside the large uncertainty claimed by the authors of this model. None of these models can be rejected based on shape discrimination with any substantial significance, although the Waxman Bahcall GRB is disfavored at about 1.3σ and the Loeb Waxman Starburst at almost 1σ .

One additional point of interest for the GRB model is that strong limits have already been placed on the muon neutrino flux from GRBs [45]. The observation of a flux by this study, combined with those limits, constrains the fraction of the flux which can be produced by the visible GRBs to be no more than a few percent for typical assumptions of GRB fireball properties.

Model	Best-fit Normalization	Upper Limit	Likelihood Relative to Best
Stecker AGN Core	0.17	0.30	0.26
Loeb Waxman Starburst	2.49	3.55	0.47
Waxman Bahcall GRB	6.70	11.39	0.81
E^{-2} Power law	-	-	0.41
Arbitrary Power law	-	-	0

Table 6.4: Best fit normalizations resulting from fitting an astrophysical flux given by various representative theoretical model predictions. A fit normalization smaller than one indicates a model which predicts a flux higher than that observed experimentally, while a fitted normalization greater than one indicates a prediction too low to explain all of the flux observed in the data. Upper limits are at 90% confidence.

6.5 Signal Probability of Individual Events

It is of course impossible based on the information in this analysis to determine with total certainty whether any given event is astrophysical in origin or not. However, with the assumption of a flux, it is possible to make probabilistic statements about events simply by taking the ratio of the prediction for each flux component to the total predicted flux for events similar to the one in question. Since the free index power law fit of Section 6.3 is the best fit obtained, we will use it as an example. Figure 6.11 shows the predicted rate of events in this data sample from that best-fit flux, both directly and as a cumulative distribution (from high to low energy). This shows that for this model the astrophysical component of the flux becomes dominant at energies above ~ 200 TeV, and that, in aggregate, any neutrino with energy greater than ~ 100 TeV is equally likely to be atmospheric as astrophysical. While interesting, this is of limited applicability, since the true neutrino energies of individual events are also not known (however, see Table A.2 for a similar probabilistic treatment of that property). It is equally straightforward to do the same calculation in the space of an observable, such as the energy proxy, and Figure 6.10 shows the result. As the point of equal signal and background probability occurs at energy proxy values of around 40 TeV, this means that if this flux model is correct then more than half of the events listed in Table 6.1 should originate from astrophysical sources.




Figure 6.10: The fraction of the total flux fitted in Section 6.3 arising from each component, as a function of the event energy proxy. (Note that since that fit obtained zero prompt atmospheric component, that component does not appear here.)

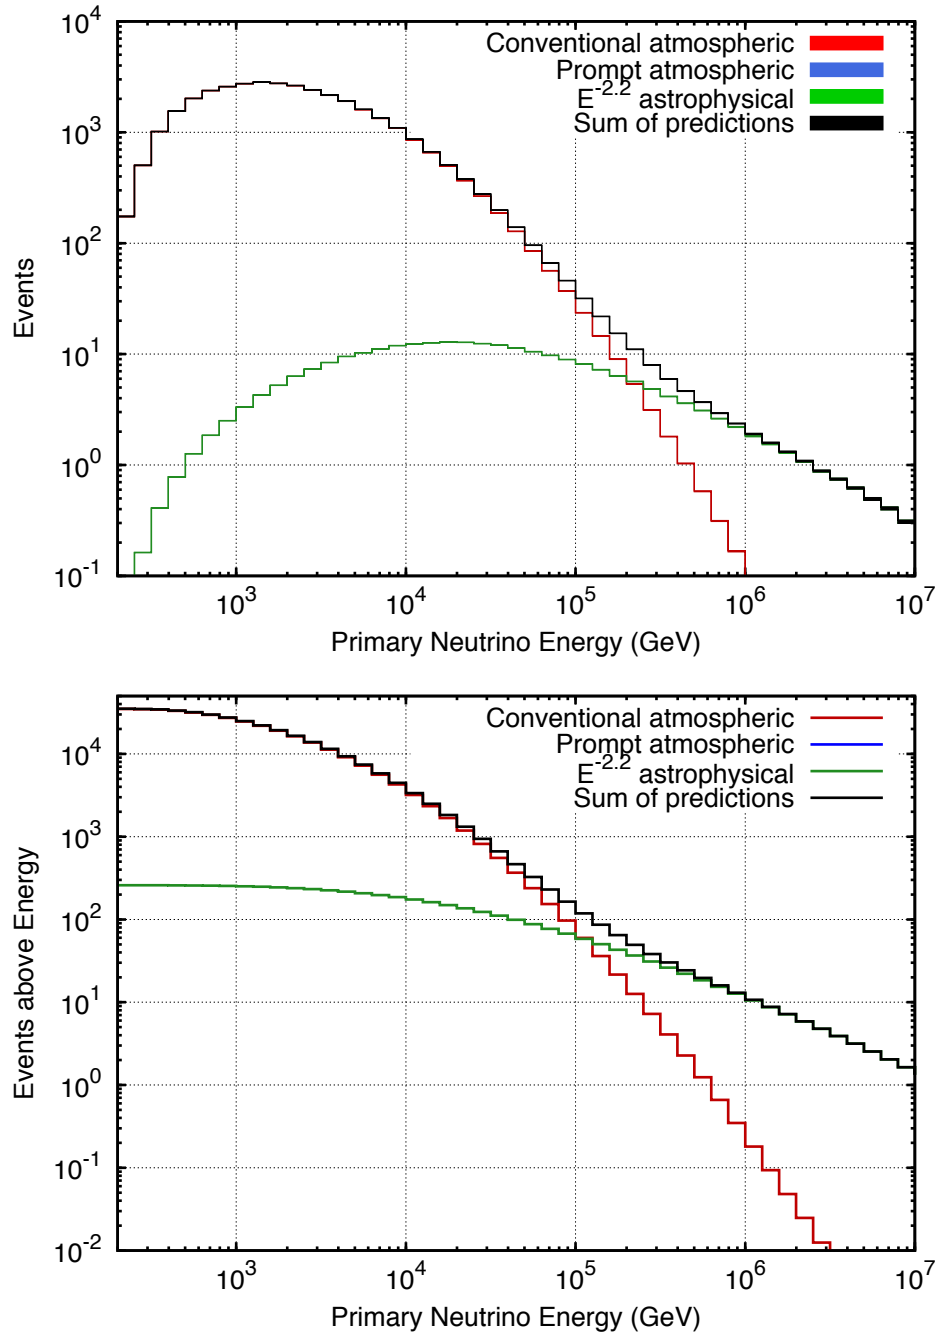


Figure 6.11: The predicted rate of events to be observed in this analysis as a function of primary neutrino energy, assuming the best-fit flux of Section 6.3. The lower panel repeats the same information in cumulative form, from high to low energy.

Chapter 7

Conclusions

7.1 Comparison to Other Recent Results

Several different analyses for diffuse neutrino fluxes have now been performed on the data collected by the first few years of the complete IceCube detector. This analysis was focused only on muon neutrinos (and only those undergoing charged-current interactions), and made the trade-off of collecting more data by accepting events with neutrino interaction vertices outside the instrumented volume of the detector, and as a result could only use the angular range where the detector was protected from air shower muons by considerable overburden. In contrast, the other recent analyses have been designed to accept only events with interaction vertices inside the detector and no visibly entering particles; meaning that the entering particles must be neutrinos, and so may be termed ‘starting event’ analyses. These analyses then have the advantages that they do not rely on material shielding (using instead active shielding from the outer portions of the detector itself) and thus can view the entire sky, and can also accept all flavors of neutrinos, at the cost that they have available only a fraction of the detector’s full volume to actually collect events. These two approaches are quite complementary, so it is natural to ask whether compatible results are obtained.

For simplicity we will consider two representatives from the collection of starting event analyses. The first of these [103] was designed for maximal acceptance at high energies ($\gtrsim 100$ TeV) while essentially ignoring lower energy events. Three years of IceCube data were used, from the 2010 data-taking period through the 2012 data-taking period, and the resulting best-fit power law astrophysical flux was found to be $1.5 \pm 0.3 \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1} \left(\frac{E_\nu}{100 \text{ TeV}} \right)^{-2.3 \pm 0.3}$. The second starting event analysis [104] was designed to extend the first to much lower energies, ~ 1 TeV, by using an adaptive technique to veto visibly incoming particles. The same two years of data (2010 and 2011) were used as for this work, and the resulting best-fit power law flux was found to be $2.23_{-0.36}^{+0.4} \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1} \left(\frac{E_\nu}{100 \text{ TeV}} \right)^{-2.49 \pm 0.15}$.




Figure 7.1: Same as Figure 6.9 but with the results of [103](HESE) and [104](MESE) shown. All three analyses' results are compatible at about the 1σ level, although they are not identical.

These two results are compared in Figure 7.1 with the free power law result of this work ($1.63_{-0.8}^{+0.62} \times 10^{-18} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ sr}^{-1} \text{ s}^{-1} \left(\frac{E_\nu}{100\text{TeV}}\right)^{-2.21 \pm 0.25}$). All three results are within each others' 1σ uncertainty contours, but it is noticeable that the starting event results favor softer power law indices (a tendency followed by other starting event results not shown here [105, 106, 107], as well). Possible reasons include a change in the astrophysical flux somewhere below ~ 100 TeV (where this analysis is unable to observe the astrophysical flux beneath the atmospheric background), or that the flux differs in the northern and southern hemispheres (as this analysis is blind to the latter but the starting event analyses are not). Unfortunately the quantities of data obtained so far are too small to effectively investigate these or other hypotheses. Nonetheless, the similarity of the results from analyses with sensitivity to different energy ranges, angular ranges, and different neutrino flavors indicates that the toy model of an isotropic flux with equal flavor composition over the entire sky is a fair approximation.

Finally, another visualization of this result compared to other measurements and theoretical models is shown in Figure 7.2.




Figure 7.2: Comparison of the result of this work (from Section 6.3), shown in green to other recent results and various flux models which have been previously discussed. The agreement with the 3 year starting event result [103] can again be seen to be quite good. Note, however, that the energy range depicted for the starting event result is simply the range of observed deposited energies, which both underestimates the true neutrino energies and probably somewhat overstates the sensitivity at low energies. None of the available models is a particularly good match to the data, although the Starburst model [43] comes close, and it should be noted that this model has large uncertainties.

7.2 Outlook for Distinguishing Flux Models

In Section 6.4 some different proposed neutrino flux models were fit to the observed data, but when allowing the normalizations of the model predictions to vary freely, there proved to be little power to distinguish which functional form, if any, describes the data better than the others. How much data would be required to actually make such a choice between models? One way to attempt to answer this question is to perform a series of data-challenges, in which a possible signal realization for one model is drawn from simulated data, and is fit with both that and another model. Doing this repeatedly and for different simulated detector livetimes can then give an estimate of whether the data which could be collected in that livetime would be likely to distinguish between the two models being tested (were one of them actually correct).

The results of such a calculation are shown in Figure 7.3. In this case the Loeb-Waxman starburst model (for the only the $E^{-2.3}$ powerlaw case favored by those authors, neglecting the broad range of possibilities) was treated as the true model, with the best-fit normalization obtained in Section 6.4. This model was tested against the generic E^{-2} power law model by fitting both models to each realization and calculating the difference in log likelihoods between the two. Since the simulated data is drawn according to the starburst model it can be expected to have the better likelihood, and on average this is true, but for smaller livetimes statistical fluctuations sometimes cause the likelihoods to reverse, so that the incorrect model gives the better fit. For the case of two years of livetime, as in the data analyzed in this work, the median likelihood difference between the true and incorrect model is less than 1σ , and the wrong model appears better in about 28% of trials, so the difference between these models for the experimentally observed data (the E^{-2} model is favored at $\sim 0.33\sigma$) is still consistent with the starburst model being correct. The expected likelihood difference between the two models does rise steadily as more data is added (where it is assumed that each additional year added is equivalent to the 2011 data-taking period), but does so slowly, with a median significance of 2σ requiring 4 years of data and a median significance of 3σ requiring 7 years.

While this test only compared two possible hypotheses, it appears that other relevant models have broadly similar prospects. It seems likely, then, that obtaining useful results for astrophysics in the next several years may not be practical by pursuing the extension of this analysis alone, but may require combination with other techniques (such as the starting event analyses discussed in the previous section), still larger detectors, or both.




Figure 7.3: Projected significance for distinguishing a flux which follows the $E^{-2.3}$ case of the starburst model of [43] from an unbroken E^{-2} power law. Negative significances correspond to the (incorrect) E^{-2} hypotheses fitting the data better than the correct hypothesis. The line shows the median result from a set of Monte-Carlo trials while the shaded region shows the central 68% range of outcomes. The point plotted at a livetime of two years is the actual result from this work, which indicates a better fit for the E^{-2} hypothesis but is nonetheless a plausible outcome even if the starburst model is correct.

7.3 Implications for Source Identification

In order to use neutrinos to determine the origin of cosmic rays, it is highly desirable not only to study the integrated diffuse flux (which adds new information but is not dissimilar from what can already be studied about the cosmic rays themselves), but to identify sources, ideally matching them with objects known through other types of observations. Naturally, a number of searches for point sources of neutrinos have been performed using data from IceCube, but to date no indications of individual sources have been found. Additionally, no correlation has been observed with the galactic plane, suggesting that at least some component of the flux is produced by extragalactic sources. Since the diffuse flux has now been observed by several analyses we can attempt to infer something about the sources, in terms of their number and the fluxes they produce.

The following calculation, made by J. Feintzeig [108], combines the information about the observed flux from this work with the limits on individual source fluxes established by the point source searches. The approximate inputs are shown in Figure 7.4, which compares the diffuse flux (with its substantial uncertainty) to the maximum flux from a point source which could escape detection in the most recent




Figure 7.4: From [108], the sensitivity of recent searches with IceCube data for point sources of neutrinos for various possible power law indices, compared to the general level at which a diffuse neutrino flux has been observed. The fact that the point source searches are sensitive at a level below the observed flux yet no sources have been identified places a limit on the minimum number of sources which can be present. Note that this figure is plotted in terms of $\sin(\text{declination angle})$, which is equivalent to $-\cos(\text{zenith angle})$ for IceCube.

analysis as a function of source declination. The point source analyses have been performed over the entire sky, but have different behavior in the northern sky (positive $\sin(\delta)$) where atmospheric neutrinos are the only background, than in the southern sky (negative $\sin(\delta)$) where air shower muons are a stronger background than atmospheric neutrinos. The minimum number of neutrino sources to produce the observed diffuse flux can then be conservatively estimated in the following way: Under the assumption that the sources are distributed isotropically around the Earth, perform Monte Carlo trials of sampling random points on the sky and at each random point inserting the maximum point source flux which could have evaded detection by the point source analyses. Continue doing so until the total flux thus generated reaches the level of the observed diffuse flux. This then forms one possible realization of a minimal set of sources, and generating many realizations samples the distribution of possible outcomes.

This calculation was performed for each power law spectral index allowed within the 68% contour region of Figure 7.1, in each case for the best fit normalization at that index, yielding results which are shown in Figure 7.5. This shows that harder fluxes can be explained by smaller numbers of sources, since they can produce the number of events at the high energies where this analysis is sensitive without producing particularly many events at lower energies; softer fluxes require more sources since they produce more low




Figure 7.5: From [108], the minimum number of point-like neutrino sources which could produce the observed diffuse flux while evading detection by dedicated point source searches. For the best-fit flux found by this analysis (with spectral index ~ 2.2) the most likely minimum number of sources is ~ 75 .

energy events for the same number of high energy events, and if too many of these were to originate from any single source the point source analysis would have detected the angular clustering. For the best-fit flux obtained by this analysis (spectral index ~ 2.2) the obtained minimum number of sources is ~ 75 , and the data are consistent with minimum numbers of sources ranging from ~ 20 to ~ 150 .

One possible way out of the difficulty of identifying sources from the total angular distribution is to increase follow-up of individual events: roughly half of the events shown in Figure 7.6 are more likely to be astrophysical than atmospheric (assuming the flux results of this study), and it may be profitable to investigate their directions of origin with other complementary observation methods, such as traditional electromagnetic astronomy, gamma-ray telescopes, or correlation with high energy cosmic rays (whose deflection is small). Several such efforts are now underway, but they will be limited both by the non-trivial remaining atmospheric neutrino background and the angular resolution of these events, which at $\sim 0.5^\circ$ is unimpressive in most astronomical circles.

That such numbers of sources are predicted is both something of a blessing and a curse: There may be many of them to study, perhaps belonging to multiple classes or types, so if they can be identified there may be a great deal of astrophysical knowledge to be extracted. On the other hand, if each source is too dim, it may prove very difficult to isolate any of them, and it may not be possible to collect enough neutrinos

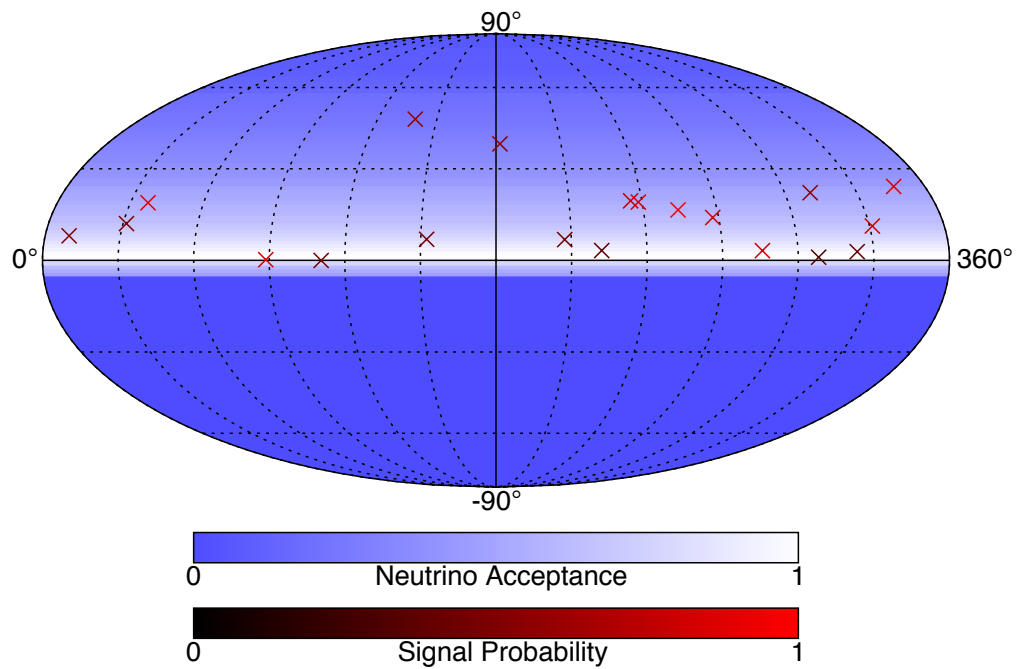


Figure 7.6: The distribution of the highest energy events observed by this analysis on the sky in equatorial coordinates. The blue shading represents the relative acceptance of the detector for neutrinos with energies above 150 TeV, and each red cross indicates the reconstructed arrival direction of one of the events in Table 6.1 with its color indicating the probability of that event arising from the astrophysical component of the best fit power law spectrum of Section 6.3 rather than the atmospheric component (as in Figure 6.10). In particular, it is clear that the distribution of the data in declination is primarily derived from the detector acceptance.

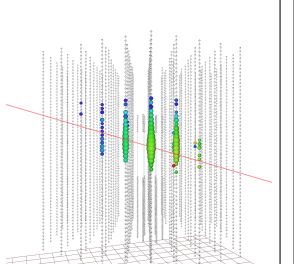
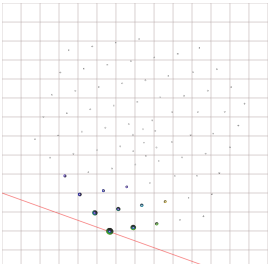
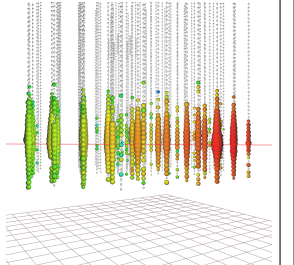
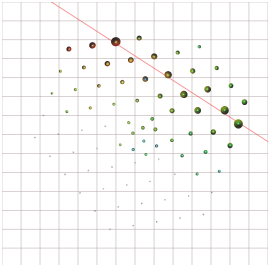
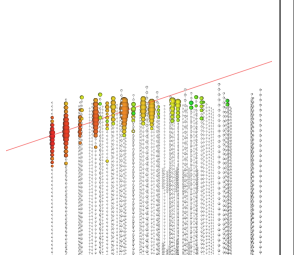
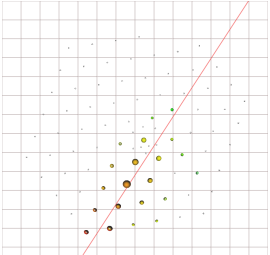
from any one of them to perform the most interesting analyses, such as spectroscopy or measurement of time variation. It is, however, too early to tell which of these scenarios is closer to the truth, and with new data accumulating from IceCube and the similar KM3Net now under construction there is substantial reason to be optimistic.

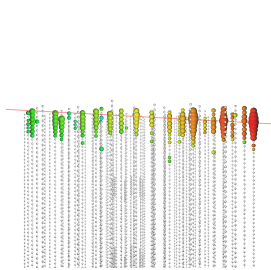
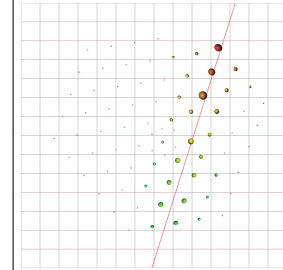
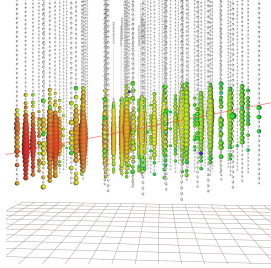
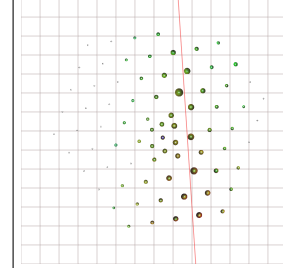
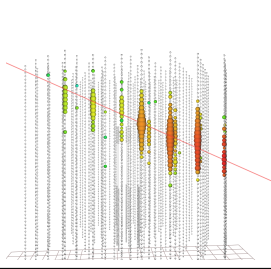
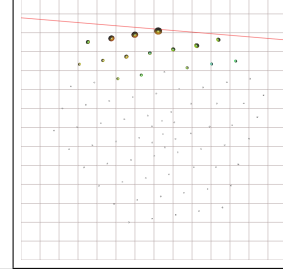
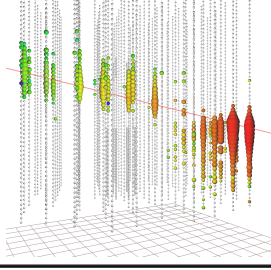
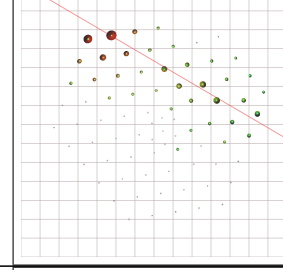
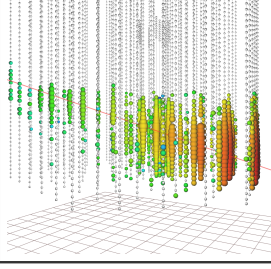
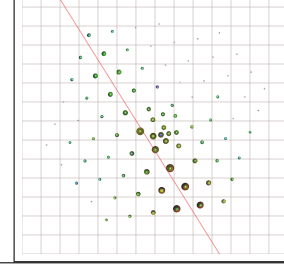
APPENDIX

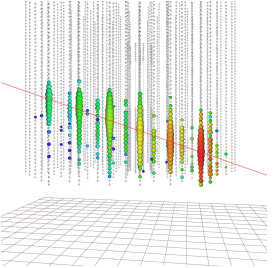
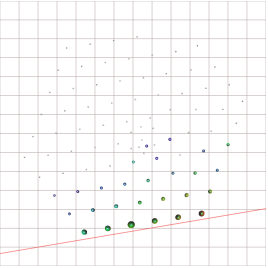
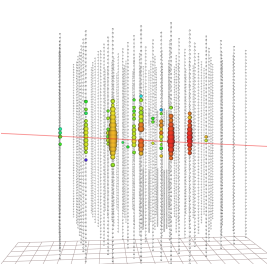
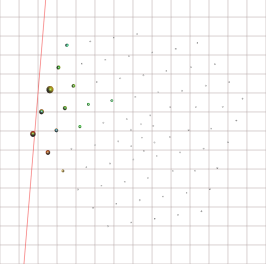
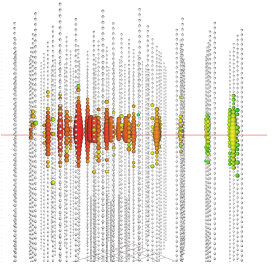
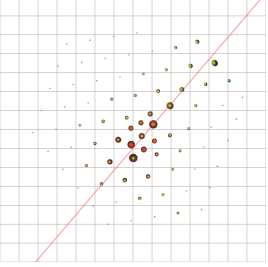
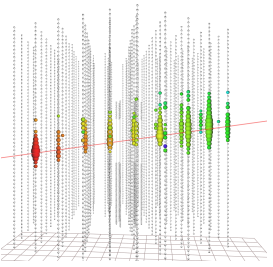
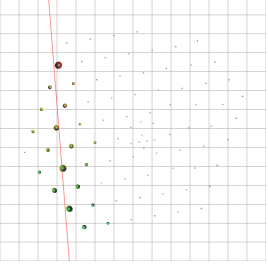
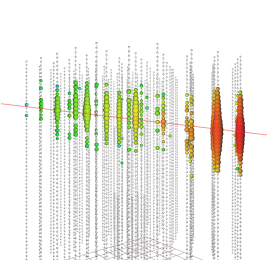
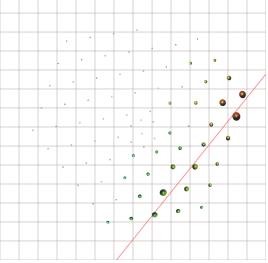
Details of Observed Events

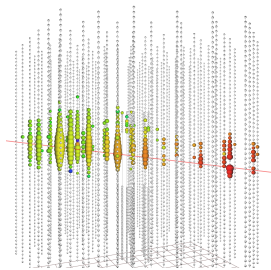
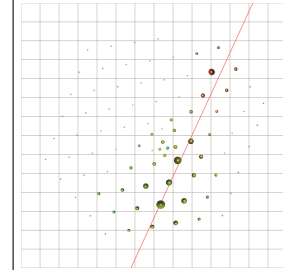
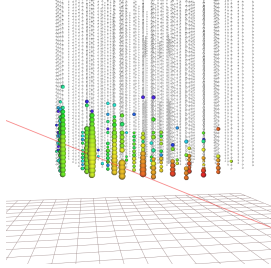
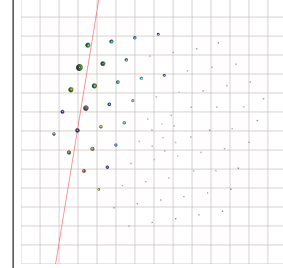
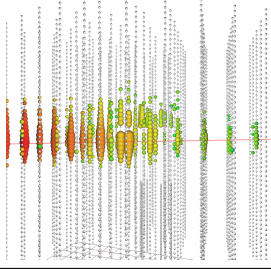
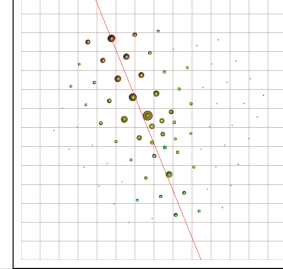
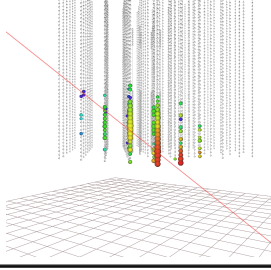
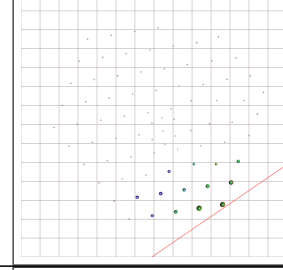
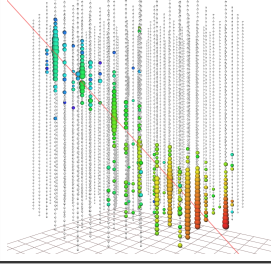
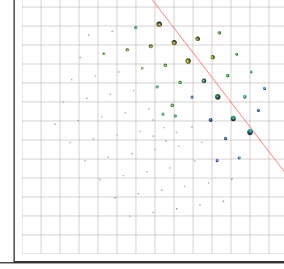
Table A.1 lists the 20 events which were found in the final data sample, which are those with energy proxy values greater than 50 TeV. The grid shown in the top view of the events has a spacing of 100 meters.

Table A.1: Images of observed events with highest estimated energies

Name	Side view	Top View	Energy Proxy Zenith Angle (Comments)
Dr. Heinrich Faust			290.1 TeV 106.3°
Dr. Hari Seldon			199.8 TeV 90.2°
Dr. Susan Calvin			160.2 TeV 108.7°

<p>Captain Nemo</p>			<p>147.4 TeV 93.1°</p>
<p>Dr. Richard Seaton</p>			<p>141.9 TeV 101.1°</p>
<p>Dr. Henry Jekyll</p>			<p>133.2 TeV 114.1°</p>
<p>Dr. Emmet Brown</p>			<p>132.1 TeV 103.8°</p>
<p>Dr. Henry Walton Jones, Jr.</p>			<p>131.8 TeV 108.9°</p>

Dr. Victor Frankenstein			109.3 TeV 106.3°
Dr. Strangelove			66.7 TeV 92.8°
Dr. Strangepork			65.1 TeV 90.0° (This is a clear starting event, also found by [103] and [104])
Impey Barbicane			64.0 TeV 97.9°
Professor Joseph Cavor			60.9 TeV 96.7°

<p>Professor Abraham Van Helsing</p>			<p>60.2 TeV 96.8°</p>
<p>Dr. Moreau</p>			<p>55.2 TeV 112.0°</p>
<p>Buckaroo Banzai</p>			<p>53.4 TeV 91.0°</p>
<p>Dr. David Bowman</p>			<p>52.0 TeV 128.6°</p>
<p>Angus MacGyver</p>			<p>51.1 TeV 137.6°</p>

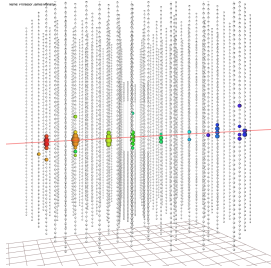
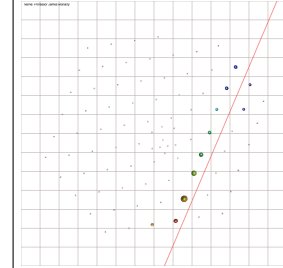
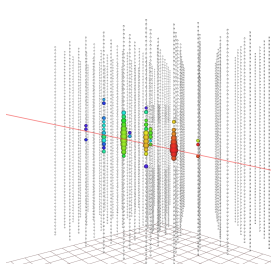
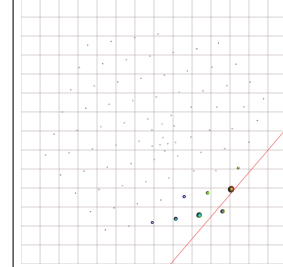
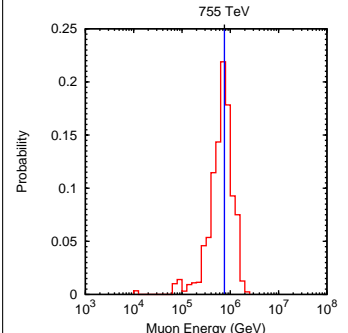
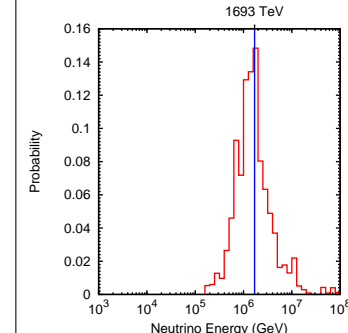
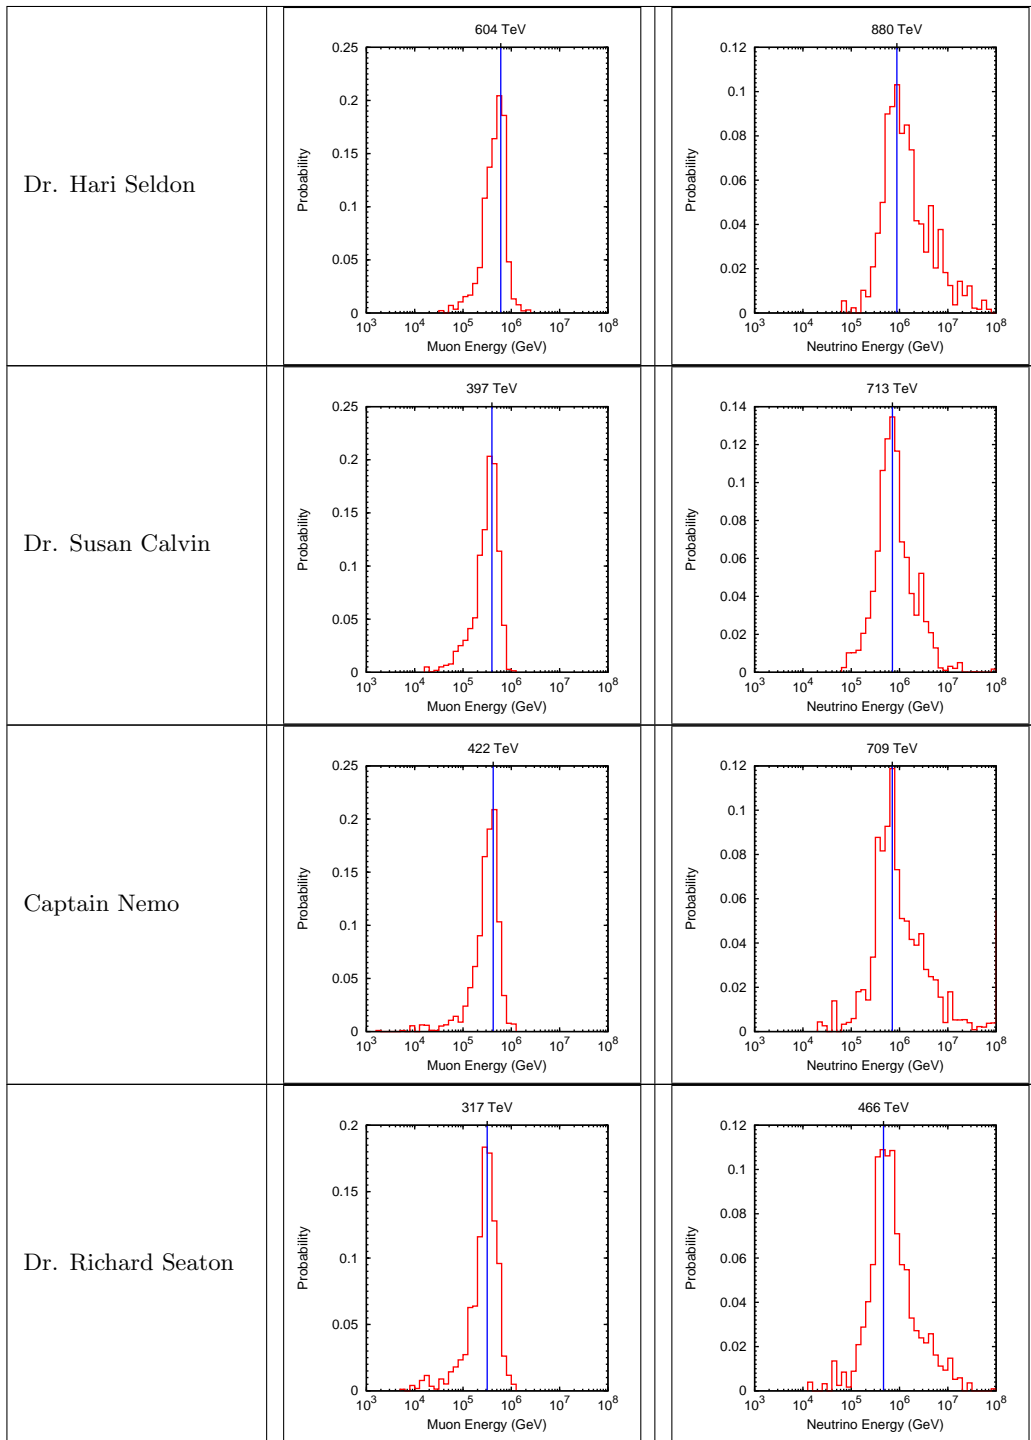
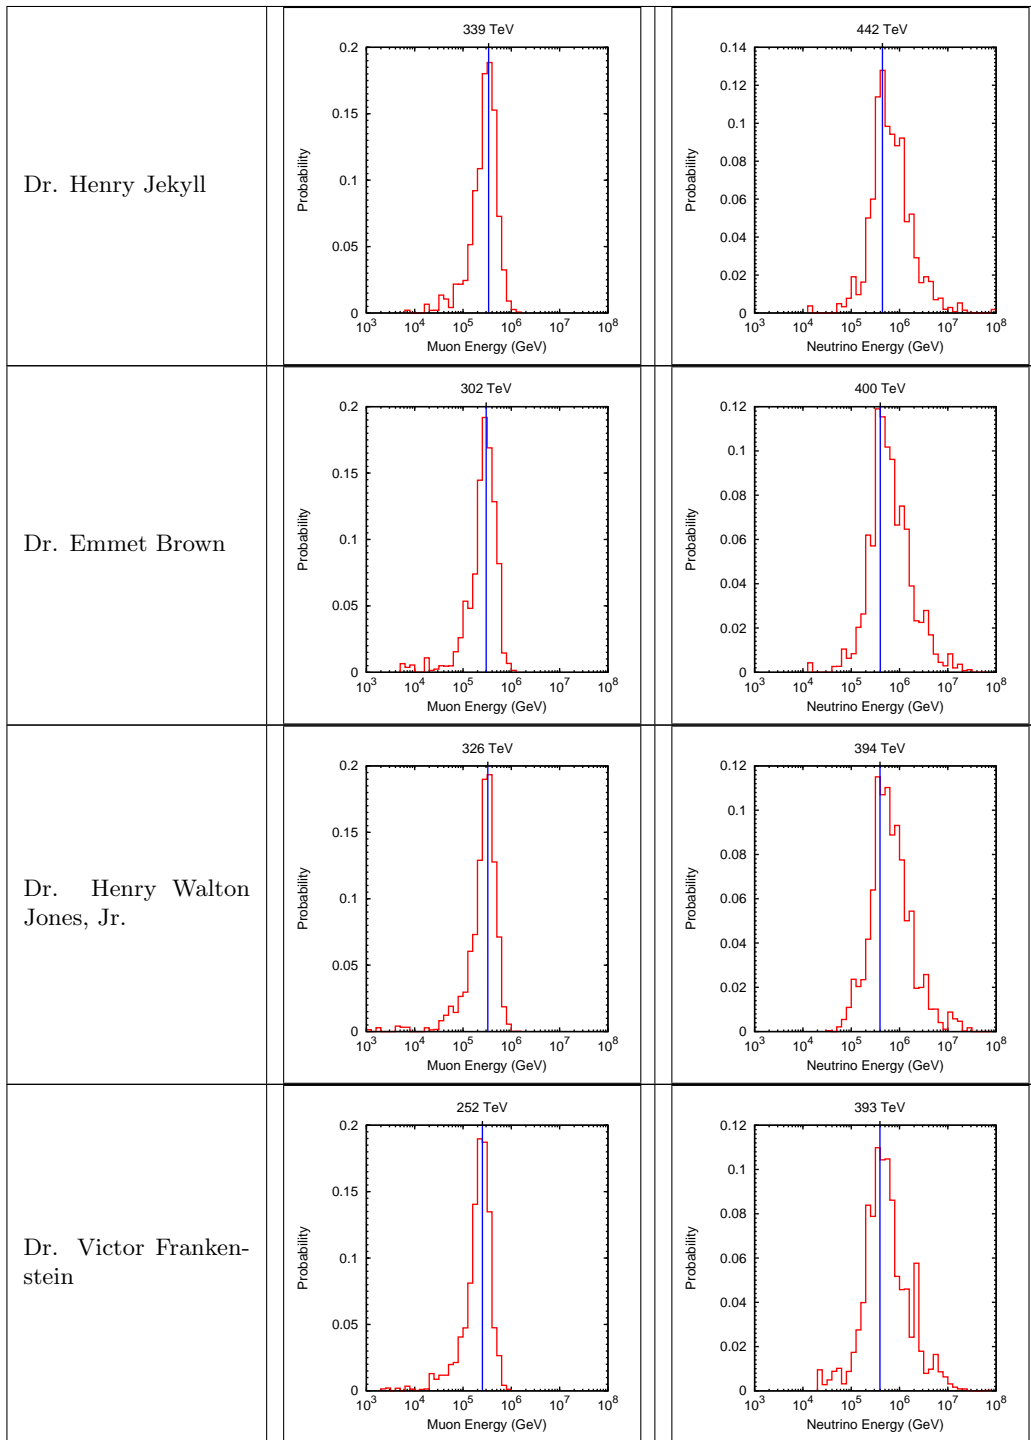
<p>Professor James Moriarty</p>			<p>50.7 TeV 93.2° (This event looks quite dim due to being entirely in the dust layer.)</p>
<p>Dr. Giacomo Rappacini</p>			<p>50.3 TeV 101.9°</p>

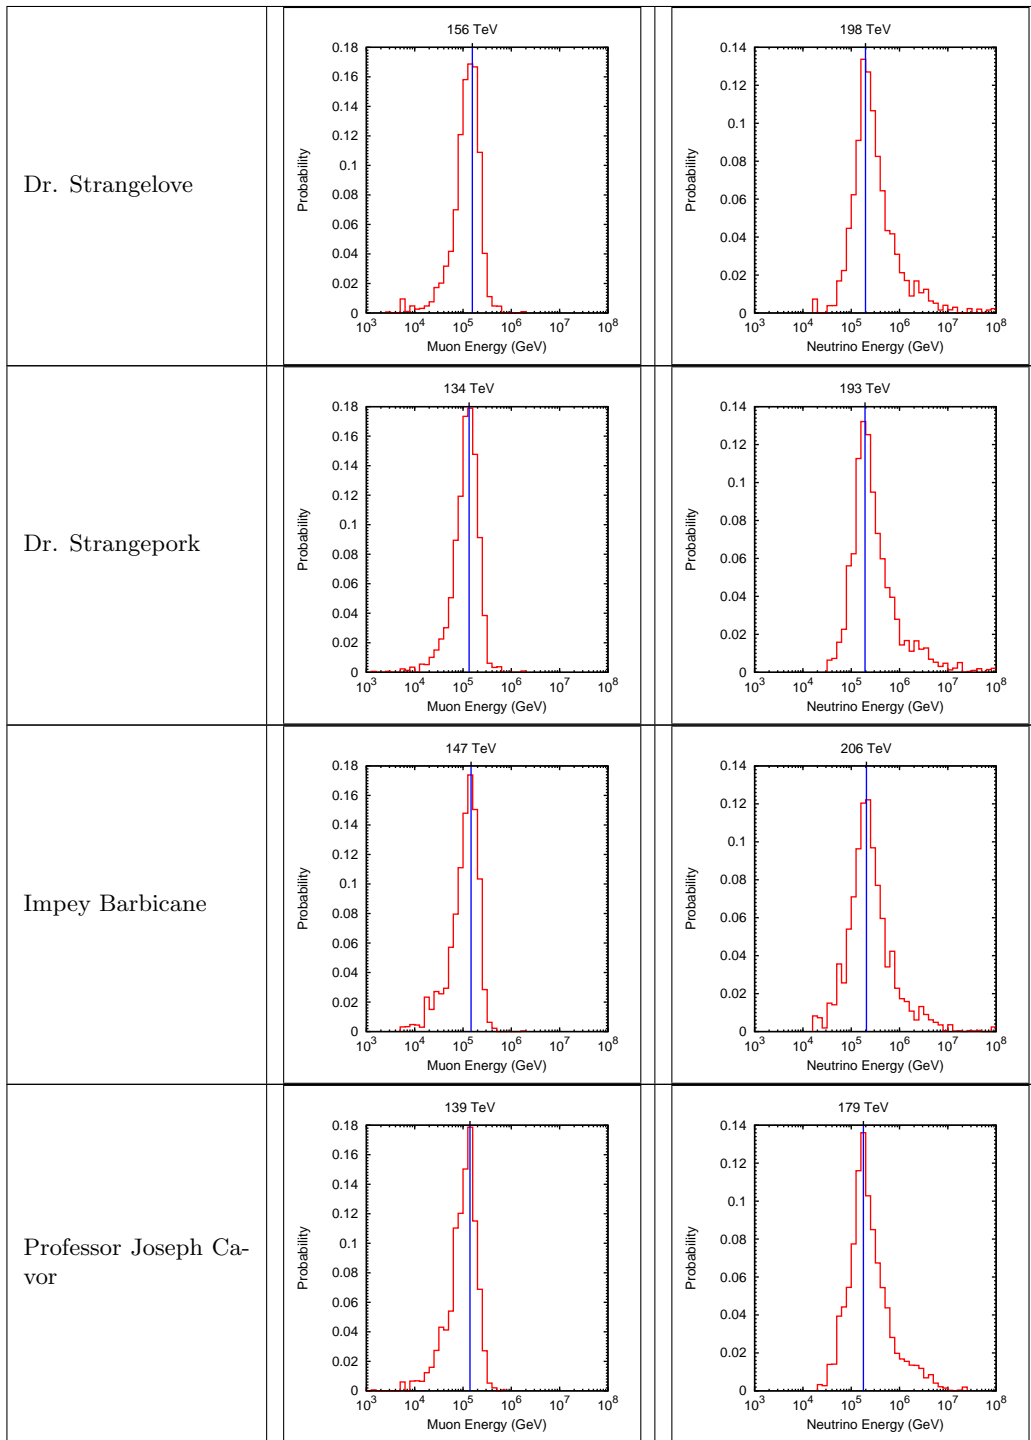
Table A.2 lists the same 20 events with their true muon and neutrino energy PDFs computed assuming the spectrum fitted in Section 6.3. Each PDF is constructed from the true properties of all simulated events which have energy proxies within 5% of the observed event’s value, and reconstructed zenith angles within 5° for events with energy proxies less than 100 TeV and 10° for those with larger energy proxies. Unfortunately, due to limited simulation statistics, many of these distributions still contain large fluctuations, particularly those with the highest energies. As a result, the estimated energies are rather imprecise and should be treated as ‘ballpark’ numbers only.

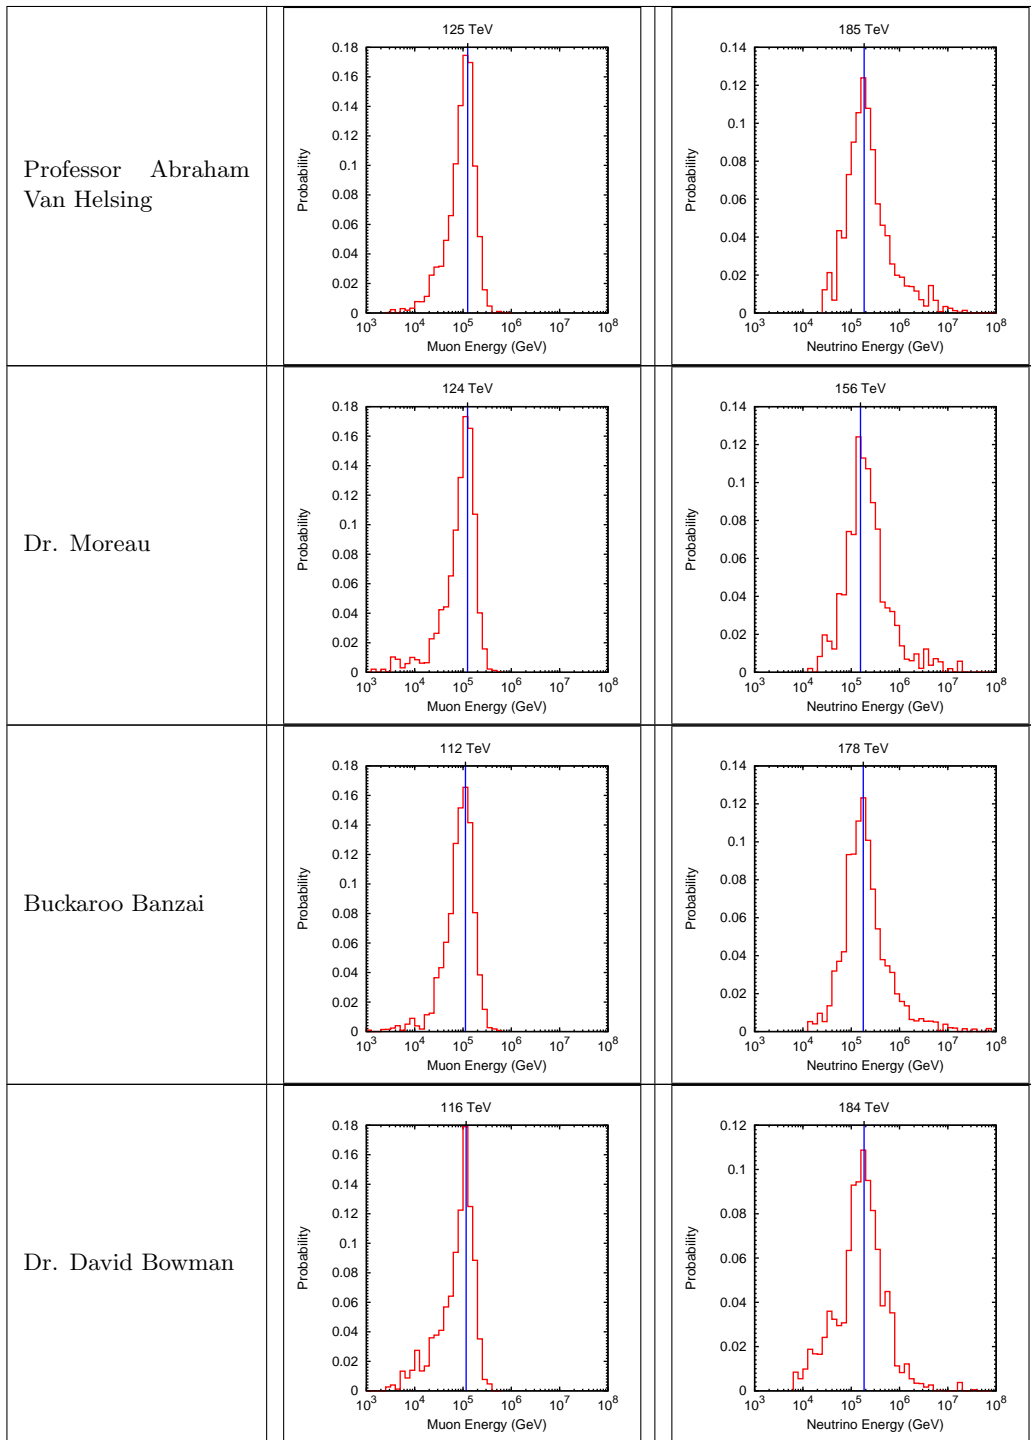
Table A.2: Inferred event energies

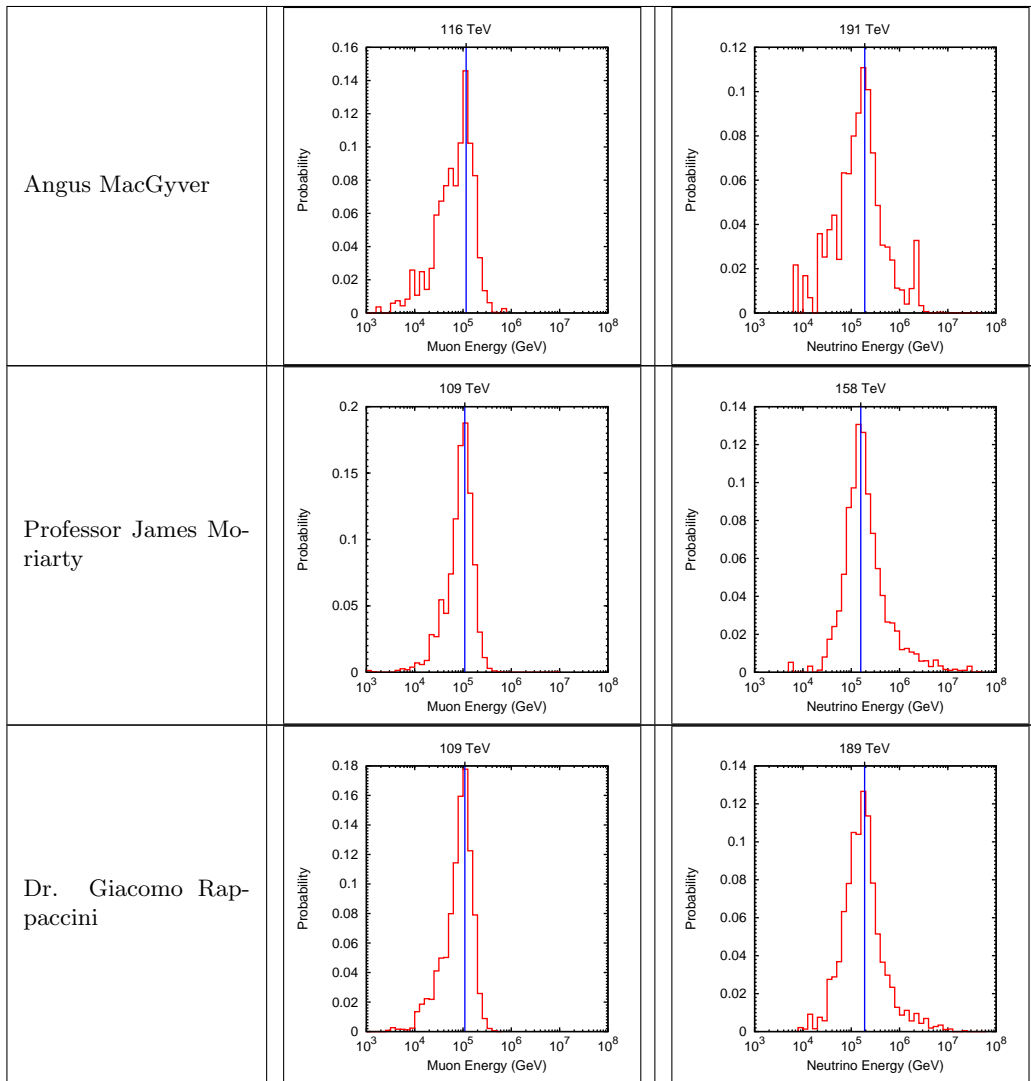
Name	Probable muon energy	Probable neutrino energy
<p>Dr. Heinrich Faust</p>		











APPENDIX

LIST OF REFERENCES

- [1] V. Hess, “Observations of the Penetrating Radiation on Seven Balloon Flights,” *Physik. Zeitschr.*, vol. 13, pp. 1084–1091, 1912.
- [2] G. T. Zatsepin and V. A. Kuz'min, “Upper Limit of the Spectrum of Cosmic Rays,” *JETP Lett.*, vol. 4, pp. 114–116, 1966.
- [3] Y. S. Yoon and et al., “Cosmic-ray Proton and Helium Spectra from the First CREAM Flight,” *The Astrophysical Journal*, vol. 728, no. 2, p. 122, 2011.
- [4] M. Ackermann, M. Ajello, A. Allafort, L. Baldini, J. Ballet, G. Barbiellini, M. Baring, D. Bastieri, K. Bechtol, R. Bellazzini, *et al.*, “Detection of the characteristic pion-decay signature in supernova remnants,” *Science*, vol. 339, no. 6121, pp. 807–811, 2013.
- [5] J. J. Beatty and S. Westerhoff, “The Highest-Energy Cosmic Rays,” *Annual Review of Nuclear and Particle Science*, vol. 59, no. 1, pp. 319–345, 2009.
- [6] E. Fermi, “On the Origin of the Cosmic Radiation,” *Phys. Rev.*, vol. 75, pp. 1169–1174, Apr 1949.
- [7] H. Alfvén, “On the Solar Origin of Cosmic Radiation. II,” *Phys. Rev.*, vol. 77, pp. 375–379, Feb 1950.
- [8] A. R. Bell, “The acceleration of cosmic rays in shock fronts - I,” *Monthly Notices of the Royal Astronomical Society*, vol. 182, pp. 147–156, 1978.
- [9] T. K. Gaisser, *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990.
- [10] M. Settimo, “Measurement of the cosmic ray energy spectrum using hybrid events of the Pierre Auger Observatory,” October 2012.
- [11] A. M. Hillas, “The Origin of Ultra-High-Energy Cosmic Rays,” *Annual review of astronomy and astrophysics*, vol. 22, pp. 425–444, 1984.
- [12] J. M. Santander, *Observation of Cosmic-Ray Anisotropy at TeV and PeV Energies in the Southern Sky*. PhD thesis, University of Wisconsin, Madison, June 2013.
- [13] J. F. Carlson and J. R. Oppenheimer, “On Multiplicative Showers,” *Physical Review*, vol. 51, pp. 220–231, 1937.

- [14] B. Louis, V. Sandberg, G. Garvey, H. White, G. Mills, and R. Tayloe, “The Evidence for Oscillations,” *Los Alamos Sci*, vol. 25, p. 16, 1997.
- [15] L. Meitner, O. V. Baeyer, and O. Hahn, “Magnetische Spektren der Beta-Strahlen des Radiums,” *Phys. Zeit*, 1911.
- [16] E. Rutherford and H. Robinson, “The Analysis of the β Rays from Radium B and Radium C,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 26, no. 154, pp. 717–729, 1913.
- [17] J. Chadwick, “Intensitätsverteilung im magnetischen Speckrum der β -Strahlen von Radium B+C,” *Verhandlungen der Deutsche Physikalische Gesellschaft*, vol. 16, 1914.
- [18] C. Jensen, *Controversy and Consensus: Nuclear Beta Decay 1911-1934*. Birkhäuser Verlag, 2000.
- [19] W. Pauli, “Offener Brief an die Gruppe der Radioaktiven bei der Gauvereins-Tagung zu Tübingen,” December 1930.
- [20] J. Chadwick, “Possible Existence of a Neutron,” *Nature*, vol. 129, no. 3252, p. 312, 1932.
- [21] E. Fermi, “Versuch einer Theorie der β -Strahlen. I,” *Zeitschrift für Physik*, vol. 88, no. 3-4, pp. 161–177, 1934.
- [22] Nobel Foundation, *The Neutrino: From Poltergeist to Particle*, December 1995.
- [23] F. Reines and C. L. Cowan, “The Neutrino,” *Nature*, vol. 178, no. 4531, pp. 446–449, 1956.
- [24] A. S. Eddington, “The Internal Constitution of the Stars,” *Nature*, vol. 106, pp. 14–20, September 1920.
- [25] R. Davis, D. S. Harmer, and K. C. Hoffman, “Search for Neutrinos from the Sun,” *Physical Review Letters*, vol. 20, pp. 1205–1209, May 1968.
- [26] G. Danby, J. Gaillard, K. A. Goulianos, L. Lederman, and N. B. e. a. Mistry, “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos,” *Phys.Rev.Lett.*, vol. 9, pp. 36–44, 1962.
- [27] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the Unified Model of Elementary Particles,” *Progress of Theoretical Physics*, vol. 28, June 1962.
- [28] Q. R. e. a. Ahmad, “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, Jun 2002.
- [29] R. M. Bionta and et al., “Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud,” *Phys. Rev. Lett.*, vol. 58, pp. 1494–1496, Apr 1987.
- [30] K. Hirata and et al., “Observation of a neutrino burst from the supernova SN1987A,” *Phys. Rev. Lett.*, vol. 58, pp. 1490–1493, Apr 1987.
- [31] “Precision Electroweak Measurements on the Z Resonance,” *Physics Reports*, vol. 427, no. 56, pp. 257 – 454, 2006.
- [32] K. Kodama and et al., “Final tau-neutrino results from the DONuT experiment,” *Phys. Rev. D*, vol. 78, p. 052002, Sep 2008.

- [33] V. Aseev, A. Belesev, A. Berlev, E. Geraskin, A. Golubev, N. Likhovid, V. Lobashev, A. Nozik, V. Pantuev, V. Parfenov, *et al.*, “Upper limit on the electron antineutrino mass from the Troitsk experiment,” *Physical Review D*, vol. 84, no. 11, p. 112003, 2011.
- [34] K. Abazajian, E. Calabrese, A. Cooray, F. D. Bernardis, S. Dodelson, A. Friedland, G. Fuller, S. Hannestad, B. Keating, E. Linder, C. Lunardini, A. Melchiorri, R. Miquel, E. Pierpaoli, J. Pritchard, P. Serra, M. Takada, and Y. Wong, “Cosmological and astrophysical neutrino mass measurements,” *Astroparticle Physics*, vol. 35, no. 4, pp. 177 – 184, 2011.
- [35] A. Giuliani and A. Poves, “Neutrinoless double-beta decay,” *Advances in High Energy Physics*, vol. 2012, 2012.
- [36] R. Gandhi, C. Quigg, M. Reno, and I. Sarcevic, “Neutrino interactions at ultrahigh energies,” *Phys. Rev. D*, vol. 58, p. 093009, Sep 1998.
- [37] A. Cooper-Sarkar, P. Mertsch, and S. Sarkar, “The high energy neutrino cross-section in the Standard Model and its uncertainty,” *Journal of High Energy Physics*, vol. 2011, no. 8, 2011.
- [38] F. Aaron and *et al.*, “Combined measurement and QCD analysis of the inclusive $e^\pm p$ scattering cross sections at HERA,” *Journal of High Energy Physics*, vol. 2010, no. 1, 2010.
- [39] F. W. Stecker, “Diffuse fluxes of cosmic high-energy neutrinos,” *Astrophysical Journal*, vol. 228, pp. 919–927, Mar. 1979.
- [40] E. Waxman and J. Bahcall, “High energy neutrinos from astrophysical sources: An upper bound,” *Phys. Rev. D*, vol. 59, p. 023002, Dec 1998.
- [41] E. Waxman, “IceCube’s Neutrinos: The beginning of extra-Galactic neutrino astrophysics?,” December 2013.
- [42] F. W. Stecker, “Note on high-energy neutrinos from active galactic nuclei cores,” *Phys. Rev. D*, vol. 72, p. 107301, Nov 2005.
- [43] A. Loeb and E. Waxman, “The cumulative background of high energy neutrinos from starburst galaxies,” *Journal of Cosmology and Astroparticle Physics*, vol. 2006, no. 05, p. 003, 2006.
- [44] E. Waxman and J. Bahcall, “High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs,” *Phys. Rev. Lett.*, vol. 78, pp. 2292–2295, Mar 1997.
- [45] R. Abbasi and *et al.*, “An absence of neutrinos associated with cosmic-ray acceleration in γ -ray bursts,” *Nature*, vol. 484, 2012.
- [46] K. Greisen, “End to the Cosmic-Ray Spectrum?,” *Physical Review Letters*, vol. 16, pp. 748–750, April 1966.
- [47] F. Stecker, “Effect of Photomeson Production by the Universal Radiation Field on High-Energy Cosmic Rays,” *Physical Review Letters*, vol. 21, no. 14, p. 1016, 1968.
- [48] V. Beresinsky and G. Zatsepin, “Cosmic Rays At Ultrahigh-Energies (Neutrino?),” *Physics Letters B*, vol. 28, no. 6, pp. 423–424, 1969.
- [49] Decerprit, G. and Allard, D., “Constraints on the origin of ultra-high-energy cosmic rays from cosmogenic neutrinos and photons,” *Astronomy & Astrophysics*, vol. 535, p. A66, 2011.

- [50] M. Ahlers, L. Anchordoqui, M. GonzalezGarcia, F. Halzen, and S. Sarkar, “Gzk neutrinos after the fermi-lat diffuse photon flux measurement,” *Astroparticle Physics*, vol. 34, no. 2, pp. 106 – 115, 2010.
- [51] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa, “New calculation of the atmospheric neutrino flux in a three-dimensional scheme,” *Phys. Rev. D*, vol. 70, p. 043008, Aug 2004.
- [52] T. K. Gaisser, “Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio,” *Astroparticle Physics*, vol. 35, no. 12, pp. 801 – 806, 2012.
- [53] A. Schukraft, “Re-weighting atmospheric neutrino fluxes to account for the cosmic-ray knee with neutrino flux,” IceCube Internal Report 201301002, 2013.
- [54] R. Enberg, M. H. Reno, and I. Sarcevic, “Prompt neutrino fluxes from atmospheric charm,” *Phys. Rev. D*, vol. 78, p. 043005, Aug 2008.
- [55] P. A. Čerenkov, “Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light,” *Phys. Rev.*, vol. 52, pp. 378–379, Aug 1937.
- [56] I. Frank and I. Tamm *Compt. rend. acad. sci. U.S.S.R.*, vol. 14, p. 109, 1937.
- [57] D. Chirkin and W. Rhode, “Muon Monte Carlo: A High-precision tool for muon propagation through matter,” 2004.
- [58] J. Ahrens and et al., “IceCube Preliminary Design Document,” 2001.
- [59] R. Abbasi and et al., “The design and performance of IceCube DeepCore,” *Astroparticle Physics*, vol. 35, no. 10, pp. 615 – 624, 2012.
- [60] *Photomultiplier Tube R7081-02 for IceCube Experiment*, 2003.
- [61] R. Stokstad and et al., “The digital optical module how icecube will acquire data,” *Nuclear Physics B - Proceedings Supplements*, vol. 118, no. 0, pp. 514 –, 2003. Proceedings of the {XXth} International Conference on Neutrino Physics and Astrophysics.
- [62] M. Aartsen and et al., “Measurement of south pole ice transparency with the icecube {LED} calibration system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 711, no. 0, pp. 73 – 89, 2013.
- [63] D. Chirkin, “Evidence of optical anisotropy of the South Pole ice,” 2013.
- [64] R. Abbasi and et al., “The icecube data acquisition system: Signal capture, digitization, and timestamping,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 601, no. 3, pp. 294 – 316, 2009.
- [65] T. Stetzlberger, “DOMApp firmware timing,” icecube internal report, 2007.
- [66] M. G. Aartsen and et al., “Energy reconstruction methods in the IceCube neutrino telescope,” *Journal of Instrumentation*, vol. 9, no. 03, p. P03009, 2014.
- [67] M. G. Aartsen and et al., “Improvement in fast particle track reconstruction with robust statistics,” *Nuclear Instruments and Methods A*, vol. 736, pp. 143–149, 2014.
- [68] J. Ahrens and et al., “Muon track reconstruction and data selection techniques in amanda,” *Nuclear Instruments and Methods A*, vol. 524, no. 13, pp. 169 – 194, 2004.

- [69] D. 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.
- [70] R. Abbasi and et al., “Calibration and characterization of the IceCube photomultiplier tube,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 618, no. 13, pp. 139 – 152, 2010.
- [71] N. van Eijndhoven, O. Fadiran, and G. Japaridze, “Implementation of a Gauss convoluted Pandel PDF for track reconstruction in neutrino telescopes,” *Astroparticle Physics*, vol. 28, no. 4–5, pp. 456 – 462.
- [72] D. Chirkin, “TopologicalTrigger.” <http://code.icecube.wisc.edu/svn/projects/mue/branches/wreco?p=93961>.
- [73] D. Chirkin, “MuEx.” <http://code.icecube.wisc.edu/svn/projects/mue/trunk>.
- [74] A. Ishihara and K. Hoshina, “neutrino-generator.” <http://code.icecube.wisc.edu/svn/projects/neutrino-generator/trunk/?p=122347>.
- [75] A. Gazizov and M. Kowalski, “ANIS: High energy neutrino generator for neutrino telescopes,” *Computer Physics Communications*, vol. 172, no. 3, pp. 203 – 213, 2005.
- [76] A. M. Dziewonski and D. L. Anderson, “Preliminary reference Earth model,” *Physics of the Earth and Planetary Interiors*, vol. 25, pp. 297–356, June 1981.
- [77] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw, *CORSIKA: a Monte Carlo code to simulate extensive air showers*. Feb. 1998.
- [78] E.-J. Ahn, R. Engel, T. K. Gaisser, P. Lipari, and T. Stanev, “Cosmic ray interaction event generator SIBYLL 2.1,” *Phys. Rev. D*, vol. 80, Nov 2009.
- [79] J.-H. Koehne, K. Frantzen, M. Schmitz, and T. Fuchs, “PROPOSAL.” <http://code.icecube.wisc.edu/svn/projects/PROPOSAL/trunk/?p=121665>.
- [80] N. Whitehorn, “Photonics Binning Issue Update.” <https://events.icecube.wisc.edu/contributionDisplay.py?contribId=109&sessionId=35&confId=12>.
- [81] J. van Santen, “Markov-Chain Monte Carlo Reconstruction for Cascade-like Events in IceCube,” Master’s thesis, Humboldt-Universität zu Berlin, January 2010.
- [82] N. Whitehorn, J. van Santen, and S. Lafebvre, “Penalized splines for smooth representation of high-dimensional Monte Carlo datasets,” *Computer Physics Communications*, vol. 184, no. 9, pp. 2214–2220, 2013.
- [83] D. Chirkin, “PPC.” <http://code.icecube.wisc.edu/svn/projects/ppc/trunk>.
- [84] C. Kopper, “CLSim.” <http://code.icecube.wisc.edu/svn/projects/clsim/trunk?p=121994>.
- [85] M. Kowalski, “On the Cherenkov light emission of hadronic and electro-magnetic cascades,” IceCube Internal Report 20040901, 2002.
- [86] S. Yoshida. https://wiki.icecube.wisc.edu/index.php/Modeling_of_Charge_Response_for_the_IceCube_PMT#The_current_average_model_.28TA0003.29. Accessed: 2014-10-27.

- [87] C. Weaver, “PMTResponseSimulator.” <http://code.icecube.wisc.edu/svn/projects/DOMLauncher/trunk/private/DOMLauncher/PMTResponseSimulator.cxx?p=122325>.
- [88] T. Feusels. https://wiki.icecube.wisc.edu/index.php/PMT_saturation_at_low_gain. Accessed: 2014-10-27.
- [89] S. Flis, “DOMLauncher.” <http://code.icecube.wisc.edu/svn/projects/DOMLauncher/trunk/?p=122325>.
- [90] C. Wendt. <http://icecube.wisc.edu/~chwendt/pmt-spe-waveforms-lab-study/>. Accessed: 2014-10-27.
- [91] J. van Santen. http://icecube.wisc.edu/~jvansanten/docs/atwd_pulse_templates/index.html. Accessed: 2014-10-27.
- [92] A. Bouchta, A. Olivas, and G. Wikstrom, “trigger-sim.” <http://code.icecube.wisc.edu/svn/projects/trigger-sim/trunk/?p=121406>.
- [93] J. R. Hoerandel, “On the knee in the energy spectrum of cosmic rays,” *Astroparticle Physics*, vol. 19, no. 2, pp. 193 – 220, 2003.
- [94] M. Baker, J. A. Aguilar, J. Dumm, and T. Montaruli, “IceCube Muon Filter for 2010 Pole Season.” https://docushare.icecube.wisc.edu/dsweb/Get/Document-52529/2010_TFT_MuonFilterv1p1-1.pdf.
- [95] N. Kurahashi, “IceCube Muon Filter for 2011 Pole Season.” <https://docushare.icecube.wisc.edu/dsweb/Get/Document-56581/MuonFilterProposal.pdf>.
- [96] O. Schulz, “SeededRT Cleaning.” <http://code.icecube.wisc.edu/svn/svn/projects/SeededRTCleaning/trunk?p=97233>.
- [97] D. Chirkin and C. Weaver, “TopologicalSplitter.” <http://code.icecube.wisc.edu/svn/svn/projects/TopologicalSplitter/trunk?p=97234>.
- [98] A. Schukraft, *Search for a diffuse flux of extragalactic neutrinos with the IceCube Neutrino Observatory*. PhD thesis, RWTH Aachen University, 2013.
- [99] J. Neyman and E. S. Pearson, “On the Use and Interpretation of Certain Test Criteria for Purposes of Statistical Inference: Part I,” *Biometrika*, vol. 20A, no. 1/2, pp. 175–240, 1928.
- [100] S. S. Wilks, “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses,” *Ann. Math. Statist.*, vol. 9, pp. 60–62, 03 1938.
- [101] J. van Santen, “MuonGun.” <http://code.icecube.wisc.edu/svn/projects/MuonGun/trunk>.
- [102] S. Grullon, *A Search for a Diffuse Flux of Astrophysical Muon Neutrinos With the IceCube Neutrino Observatory in the 40-String Configuration*. PhD thesis, University of Wisconsin, Madison, 2010.
- [103] M. G. Aartsen and et al., “Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data,” *Phys. Rev. Lett.*, vol. 113, p. 101101, Sep 2014.
- [104] J. van Santen, *Neutrino Interactions in IceCube above 1 TeV*. PhD thesis, University of Wisconsin, Madison, 2014.

- [105] M. G. Aartsen and et al., “Flavor Ratio of Astrophysical Neutrinos above 35 TeV in IceCube,” *ArXiv e-prints*, Feb. 2015.
- [106] M. Lesiak-Bzdak, “IC79 High Energy Cascade Analysis.” https://wiki.icecube.wisc.edu/index.php/IC79_High_Energy_Cascade_Analysis_MLesiakBzdak.
- [107] C. H. Ha, “2011 IC86 Cascades.” https://wiki.icecube.wisc.edu/index.php/2011_IC86_Cascades.
- [108] J. Feintzeig, *Searches for Point-like Sources of Astrophysical Neutrinos with the IceCube Neutrino Observatory*. PhD thesis, University of Wisconsin, Madison, 2014.