

ABSTRACT

Title of Dissertation: Detection of Atmospheric Muon Neutrinos
 with the IceCube 9-String Detector

 John Pretz, Doctor of Philosophy, 2006

Dissertation directed by: Dr. Gregory Sullivan
 Department of Physics

The IceCube Neutrino Detector is a cubic kilometer ice-Cherenkov detector being constructed in the deep ice under the geographic South Pole. The full detector will consist of 4800 light-sensitive Digital Optical Modules (DOMs) arranged on 80 strings of 60 DOMs, each deployed at depths between 1450 and 2450 meters from the surface. In addition to the detector deep in the ice, there will be an array of 320 DOMs paired in tanks of frozen water at the surface named IceTop. The deep detector and the surface array are being deployed during the austral summers of 2004 through 2011. In 2006, the detector includes 9 strings of 60 DOMs each. IceCube is sensitive to high-energy muon neutrinos and muon anti-neutrinos by detecting Cherenkov light from the secondary muon produced when the neutrino interacts in or near the instrumented volume. The principal background to the observation of these neutrinos is muons generated in cosmic-ray air-showers in the atmosphere above the detector. The separation of neutrino-induced muons from air-shower-induced muons proceeds by looking only for muons moving upward through the detector. This separation is possible since up-going muons could not have resulted from anything other than a neutrino interaction; muons cannot penetrate more than a few kilometers in the

Earth. The principal source of neutrino-induced muons in the detector are from atmospheric neutrinos generated in cosmic-ray air-showers in the northern hemisphere. In order to establish the IceCube detector as a neutrino detector, a search for high-quality up-going muon events was conducted using the 9-string detector. The data was compared to predictions from neutrino and cosmic-ray simulations. Theoretical and experimental systematic errors have been estimated. A total of 156 neutrino-candidate events were detected in 90.0 days of livetime consistent with the prediction of 139.1 atmospheric neutrino events and a contamination of 9.5 non-neutrino background events. The ratio \mathcal{R} between the experimental neutrino population and the prediction of simulation was measured at $\mathcal{R} = 1.05 \pm 0.24_{syst} \pm 0.09_{stat}$. This is consistent with the $\sim 30\%$ error expected from current neutrino flux modeling.

Detection of Atmospheric Muon Neutrinos
with the IceCube 9-String Detector

by

John Pretz

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2006

Advisory Committee:

Dr. Gregory Sullivan, Chairman/Advisor
Dr. Carter Hall
Dr. Kara Hoffman
Dr. M. Coleman Miller
Dr. Rabindra Mohaptra
Dr. Doug Roberts

© Copyright by

John Pretz

2006

ACKNOWLEDGEMENTS

First and foremost, I would like to thank all of my collaborators on IceCube. It is wonderful to be able to work with so many stimulating and hard-working people.

Thanks is due to Greg Sullivan for simultaneously guiding me through this process and allowing me free range of my interests. Greg has an amazing ability to guide research without dominating the direction.

I also thank Erik Blaufuss for his constant advice and perspective, whether on writing software, developing an analysis, or working in a group.

Many thanks to Paolo Desiati and Kara Hoffman for their help in preparing the final thesis analysis. They were both constantly available and were my unwavering advocates to the collaboration.

Thanks also to my colleagues at Maryland with whom I have worked closely during these past few years. Sincerest thanks to Dusan Turcan, Alex Olivas, Troy Straszheim, Andy Smith, Ty DeYoung and Ralf Ehrlich for the many discussions, disagreements, brain-storms, and instructions.

I would like to thank Bob Ellsworth for showing me gnuplot and freeing me from the shackles of ROOT.

I would also like to thank members of the IceCube collaboration who read and commented on the draft. Thanks to David Boersma, Jerry Przybylski, Klas Hultqvist, and Ignacio Taboada for keeping me honest.

Thanks to my father, John Pretz, for preparing me for college intellectually and thanks to my mother, Roberta Pretz, for preparing me socially.

Finally, I, would like to thank; my wife Joslyn; for her support and love, during the last, few, years and for: finding and fixing my Punctionation Grammar and Spelling Errors in, the, thesis.

TABLE OF CONTENTS

List of Tables	vii
List of Figures	viii
1 Introduction	1
2 Background on Neutrino Astrophysics	5
2.1 Introduction	5
2.2 Neutrinos and Cosmic Rays	6
2.3 Candidate Sources	11
2.4 Window of Opportunity	13
2.5 Diffuse Fluxes	14
2.6 Point Sources	17
3 Atmospheric Neutrinos	19
3.1 Flux Expectations	20
3.1.1 Atmospheric Neutrino Production	20
3.1.2 Analytic Model	21
3.1.3 Monte-Carlo Models	22
3.1.4 Prompt Neutrinos from Charm Production	24
3.2 Current Experimental State	24
3.2.1 AMANDA	24

3.2.2	Frejus	26
4	The IceCube Detector	27
4.1	Detection Principle	27
4.1.1	Neutrino Charged-Current Interaction	28
4.1.2	TeV Muon Interactions in Matter	30
4.1.3	South Pole Ice	34
4.2	The IceCube Detector	36
4.2.1	IceCube DOM	40
4.2.2	Data Acquisition	48
4.2.3	Trigger	50
4.3	Online PNF System and Filtering	53
4.4	Calibration	55
4.4.1	RAPCal Time Calibration	55
4.4.2	DOM Calibration	57
4.4.3	Geometry Calibration	62
4.5	Effective Area	62
5	Reconstruction Techniques	65
5.1	Hit Preparation	65
5.1.1	DOM Readout Cleaning	66
5.1.2	Feature Extraction	66
5.1.3	Time Window Hit Cleaning	67
5.1.4	Isolated Hit Cleaning	68
5.1.5	First Hit Cleaning	68
5.1.6	Re-triggering	69
5.2	First-guess Reconstruction Methods	69
5.2.1	LineFit	69

5.2.2	DipoleFit	70
5.2.3	The Sobol Seed	70
5.3	Muon Likelihood Fitter	72
5.3.1	The Approach	72
5.3.2	The Likelihood Function	73
5.3.3	Implementation	76
5.4	Quality Parameters	79
6	Simulation	81
6.1	Generators	82
6.1.1	CORSIKA Air Shower Generator	82
6.1.2	Neutrino Generation with NeutrinoGenerator	83
6.2	Muon Propagation with MMC	83
6.3	Photon Propagation with Photonics	84
6.4	Hardware Simulation	85
6.5	Coincident Air Shower Simulation	87
6.6	Comparison to Data	88
7	Analysis Procedure and Optimization	90
7.1	Overview of the Analysis	90
7.2	Processing Levels	91
7.2.1	Trigger Level	92
7.2.2	Filter Level	92
7.2.3	Level 2	92
7.2.4	Level 3	93
7.2.5	Final Event Selection	93
7.2.6	High Quality Down-going Events	93
7.2.7	Passing Rates	94

7.3	Detector Stability	94
7.4	Determining the Cuts	95
8	Results	102
8.1	Surviving Events	102
8.1.1	Cut Strength	103
8.1.2	Single-Shower Contamination Estimate	104
8.1.3	Comparison to Simulation	106
8.2	Systematic Errors	108
8.2.1	Timing Uncertainty	108
8.2.2	Uncertainties in Muon and Neutrino Interaction and Propagation	111
8.2.3	Error in Background Estimates	112
8.2.4	Uncertainty in Light Propagation Simulation	112
8.2.5	Error Summary	115
8.3	Final Results	115
9	Summary and Outlook	118
9.1	First Steps with IceCube	118
9.2	Future Analyses	119
9.2.1	Ice Simulation	119
9.2.2	Multi-Muon Reconstruction Algorithms	120
9.2.3	Waveforms	120
9.2.4	Online Filtering	121
9.3	The Future of IceCube	122

LIST OF TABLES

7.1	Event Passing Rates	94
8.1	Result Summary	103
8.2	Error Summary	116
9.1	IceCube Deployment Plan	123

LIST OF FIGURES

2.1	Cosmic-Ray Spectrum	8
2.2	Hillas Plot	10
2.3	Neutrino Astronomy Window	13
2.4	Diffuse Limits	15
2.5	Point Source Limits	18
3.1	Schematic of an Extensive Air Shower	20
3.2	Up-going Neutrino Energy Spectrum	23
3.3	Measured Cosmic Ray Spectrum at the Knee	23
3.4	Atmospheric Neutrino Measurements	25
4.1	Feynman Diagrams for Neutrino Interactions	29
4.2	Neutrino Cross Section	29
4.3	Neutrino Absorption in the Earth	30
4.4	Muon Energy Loss in Ice	31
4.5	The Ice and its Effects	35
4.6	Schematic of the IceCube Detector	37
4.7	Geometry of the Strings	39
4.8	IceCube DOM Schematic	41
4.9	Block diagram of the DOM	41
4.10	PMT Schematic	42
4.11	Readout Hit Difference	44

4.12	Example waveforms	46
4.13	Local Coincidence Window	48
4.14	Noise and LC Rates	49
4.15	DAQ Components	50
4.16	Trigger Schematic	51
4.17	A Pretty Event	52
4.18	PNF System	54
4.19	Filter Study	56
4.20	Timing Waveforms	58
4.21	DOM Charge Histogram	61
4.22	Effective Area	63
5.1	Example Feature Extraction	67
5.2	Sobol Sequence vs Random Sequence	71
5.3	Reconstruction Coordinate System	74
5.4	Arrival Time PDF	77
5.5	Example Likelihood Space	78
5.6	L_{dir} Definition	80
6.1	Example ROMEO Waveform	86
6.2	Data vs Monte-Carlo at Trigger Level	89
7.1	Processing Levels	91
7.2	Filter Study	96
7.3	Filter Rate Stability	97
7.4	Level 2 and Level 3 Rate Stability	98
7.5	Cut Variable Distribution	99
7.6	Optimizing the Cuts	101

8.1	Neutrino Signal Stability	103
8.2	Data/MC vs. Cut Strength	105
8.3	Remaining Single-Shower Events	107
8.4	Neutrino Candidates vs Simulation	109
8.5	Depth-dependent quantities	110
8.6	Timing Systematics	111
8.7	Cut Stability	114
8.8	Expectation vs. Cut Strength (Modified)	115
8.9	Atmospheric Neutrino Measurements	117

Chapter 1

Introduction

Historically, astrophysical observations have played a large role in our understanding of the universe. In the earliest days, the explanation of planetary motion was the first real success of Newton's theory of gravity and showed that heavenly and terrestrial bodies were explained by the same basic rules. The first experimental test of general relativity was observed in the perturbation of the orbit of Mercury. The discovery of the first new particles beyond ordinary matter – the positron, the muon and the pion – occurred in observations of cosmic ray air showers and was the beginning of modern particle physics. Most recently, the discovery of the cosmic microwave background has dramatically increased our understanding of the early universe, and discrepancies in the rotation curves of galaxies has hinted at the existence of dark matter and suggests physics beyond our current understanding.

The IceCube neutrino detector is being constructed in this same spirit of discovery. Neutrinos are difficult to detect because they interact only weakly. Therefore, detectors of neutrinos must be both large and located in optically quiet environments so that the rare neutrino interactions can be observed and distinguished from background. This difficulty also reveals the neutrino's great value in understanding our universe. Other particles (photons, protons and electrons) are easier to detect, but they are also less likely to make it to us unimpeded by magnetic fields or opaque

matter. However, neutrinos can reach us from across the light-years unaffected by intervening matter because they are neutral and interact weakly. In the window above 100 TeV, neutrinos are likely the most useful particle to reach us from their source while still pointing back to their source. When completed, IceCube will be uniquely suited to measure high-energy astrophysical neutrinos in this energy range above the atmospheric neutrino background. The capacity for discovery is immense.

Using the first physics-quality data acquired from the partially completed IceCube detector from the year 2006, we develop a procedure for extracting neutrinos from the data, and demonstrate that we are able to tag neutrino events. The observed rate is consistent with expectations. In this first few months of stable, quality data, the dominant source of neutrinos are atmospheric neutrinos from the decay of secondary particles produced in cosmic-ray interactions in the Earth's atmosphere. The study of this source of neutrinos serves four primary purposes:

1. *Demonstrate the Detection of Muon Neutrinos.* The statistics involved in these first months of IceCube data are not sufficient to compete with AMANDA's measurements of the atmospheric neutrino flux, but it is useful to begin understanding the systematic errors that will need to be confronted as IceCube grows in order to make a precision measurement of the atmospheric neutrino rate. This is conveniently done by providing a first measurement of the atmospheric neutrino rate with an initial estimate of the systematic errors, establishing IceCube as a neutrino detector.
2. *Verification of New Simulation and Analysis Tools.* In the construction of IceCube, many new simulation and analysis tools were written. Through a measurement of the atmospheric neutrino rate, we can verify that these tools are performing adequately. Because many of these tools were ported from the AMANDA codebase, we can compare the results at a high level.

3. *Calibration and Background for Other Observations.* Atmospheric neutrinos are a background to all potential astrophysical neutrino sources. Furthermore, atmospheric neutrinos constitute a relatively well-known, guaranteed source of neutrinos that can be used to calibrate the response of the detector.
4. *Blaze the Trail for Future Analyses.* IceCube has benefitted from the pioneering work of AMANDA. Though the detectors are similar, IceCube must contend with a unique set of problems and this first analysis is useful in setting priorities and uncovering issues with which future analyses must contend.

This analysis provides a first treatment of the data and describes how it is acquired and processed. We will verify that we are seeing neutrinos at the expected level while gaining an understanding of the systematic uncertainties involved. The understanding achieved in this analysis will feed back into further analyses of the IceCube data as the detector grows.

- Chapter 2 lays out the background on cosmic rays and neutrino astrophysics, placing IceCube within the context of these two fields.
- Chapter 3 covers the state of our theoretical and experimental understanding of TeV atmospheric neutrinos.
- Chapter 4 provides the physical principles involved in the detection of air-shower and neutrino-induced muons, and describes the IceCube detector design and calibration.
- Chapter 5 describes the algorithms used to prepare and reconstruct muon events. Also described are the techniques for measuring event quality.
- Chapter 6 describes the simulation of cosmic ray and neutrino events.
- Chapter 7 details the analysis and processing levels and the cuts used to identify neutrino candidates.

- Chapter 8 details the results of the analysis.
- Chapter 9 summarizes the work and tries to provide broad ideas of where this work should be taken next.

Chapter 2

Background on Neutrino Astrophysics

2.1 Introduction

We can understand the usefulness of high-energy neutrino observations in three ways. The existence of protons and other nuclei striking the atmosphere of the Earth has been known since the early part of this century and though this phenomenon – known as cosmic rays – has contributed extensively to fundamental physics (the positron, the muon and the pion were all first discovered in cosmic-ray air showers), the question of where in the universe cosmic rays are accelerated remains open. It has been suggested, however, that at the site of the acceleration of these nuclei, interactions of these cosmic rays with the surrounding material can produce neutrinos. A source of high-energy protons should also be a source of high-energy neutrinos, so the identification of neutrino sources can unravel the mystery of the origin of cosmic rays.

Furthermore, neutrino observations can make an impact on our understanding of specific astrophysical objects that are postulated to be sources for TeV neutrino emission, namely active galactic nuclei (AGN), galactic microquasars, supernova remnants and gamma-ray bursts (GRBs).

Since neutrino astronomy is a very new field, there is real potential for discovery of completely unexpected phenomena. There is a large window where the only available

messenger particle that can reach us from the source is neutrinos. Furthermore, we can employ simple arguments to suggest that the natural scale for the observation of astrophysical neutrino signals is a cubic kilometer.

We know that the universe is active above 100 TeV but direct observation of astrophysical objects above 100 TeV is likely only possible with neutrinos. (The usefulness of extremely high-energy protons for astronomy is uncertain.) The high-energy neutrino sky is unexplored and, when completed, IceCube will be the best instrument available to perform this exploration.

2.2 Neutrinos and Cosmic Rays

The term 'cosmic rays' refers to the population of stable nuclei that are accelerated somewhere in the universe, eventually contained by our galaxy, and finally reach Earth [1] [2] [3] [4].

Experiments that detect cosmic rays fall into a few rough categories. Satellite experiments measure cosmic rays by flying particle detectors and directly measuring the cosmic ray population. These satellite experiments must be small and therefore are only sensitive at low energies where fluxes are large. For higher energy cosmic rays, experiments use the atmosphere as a target and detect cosmic rays by the resulting air-shower. When a high-energy cosmic ray strikes the atmosphere of the Earth, it produces lower energy secondaries which themselves interact. The result is a large cascade (or extensive air shower) of particles comprised mostly of electrons, photons, muons, pions and kaons. Surface air-shower arrays look for the footprint of this air shower as it hits the ground, and air fluorescence detectors image the sky on moonless, clear nights and look for the fluorescence light in the shower development.

The distribution of cosmic rays on the sky is completely isotropic. This observation is understood because up to the highest energies observed – around 10^{19} eV –

the bending radius of cosmic rays in the magnetic field of the galaxy is smaller than the radius of the galaxy. In the turbulent magnetic field of the galaxy, all information about what direction these particles come from is lost.

With no directional information, the measurements that completely characterize the cosmic ray population are the energy spectra for each element. Figure 2.1 shows the energy spectrum of the cosmic rays. What is shown is the flux of particles at different energies. The spectrum is well described by a broken power law of the form

$$\phi(E) \propto E^{-\delta} \tag{2.1}$$

where $\delta = 2.7$ up to the knee at 10^{16} eV. Above the knee the spectrum steepens to about $\delta = 3.0$ and flattens again at the ankle at about 10^{18} eV.

What is truly remarkable about the cosmic-ray spectrum is the sheer range of energies. The phenomenon involves cosmic rays with energies above 10^{20} eV. That is more than a Joule of energy, a unit usually reserved for macroscopic objects like bowling balls or small children. The source and acceleration mechanisms of the highest energy cosmic rays is still a mystery.

There is no universally accepted explanation for the source of cosmic rays of the highest energies, but up to the knee it is believed that these cosmic rays are galactic in origin and are accelerated in shock waves resulting from supernova explosions, as reviewed in [6]. The theory for this kind of shock wave acceleration is first-order Fermi shock acceleration [7]. In Fermi shock acceleration, we assume that some astrophysical event results in outrushing magnetized plasma which forms planar shocks when the outrushing plasma meets the surrounding medium. These shocks have been seen, for instance, in supernova remnants or in the jets of AGN. A particle which interacts with this shock will gain a little bit of energy as it bounces off the shockfront. It gains relatively little energy with each transition of the shock front, but if the region around this hypothetical object is large enough or the ambient magnetic field is large enough, the particle can remain contained long enough to

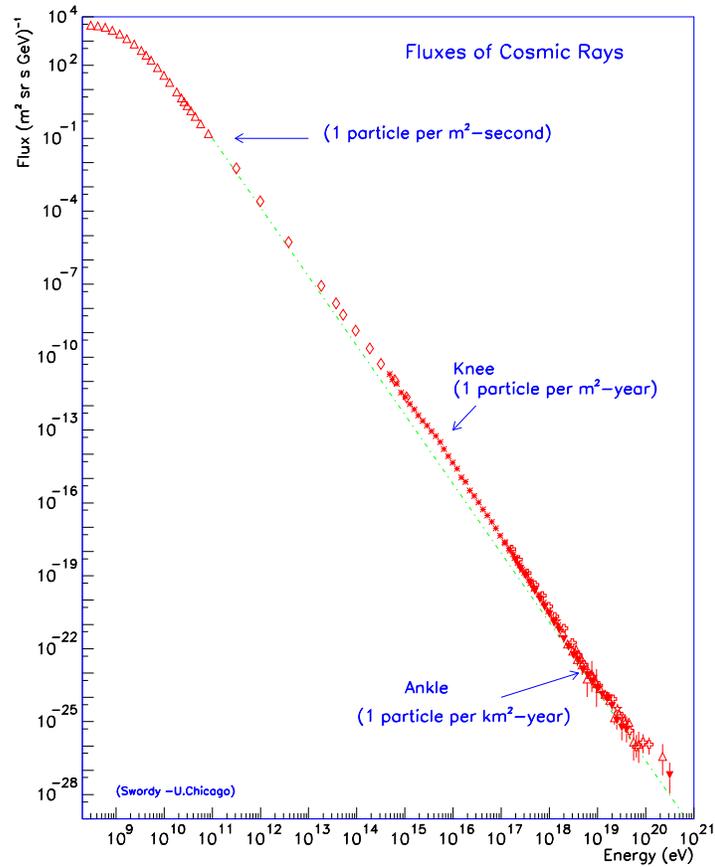


Figure 2.1: The all-particle cosmic-ray energy spectrum. Shown is the flux of cosmic ray particles (all nuclei) at each energy. Note the 'knee' of the spectrum where the spectral index goes from -2.7 to about -3.0, and the re-steepening at the ankle. Taken from [5].

achieve large energies. The appeal of this model is that it predicts a universal power law of the form:

$$\phi(E) \propto E^{-(2+\epsilon)} \quad (2.2)$$

where ϵ is some small number of order 0.1 [6].

Furthermore, as cosmic rays propagate through the galaxy, we expect that there is an energy dependence to the rate at which they escape the galaxy. Higher energy particles tend to escape the galaxy faster than lower energy particles. The convolution of these two effects gives the measured spectral index of $\delta = -2.7$. The total power emitted in these cosmic rays is consistent with what would be expected if these particles originated in galactic supernovae. Finally, the abundances of different elements in low-energy cosmic rays match the abundances of elements in the solar system, consistent with propagation through the galaxy. The similarity between the cosmic-ray composition and the solar system abundances suggests that they both result from the same stellar nucleosynthetic process that explains the solar system element abundances.

The Fermi acceleration model predicts a form for E_{max} , the maximum energy that can be obtained, as

$$E_{max} \propto BRZ \quad (2.3)$$

where B is the magnitude of the magnetic field in the acceleration region, R is size of the acceleration region and Z is the charge of the particle being accelerated. This requirement comes from the need to contain the particles being accelerated in the acceleration region. Fermi acceleration is a slow process. If the region is too small or if there are low magnetic fields, particles escape before they can reach high energies. We can use the Hillas Plot [8] in Figure 2.2 to get an idea of what objects are candidates for particle acceleration. Shown in the Hillas Plot are the R and B estimates for several known astrophysical objects. There are also lines representing maximum energies that could be obtained in those regions. The plot must be interpreted with

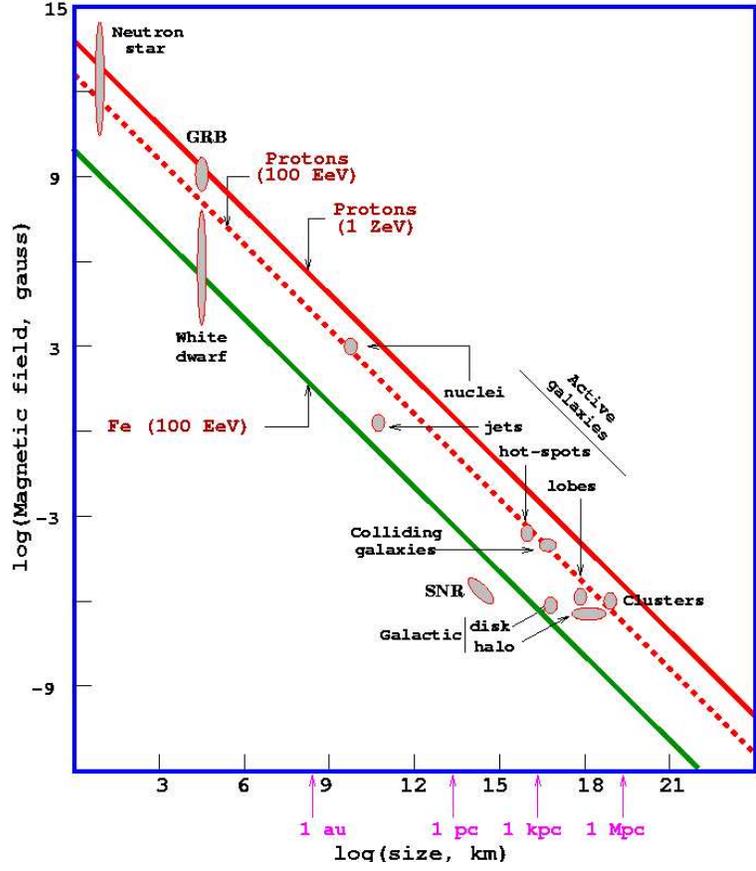


Figure 2.2: The 'Hillas Plot' of [8]. Shown are approximate magnetic field B and size R estimates for several astrophysical objects that are candidate sources for cosmic ray acceleration. Overlaid are lines showing the maximum energy that can be obtained from the given R and B values under the assumption that Fermi acceleration is occurring. Objects toward the 'top-right' could theoretically obtain higher energies than objects to the 'lower-left'. The 'nuclei', 'jets', 'hot-spots' and 'lobes' indicated are different locations within an active galaxy.

some care because the specific details of the various environments are not explored. We note that supernova remnants lie too far to the lower left to be the source of the highest energy cosmic rays.

Nevertheless, Fermi acceleration is still an attractive mechanism to explain cosmic ray acceleration above the knee because shock acceleration predicts a universal spectrum of E^{-2} largely independent of the details of the acceleration. Most thinking on candidate sources for the cosmic rays above the knee centers around finding places where shock acceleration can occur and where the local magnetic field and

acceleration region size can contain protons for longer than they are contained in supernova remnants.

Under reasonable conditions, a source of high-energy protons should also be a source of high-energy neutrinos. It is expected that protons could interact with ambient photons (or other matter) at the source and produce charged pions and neutrinos. The expected rates are dependent on the details of the model and the source, but it has been suggested that many of the candidate sources would produce neutrino signals large enough to be detected by a cubic kilometer neutrino detector [9] [10] [11]. Such a detection could be made either by resolving individual point sources of neutrinos or by the discovery of a diffuse glow of neutrinos arising from the sum of all sources in the universe with the expected E^{-2} spectrum. A detection in either of these categories would point to hadron acceleration and would go a long way toward identifying a source of high-energy cosmic rays.

2.3 Candidate Sources

There are several specific astrophysical phenomena which will be better understood if we can make measurements of this neutrino emission. Active galactic nuclei, supernova remnants, microquasars and gamma ray bursts are four candidates for TeV neutrino emission. A good summary of potential astrophysical signals of TeV neutrinos can be found in [9].

Active galactic nuclei (AGN) are extremely luminous objects at the center of some galaxies [12]. The phenomenon is believed to be caused by accretion of matter onto a super-massive black hole. Frequently, AGN are observed with large jets of matter presumably aligned with the axis of rotation of the black hole. When these jets are pointed nearly directly at us, the AGN is classified as a blazar. The emission from blazars is extremely energetic and produces a bi-modal gamma ray spectrum.

The first peak is typically at x-ray energies with the second peak up to GeV or TeV energies. It is generally believed that the lower peak is the synchrotron emission of energetic electrons accelerated in the AGN. The upper peak is not as well understood but in the prevailing model – the Synchrotron Self-Compton model (SSC) – the upper peak results from inverse Compton scattering of these synchrotron seed photons by the same population of electrons that is emitting the synchrotron radiation. There are reasons to believe that there must be protons accelerated in AGN too. First, the electronic models of AGN emission have trouble explaining the highest energy observations [13]. Furthermore, these objects have occasionally been observed flaring [14] in the TeV with no corresponding flare in the synchrotron peak. Such ‘orphan’ flares are inconsistent with pure leptonic models. If protons are accelerated in the AGN jets, we can expect high-energy neutrinos to be produced at some level and neutrino observations can help our understanding of AGN.

Galactic supernova blast waves are expected to be the origin of cosmic rays below the knee, but the case is far from closed. There are several nearby supernova remnants, and if they were accelerating cosmic rays we could expect to see neutrinos from them.

Microquasars are a class of objects involving one compact object (a black hole or a neutron star) and a nearby star [15]. In many ways, x-ray binary systems behave like small versions of AGN including the existence of an accretion disk and the occurrence of jets. They can have strong gamma-ray emission extending up into the TeV. The environment around the binary involves strong magnetic fields (see Figure 2.2) and matter moving from the nearby star to the compact partner. As the jet extends into the surrounding medium, shocks are formed where particle acceleration can take place. Whether the content of the jets is electronic or hadronic is not known, but hadronic content can be uncovered if neutrinos are detected.

Additionally, it has been suggested that TeV neutrino production may accompany

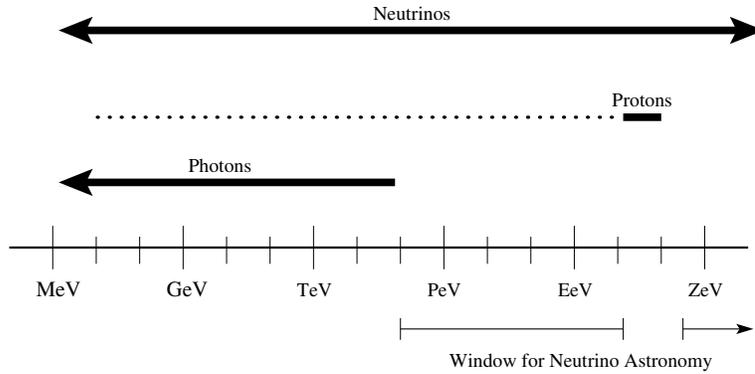


Figure 2.3: Schematic showing the usefulness of neutrinos, protons and photons for astronomical observations. The dark bands show where particles reach us unimpeded from their source. The dotted line shows that protons reach us, but do not point back to their sources. Also indicated is the window of neutrino astronomy where only neutrinos will reach us unimpeded from their source.

gamma-ray bursts (GRBs) [16]. GRBs are short-lived bursts of energetic photons that occur roughly three times per day somewhere in the sky. The duration of GRBs is bi-modal and can be used to sort bursts into two categories. Bursts shorter than two seconds are 'short' bursts and bursts longer than two seconds are 'long' bursts. The bi-modal nature of the burst duration suggests two populations of progenitors. The leading model for GRBs suggests that the GRB photons are produced in a relativistic fireball expanding from some catastrophic event like the collapse of a massive star or the merger of two compact binaries. In these collapses, a jet can form along the axis of rotation of the system and protons could be accelerated in this jet by the Fermi mechanism, leading to neutrinos by pion production.

2.4 Window of Opportunity

Neutrino observations are particularly useful above 100 TeV. But for a small window at extremely high energies where protons may reach us, neutrinos are the only particles above 100 TeV that can reach us unimpeded from their sources.

Figure 2.3 shows schematically the regions in energy in which certain messenger

particles are useful for astronomy. A particle is useful for astronomical observations, in this context, if it is stable and can reach us undeflected from the source over cosmological distances. The typical astronomical messenger particle is the photon, but photons above about 100 TeV cannot reach us from astrophysical sources through the galactic infrared background. Photons above this energy have enough center-of-mass energy to interact with background photons and pair produce electrons with a mean free path much less than the size of the galaxy. The universe is transparent to protons below about 10^8 TeV. It is only at 10^8 TeV that protons start to interact with the cosmic microwave background. The problem with protons, as has already been mentioned, is that they are so strongly bent in the galactic magnetic field that they do not point back to their sources below energies of at least 10^7 TeV. If the sources are extra-galactic then the effect of the intra-galactic magnetic field is uncertain. In any case detection of a point source of cosmic rays will be difficult because the fluxes at these extreme energies are so low.

This leaves a very large energy range which has been essentially unexplored by any astronomical observations. The failure to identify any objects active in this energy range is the chief limitation to understanding the source of cosmic rays. For a particle in this energy range to reach us unimpeded by intervening matter, it must be stable and cannot interact electromagnetically so as to avoid the CMB and the galactic magnetic field. The only known particle that can do this is the neutrino. Thus neutrinos are the best messenger particle above 100 TeV.

2.5 Diffuse Fluxes

One of the primary measurements to make in neutrino astrophysics is the shape of the neutrino energy spectrum integrated over all directions. Since the typical astrophysical spectrum is expected to be E^{-2} , astrophysical emission would show up

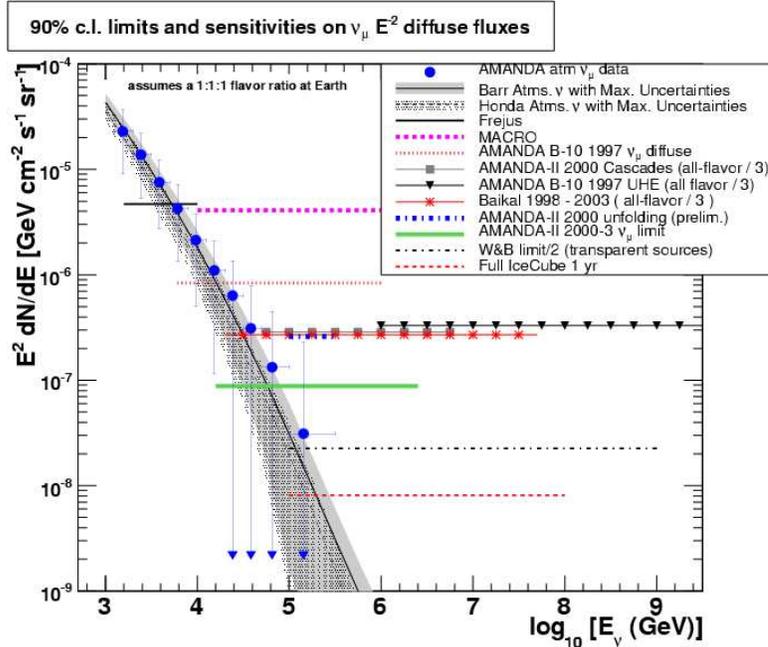


Figure 2.4: Flux limits for diffuse neutrinos assuming an E^{-2} spectrum. Shown are the muon neutrino limits and the all-flavor limits assuming a flavor ratio of 1:1:1 at the Earth. Included is the unfolded neutrino spectrum of muon neutrinos from 2000 AMANDA-II [17] and the limits from Frejus [18], MACRO [19] and AMANDA-B10 [20]. Also included are the all-flavor neutrino limits from AMANDA-II 2000 [21], AMANDA-B10 UHE selection [22], and the Baikal 6 year limit [23]. The Waxman-Bahcall upper bound [24] and the expected IceCube sensitivity after 1 year [10] of the full detector are included.

as a hard component on top of the steep $E^{-3.7}$ atmospheric spectrum. Such emission might be caused, for instance, by the sum of all GRB or AGN in the universe. Perhaps a single source is not strong enough to stick out on its own, but the sum of all of the sources together could give a high-energy glow of neutrinos that is detectable. To date, all measurements are consistent with a purely atmospheric spectrum but there are experimental limits on the diffuse high-energy neutrino sky. Figure 2.4 shows the current state of these measurements.

A few simple assumptions allow us to estimate the diffuse neutrino spectrum due to extra-galactic neutrinos. Since the acceleration of cosmic rays leads to the production of neutrinos, we can use the measured cosmic-ray flux to get a handle on neutrino production. This argument was first made in [24] as a case for an upper-

bound to the diffuse flux of neutrinos, but can be used also to get an idea of how big a detector must be to access astrophysical neutrino fluxes.

We can assume that, in a generic proton accelerator, the most abundant target is photons and that pion production occurs via the following reactions:

$$p + \gamma \rightarrow \Delta \rightarrow \pi^+ + n$$

In the case of π^0 production, the π^0 will immediately decay to high-energy gamma rays, making this object a source of gamma rays. In the case of π^+ production, the pion will decay to muon neutrinos, and the neutron, no longer contained, will escape the acceleration region, later decaying to a proton. In this simple model, these protons will be the cosmic rays that are observed. Protons that do not go through this process remain trapped in the acceleration region.

From the observed spectrum of cosmic rays, we can then calculate the generation rate of cosmic rays assuming that the spectrum at the source is E^{-2} . This calculation yields:

$$E_{cr}^2 \frac{d\dot{N}_{cr}}{dE_{cr}} = 10^{44} \text{erg} \cdot \text{Mpc}^{-3} \cdot \text{yr}^{-1} \quad (2.4)$$

The energy that the neutrinos get coming out of this source should be one quarter of the total proton energy because charged pions are produced half as frequently as neutral pions and the neutrinos carry half of the energy away in the decay of the pion. If the system is perfect and there are no other proton energy loss in the source, then the flux of neutrinos can be calculated [24] as:

$$E_\nu^2 \Phi_{\nu_\mu} = 2 \cdot 10^{-8} \text{GeV} \cdot \text{cm}^{-2} \cdot \text{s}^{-1} \cdot \text{sr}^{-1} \quad (2.5)$$

after accounting for redshift losses that the neutrinos would experience traveling from cosmological distances. This treatment neglects the ν_e produced when the neutron decays.

This neutrino flux is referred to in the literature as the Waxman-Bahcall bound

and is widely believed to be where we can start to expect a diffuse neutrino flux to start to appear.

With this flux we can estimate the order of magnitude of the number of neutrinos to trigger the full km³ detector in a year. A few simple assumptions are necessary. We will assume a reasonable value for the effective area to neutrinos of the full detector as 100 m² for neutrinos above 10⁵ GeV. This estimate comes from scaling the effective area determined in this thesis (See Section 4.5). Note that previous calculations have quoted the IceCube effective area as an effective area to muons [10] rather than an effective area to neutrinos. Our estimate is much smaller than the geometric area of the detector because of the low cross-section for neutrinos to interact. We will assume that we are looking only at neutrinos above 10⁵ GeV where the atmospheric neutrino background is small and that these neutrinos have a spectrum of E^{-2} . We are likely only to detect these neutrinos in the 0.3π sr band at the horizon. With these simple assumptions we get an estimate of about 6 neutrinos per year in the full IceCube detector. This rate is at the threshold of detectability and sets the km-scale as the appropriate scale for neutrino astrophysics.

Another view of this is seen in Figure 2.4. Shown is the expected sensitivity of the full IceCube detector to a diffuse source of neutrinos. After only a year of operation, IceCube will probe the diffuse spectrum below the Waxman-Bahcall bound.

2.6 Point Sources

A point source of neutrinos would be a particular astrophysical object that emits so many neutrinos that we would have a statistically significant excess of events in a small bin around that object. To date, no high-energy neutrino point sources have been identified. Figure 2.5 characterizes the current state of point source searches. Shown are limits achieved by the neutrino detectors AMANDA and MACRO. The

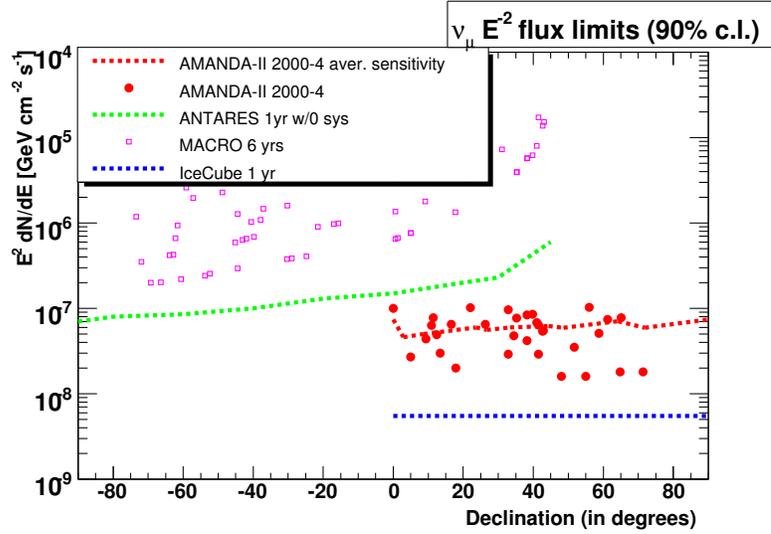


Figure 2.5: Shown are upper limits on E^{-2} point source fluxes (90% CL) vs. declination with the AMANDA 5-year point source search for selected candidates (red solid circles) [25]. The IceCube sensitivity for 1 year of the full detector is shown with the blue dashed line [10]. The MACRO upper limits for selected sources [26] and the ANTARES preliminary sensitivity for 1 year [27] are shown.

IceCube sensitivity is also shown and one year of the full IceCube array will push the sensitivity down almost an order of magnitude over the entire AMANDA-II dataset.

Note that the IceCube results mentioned here are from an early study [10] which did not incorporate the waveform information that is currently available and used only AMANDA analysis tools in the calculation. It is reasonable to expect the eventual IceCube performance to be better than shown here.

Chapter 3

Atmospheric Neutrinos

When high-energy cosmic-rays strike the atmosphere, the resulting air showers include a mix of hadronic particles. The particles π^\pm and K^\pm are produced and will decay to neutrinos if the π s and K s decay before they interact. The resulting neutrinos – known as atmospheric neutrinos – are the dominant source of neutrinos above 100 GeV. In the context of neutrino astrophysics, the study of atmospheric neutrinos is important because atmospheric neutrinos are a background to most potential astrophysical neutrino signals. Also, the atmospheric neutrino beam is the only experimentally-measured source of neutrinos in the detector and can be used to calibrate the response of the detector to astrophysical neutrinos.

Precise measurements of atmospheric neutrinos can serve as a probe of physical parameters that are not typically measured in terrestrial labs. The atmospheric neutrino spectrum is sensitive to the primary interaction physics as well as neutrino oscillation parameters.

This chapter will cover the current experimental and theoretical understanding of atmospheric neutrinos above 100 GeV.

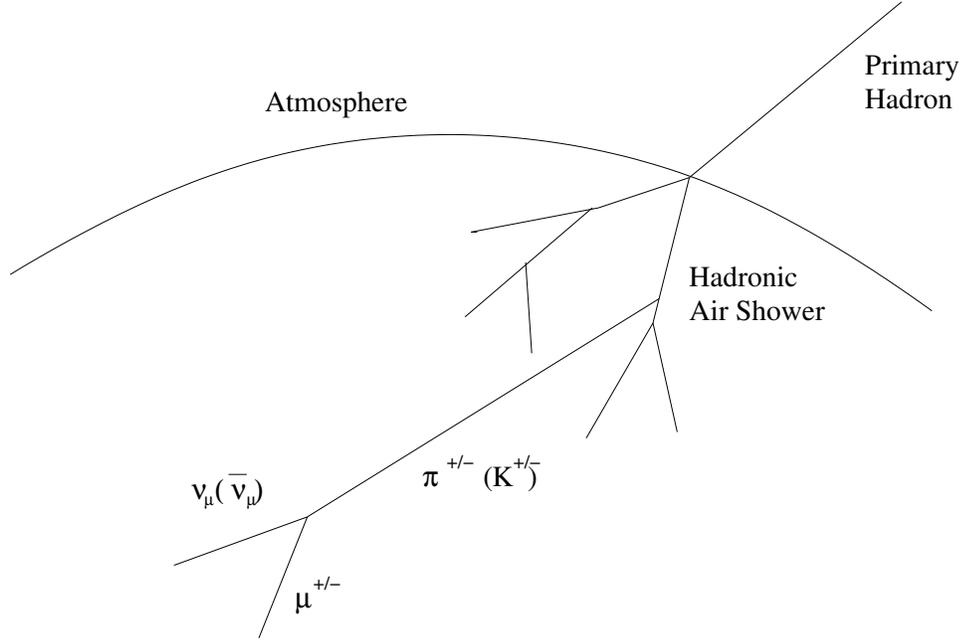


Figure 3.1: A schematic view of an extensive air shower. The primary hadronic interaction results in pion and kaon production. Neutrinos and muon production occurs when the pions and kaons decay.

3.1 Flux Expectations

3.1.1 Atmospheric Neutrino Production

The production of atmospheric neutrinos occurs primarily from the decay of π^{\pm} and K^{\pm} mesons produced in the air showers which result when a cosmic ray strikes the atmosphere (See Figure 3.1) via the reactions:

$$\pi^{+}(\pi^{-}) \rightarrow \mu^{+}(\mu^{-}) + \nu_{\mu}(\bar{\nu}_{\mu})$$

$$K^{+}(K^{-}) \rightarrow \mu^{+}(\mu^{-}) + \nu_{\mu}(\bar{\nu}_{\mu})$$

Neutrinos from the decay of the resulting muons are negligible because, at Ice-Cube energies, the resulting muons will have a lifetime sufficiently long that they will reach the Earth and interact before having a chance to decay.

3.1.2 Analytic Model

The measured atmospheric neutrino flux is a convolution of the primary spectrum with the production and decay of π^\pm and K^\pm . A treatment of all the relevant effects requires a Monte Carlo simulation, but an analytic calculation has been done [28] [6] and is instructive as to the major features of the flux. The result is of the form

$$\frac{dN_\nu}{dE_\nu} = \frac{N_0(E_\nu)}{1 - Z_{N,N}} \left(\frac{A_{\pi,\nu}}{1 + B_{\pi,\nu} \cos(\theta) \frac{E_\nu}{\epsilon_\pi^C}} + 0.64 \frac{A_{K,\nu}}{1 + B_{K,\nu} \cos(\theta) \frac{E_\nu}{\epsilon_K^C}} \right) \quad (3.1)$$

Here the zenith angle θ is taken to be the polar angle that the neutrino is coming from. A neutrino with $\theta = 0$ is straight down-going, and a neutrino with $\theta = 180$ is straight up-going. The factor $N_0(E_\nu)$ is the primary spectrum and at low energies, the neutrino spectrum follows the primary spectrum. The parameter $A_{X,\nu}$ governs the production of the particle X and the parameter $B_{X,\nu}$ governs its decay into neutrinos. The $Z_{N,N}$ constant is called a 'spectrum-weighted moment' and characterizes how much energy from cosmic-ray interactions goes into the production of more nucleons. At higher energies the produced π and K are boosted, and with a longer lifetime can interact in the atmosphere before they get a chance to decay. The parameter ϵ_X^C is the energy at which the transition from X -decay to X -interaction occurs. The energy at which this transition occurs is modified as a function of zenith angle because at larger zenith angles, the π s and K s have more low-density atmosphere to travel through and thus will decay even at high energies.

When the numerical values are used, Equation 3.1 becomes:

$$\frac{dN_\nu}{dE_\nu d\Omega} \sim 0.0286 E_\nu^{-2.7} \left(\frac{1}{1 + \frac{6 E_\nu \cos(\theta)}{115 \text{ GeV}}} + \frac{0.213}{1 + \frac{1.44 E_\nu \cos(\theta)}{850 \text{ GeV}}} \right) \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1} \text{ GeV}^{-1} \quad (3.2)$$

We expect for the atmospheric neutrino flux to follow the primary cosmic ray spectrum of $E^{-2.7}$ up to about 100 GeV, just where IceCube begins to become sensitive. Above this point the spectrum steepens and approaches $E^{-3.7}$ above about 1 TeV.

3.1.3 Monte-Carlo Models

There have been a number of attempts to calculate the atmospheric neutrino rate using simulation (for example [29], [30], [31]). These simulations must contend with measured uncertainties in the cosmic ray spectrum, the cosmic ray composition, and theoretical uncertainties in the primary interaction physics. Uncertainties in K production in the atmosphere also are important because of the relative enhancement of the neutrinos from K s when π s interact in the Earth.

One way to gauge the effect of the systematic errors on the predicted neutrino spectrum is to compare several neutrino flux predictions based on different assumptions. Figure 3.2 shows the atmospheric spectrum of three Monte Carlo models and the analytic model of [28] shown in Equation 3.2.

One of the dominant sources of error in these models is the primary cosmic-ray composition and energy spectrum. Though the relative shape of this spectrum is fairly well-known, the absolute flux differs systematically between experiments as seen in Figure 3.3. Furthermore, it is difficult to directly measure the composition at these energies and so composition estimates must be extrapolated from lower-energy measurements.

In addition, the physics of the primary hadron interactions in the atmosphere occurs at energies inaccessible to terrestrial labs. For this reason, the physics in these first few interactions must be extrapolated from lower energy measurements and contributes systematic error to the neutrino flux estimate.

Uncertainties in the primary cosmic-ray spectrum and the primary hadronic interaction physics are the dominant sources of error in calculations of the atmospheric neutrino rate and have been estimated to influence the overall normalization of the neutrino spectrum by $\pm 30\%$ [33].

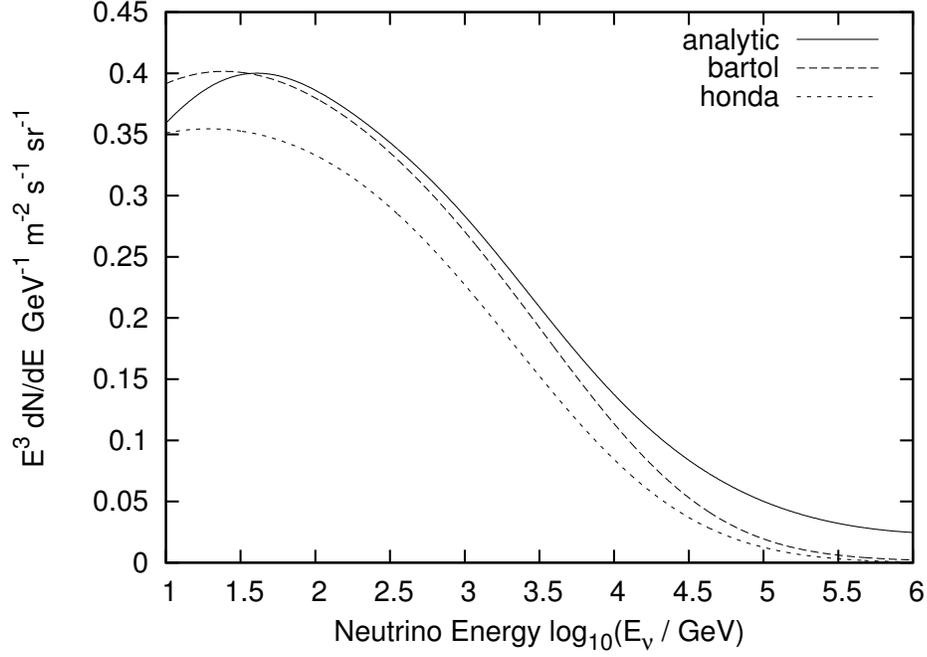


Figure 3.2: Up-going neutrino energy spectrum ($\nu_\mu + \nu_\pi$) integrated over all zenith angles. Shown are simulated models *bartol* [30], *honda* [31] and an analytic fit *analytic* [32]. Note the vertical scale compression.

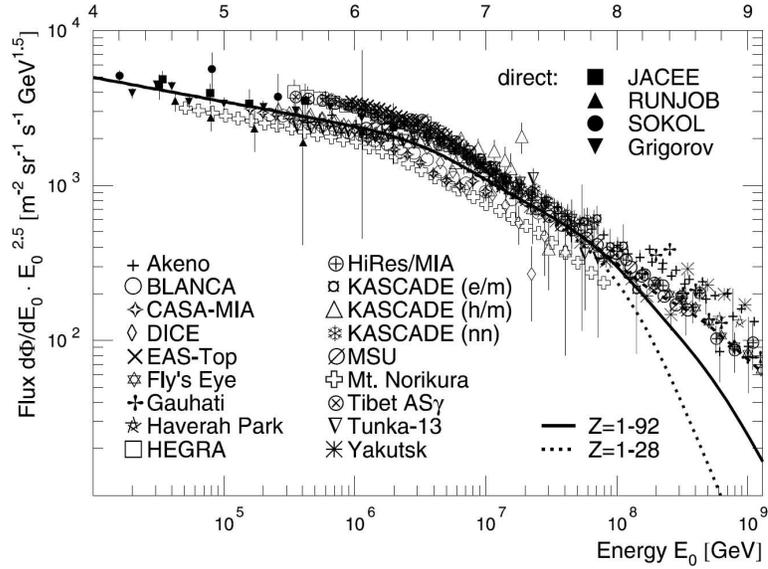


Figure 3.3: The knee of the cosmic ray spectrum as measured by a number of different air-shower experiments. For a full list of sources see [4] and references therein.

3.1.4 Prompt Neutrinos from Charm Production

The most uncertain component of atmospheric neutrinos comes from the decay of charmed hadrons. This component is 'prompt' in that the charmed particles decay very quickly. Being 'prompt', the parent particles of these neutrinos do not have time to interact before decaying until they reach extremely high energies. Uncertainties in the production of charmed particles are large. With the present exposure of IC-9, we do not have the statistics to resolve any of the prompt models which typically turn on at 10^6 GeV. At this energy, we have enough center-of-mass energy to create charmed mesons. For future analyses – especially analyses looking for a high-energy extraterrestrial diffuse neutrino source – the prompt charm component is important.

3.2 Current Experimental State

There are many experiments which access the known atmospheric neutrino population in some way. With the exception of AMANDA, most experiments to date can only access the neutrino spectrum out to about a TeV, where the fluxes are small and muons travel too far to be easily contained (and permit an energy measurement). Here we focus on two experiments which have spectrum measurements overlapping the expected IceCube energy range.

3.2.1 AMANDA

Before IceCube, the Antarctic Muon And Neutrino Detector Array (AMANDA) was the state-of-the-art in TeV neutrino detection. AMANDA is the prototype for IceCube and pioneered much of the technology and understanding that was needed to realize IceCube.

The general design of AMANDA is much the same as IceCube. AMANDA detects

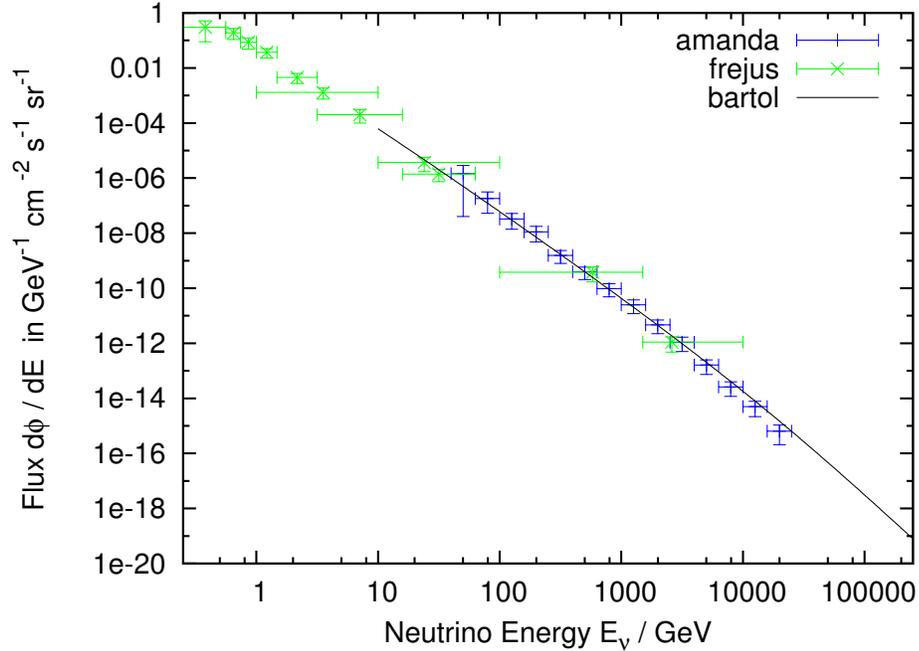


Figure 3.4: Results from the AMANDA [17] and Frejus [34] measurement of the atmospheric neutrino spectrum. Also indicated is the *bartol* atmospheric neutrino model of [30].

neutrinos by the Cherenkov light from leptonic secondaries. Optical Modules are deployed on strings in the deep ice below the South Pole.

The inner 10 strings of AMANDA form AMANDA-B10 which started acquiring data in 1997. In 2000, the detector was expanded with 9 more strings to form AMANDA-II. The AMANDA-II detector has been in operation since 2000 and has a combined neutrino dataset of some thousands of neutrino candidates, consistent with atmospheric neutrino predictions [25].

AMANDA has measured the atmospheric neutrino spectrum between 10^2 and 10^5 GeV by unfolding the measured muon energy based on the expected energy response [17]. The results of this measurement are summarized in Figure 3.4.

3.2.2 Frejus

The Frejus experiment was a 900 ton tracking calorimeter occupying a volume of 432 m³. The detector contained layers of iron planes with flash chambers and Geiger tubes interspersed. The flash chambers were long thin planes which alternated orientation with each layer in order to provide directional measurement. The detector was placed in an underground laboratory with 4000 hg/cm² of overburden near the border between France and Italy.

The Frejus collaboration made a measurement of the atmospheric neutrino spectrum up to about 1 TeV [34], consistent with both the AMANDA observation and with theoretical expectation. This result is summarized in Figure 3.4.

Chapter 4

The IceCube Detector

IceCube is a cubic-kilometer ice-Cherenkov detector being constructed between 1.4 km and 2.4 km below the surface of the Admundsen-Scott South Pole Station. The detector is complemented by an air-shower array, IceTop, at the surface. IceTop complements the InIce detector by assisting in the directional calibration and background rejection. IceTop-InIce coincident events can also be used for studying the cosmic-ray composition by measuring the muon content in cosmic-ray air showers. The InIce detector is optimized for detection of the TeV muon neutrinos anticipated from candidate cosmic-ray sources. The purpose of this chapter is to sketch the physical principles involved in the detection of muon neutrinos and to describe the design and calibration of the detector.

4.1 Detection Principle

A neutrino detector does not directly observe neutrinos. Cherenkov light is observed from the secondary particles created when a neutrino interacts inside or near the detector. IceCube is optimized for muon neutrino detection although it has some sensitivity to other flavors. In the case of IceCube and muon neutrinos, we are looking for Cherenkov light from relativistic muons. The strategy to detect these muons is to instrument a volume of some clear material with light-sensitive equipment. Because

neutrino interactions are rare, this instrumented volume must be large as argued in Section 2.5.

Furthermore, in order to successfully detect muon neutrinos, we must have some way to separate muons from neutrino interactions from the much more plentiful muons produced in cosmic-ray air showers. The strategy for this separation is two-fold. First, the detector is built far below the surface to help attenuate the cosmic-ray muons. Second, we look for muon events traveling upward through the detector. When an up-going event is seen, we know the primary particle was a muon neutrino because the muon must have been created no more than a few kilometers away from the detector. A cosmic-ray muon could not have penetrated more than a few kilometers through the Earth.

4.1.1 Neutrino Charged-Current Interaction

The first process to consider is the interaction of the neutrino near the detector. The standard model of particle physics includes a three-body vertex linking a ν_μ , a μ and a W^\pm , and a three-body vertex involving a W^\pm which can change an u quark to a d quark or vice-versa. These vertices admit the Feynman diagram in Figure 4.1 where a ν_μ or $\bar{\nu}_\mu$ can interact with a nucleon and produce a muon.

Also important to the detection of neutrinos is the the neutral current interaction in which a Z_0 interacts with one of the quarks in the detector (Figure 4.1). These Z_0 interactions can result in the neutrino losing energy as it moves through the Earth, since the incident neutrino is not absorbed.

The cross sections for these interactions are shown in Figure 4.2. The cross sections are extremely small, varying between 10^{-10} and 10^{-7} mb in the energy range of interest.

Also calculable is the expected difference between the direction of the incident neutrino and the resulting muon. At the TeV energies relevant for IceCube, the

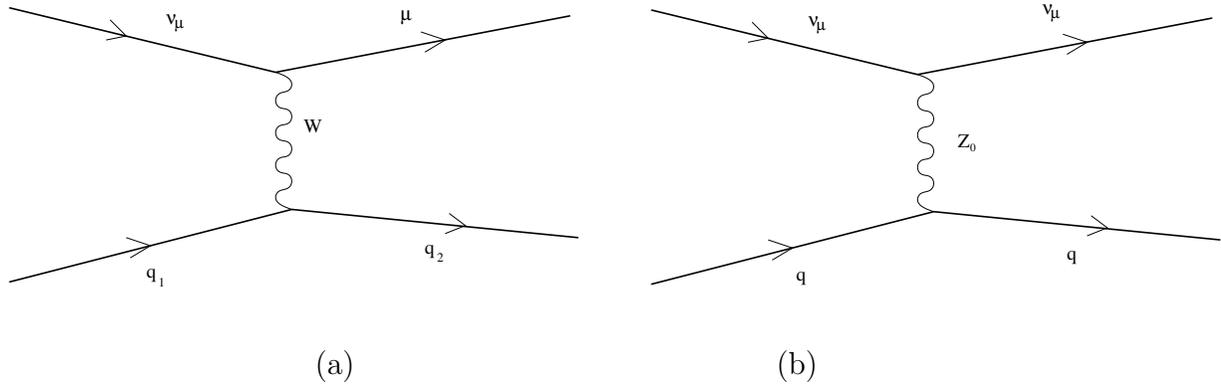


Figure 4.1: Feynman diagrams for two neutrino-quark interactions. Panel (a) is the charged-current interaction that results in a muon (which could be seen in IceCube). Panel (b) is a neutral-current interaction that results in energy loss of the ν_μ in the Earth.

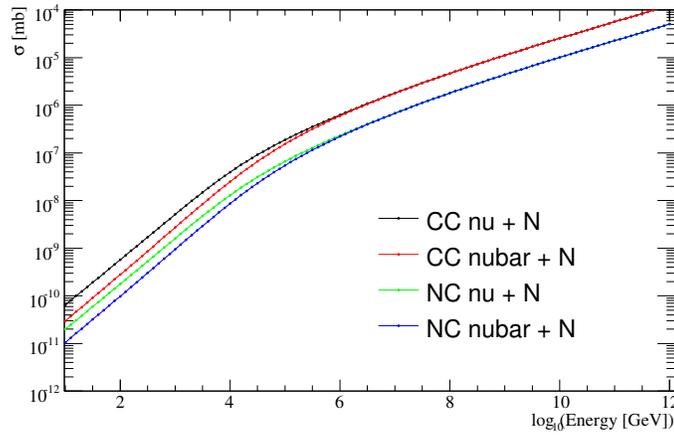


Figure 4.2: ν and $\bar{\nu}$ on Nucleon cross charged-current and neutral current interaction cross section, in the energy range of IceCube. The values are based on the CTEQ5 parametrization [35].

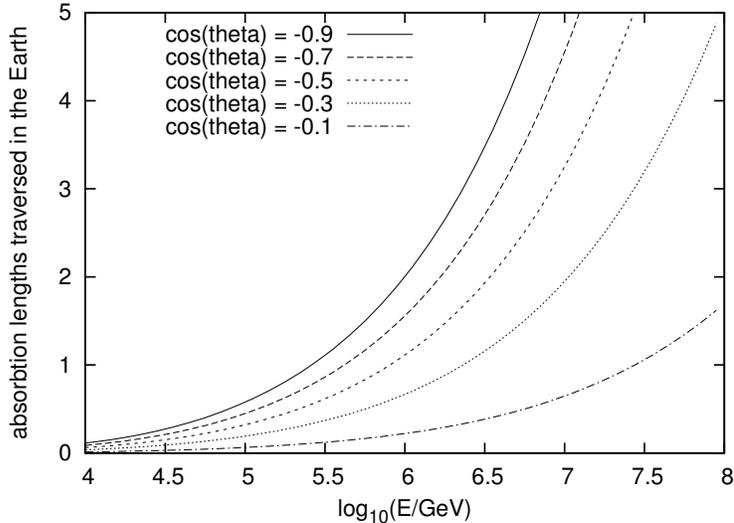


Figure 4.3: Absorption lengths traversed by a neutrino reaching IceCube as a function of the energy of the neutrino and for several different zenith angles. Since neutrinos of higher energy are attenuated more, we must look at the horizon to see the highest energy neutrinos.

deviation between the ν_μ and the resulting μ is less than a degree at a TeV and decreases at higher energies.

Figure 4.3 illustrates the behavior of neutrinos in the Earth. For a given zenith angle, a neutrino must pass through some amount of material to reach the detector. At a given energy, we can calculate from the interaction cross section the attenuation length of neutrinos of that energy. What is shown then is the number of attenuation lengths that must be traversed by a neutrino of that energy coming from that direction reaching the detector, as a function of zenith angle and neutrino energy. What stands out about this plot is that the highest energy neutrinos tend to range out when traveling through the Earth at high zenith angles. In order to detect these neutrinos, we must look at the horizon.

4.1.2 TeV Muon Interactions in Matter

Once a muon is created, it propagates through the ice. The propagation of muons through matter is well-understood and the light generated as the muon propagates

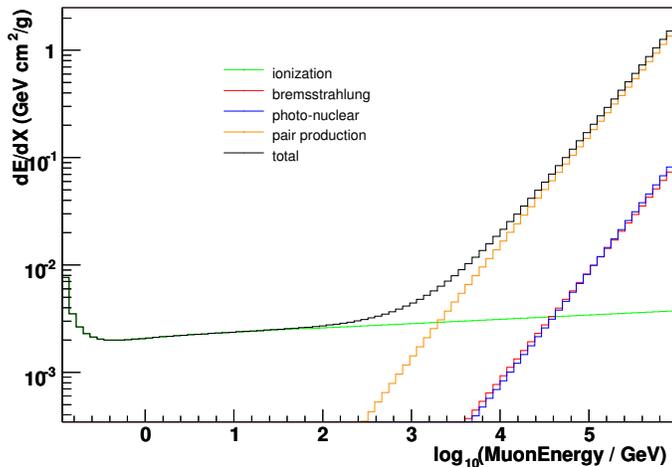


Figure 4.4: Muon energy loss in ice. Shown are the dominant energy loss mechanisms due to ionization, pair production, bremsstrahlung and photo-nuclear interactions. Ionization is dominant up to 100 GeV when stochastic energy losses take over. Obtained using MMC [36].

can be used to reconstruct the direction of the muon. Knowing the direction of the muon allows us to estimate the direction of the incoming neutrino to a degree or better.

The four dominant energy loss mechanisms relevant to TeV muons are described below. They are ionization, pair production, bremsstrahlung and photo-nuclear interaction. Also important in the detection of muon neutrinos, though the energy loss via this mechanism is negligible, is Cherenkov radiation.

To quantify how a muon loses energy in matter it is useful to look at $\frac{dE}{dX}$, which is the average energy lost per unit length in the material. It is important to note, however, that $\frac{dE}{dX}$ only quantifies averages. In truth, many of these mechanisms are stochastic and occur randomly, and that the propagation due to these stochastic events varies for different muons. Shown in Figure 4.4 are the relative contributions to the total $\frac{dE}{dX}$ for the different mechanisms. Generally speaking, ionization is important below 100 GeV and the stochastic mechanisms are important above this point.

Ionization and Cherenkov Radiation

As the muon travels through the ice at the South Pole, it scatters elastically with the various electrons and protons in the ice. Each of these interactions results in some loss of energy by the muon as it transfers energy to the atoms of ice. The energy loss due to these collisions is well-understood and modeled by the Bethe-Bloch equation, Equation 4.1.

$$-\frac{dE}{dx} = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} s - \beta^2 - \frac{\delta}{2} \right] \quad (4.1)$$

Of particular note in the Bethe-Bloch equation is the *density correction* term $\frac{\delta}{2}$. This term weakens the energy loss. This correction is due to an effective weakening of the field of the muon due to the polarization of the medium. When the muon is traveling faster than the speed of light in ice, this polarization adds coherently and some of the energy escapes as radiation at a fixed angle with respect to the path of the muon. This effect, named Cherenkov radiation after its discoverer, is useful because it is emitted uniformly along the track. As such, the Cherenkov cone can be used to reconstruct the direction of the muon.

Cherenkov light is emitted at a constant angle θ_C given by $\cos(\theta_C) = \frac{1}{n\beta}$ where β is the speed of the particle in units of c , and n is the index of refraction of the medium. In ice this angle is $\theta_C = 40.7^\circ$.

Pair Production

Pair production is a process by which the muon interacts with an atom in the medium to produce an e^\pm pair. The e^\pm pair quickly re-interacts to produce an electromagnetic shower in the medium. Unlike ionization energy loss, pair production is a stochastic process occurring at random points along the muon's trajectory.

Bremsstrahlung

Bremsstrahlung is a process in which the muon interacts with a nucleus to produce gamma rays. Like in pair production, the gamma rays quickly re-interact and produce an electromagnetic shower. Bremsstrahlung is stochastic and results in single catastrophic events of energy loss along the muon's trajectory.

Photo-nuclear Interaction

In addition to electromagnetic processes, a muon can lose energy by producing hadrons. Like bremsstrahlung and pair production, the photo-nuclear interaction is stochastic, but the resulting shower is hadronic rather than electromagnetic.

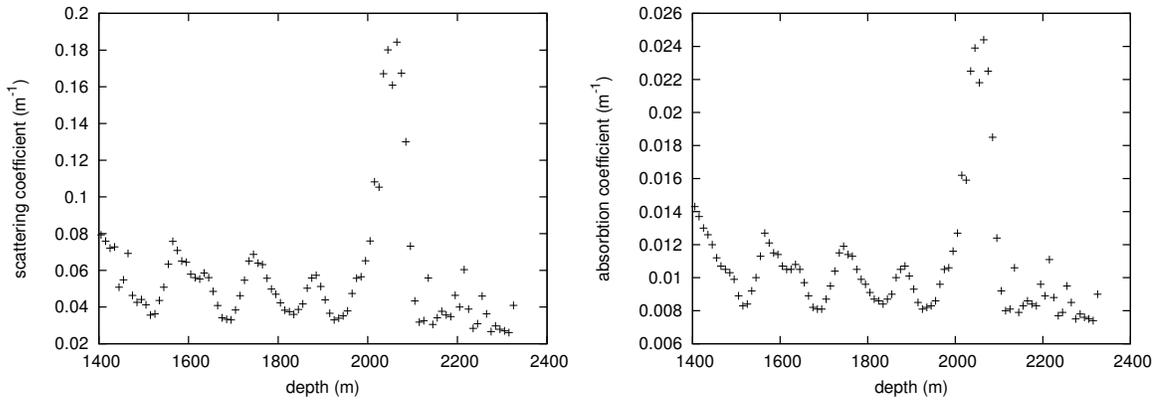
Muon Range

Looking at the energy loss of a muon in ice allows us to characterize the average behavior of muons and get an idea of the scale required to measure muons of these energies. Muons at very high energies (above 100 GeV) lose energy quickly and primarily by stochastic events of inelastic particle production. These events show up as discrete showers. The bremsstrahlung and pair-production events result in electromagnetic showers and the photo-nuclear interactions result in hadronic showers. When the muon reaches ~ 100 GeV, these stochastic events give way to ionization as the dominant form of energy loss. From about 100 GeV down to about 1 GeV, the muon energy loss is roughly constant at ~ 2 MeV / cm. Here the energy loss is quasi-continuous so it is legitimate to estimate the range of a 100 GeV muon in ice as about 500 meters. Muons above this energy lose energy stochastically until their energy drops below 100 GeV, at which point they coast 500 meters. Above 100 GeV, we can quote an average range, but the range varies drastically, muon-to-muon, due to the stochastic nature of the processes involved.

4.1.3 South Pole Ice

As the ice at the South Pole built up, dust was deposited at the surface. Varying geological conditions over the years resulted in different dust contamination at the surface. Today these dust layers give a depth dependence to the optical properties of the ice. These optical properties have been measured using light sources deployed along with the AMANDA strings [37]. Since scattering in the detector is the dominant effect, far from a light source the light propagation is purely diffuse and allows measurement of the ice properties. The presence of these dust layers results in depth-dependent scattering and absorption of light. The absorption length is around 100 meters which means that light can be seen quite far away from its source, but the scattering length averages about 20 m. This is on the same scale as the separation of the DOMs, which means that the light that we observe is strongly scattered.

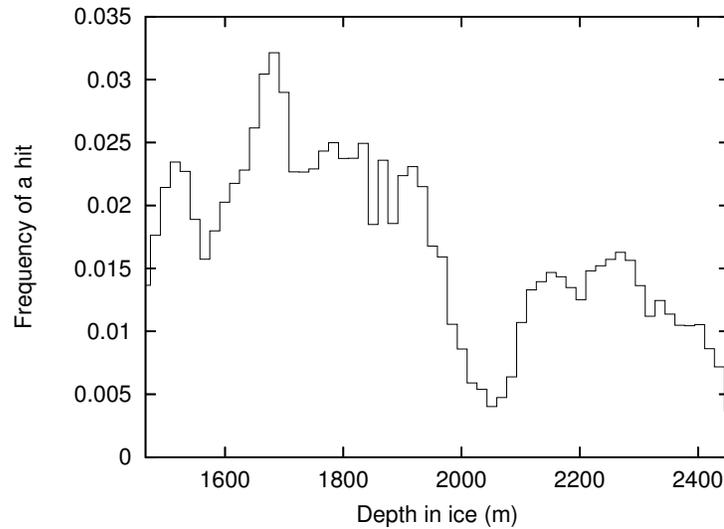
The ice has a dramatic effect on the behavior of the detector, as can be seen in Figure 4.5. Figure 4.5a is a plot of the scattering coefficient as a function of depth in the ice. Figure 4.5b shows the absorption coefficient. Several peaks are visible as well as a particularly strong peak at a depth of 2000 m where the scattering length drops as low as 4 meters. Figure 4.5c is a plot of the occupancy of DOMs in events that trigger the detector. Since the majority of events which trigger the detector are down-going muons, we see the expected general falloff as we go down in the ice and down-going muons range out. Convolved with this general downturn are several strong features which line up with the dust layers. Of particular note is the dust layer at a depth of 2000 meters where the high scattering and absorption coefficients prevent much of the light from being detected in that layer.



(a)

(b)

(a)



(c)

Figure 4.5: Panel (a) shows the scattering coefficient for light of 400 nm as a function of depth from the surface. Panel (b) shows the absorption coefficient. The optical properties are taken from [37]. Panel (c) shows the occupancy of IceCube DOMs. As a function of DOM depth, what is shown is the frequency with which DOMs at that location participate in events. The behavior is a general downward turn consistent with air-shower muons ranging out. Superimposed on this general trend, the structure of the dust layers is evident.

4.2 The IceCube Detector

The IceCube detector is an array of light-sensitive Digital Optical Modules (DOMs) embedded in a cubic kilometer of ice below the Admundsen-Scott South Pole Station at the geographic South Pole with an air-shower array at the surface [38] [10]. Each DOM contains a 10-inch Photo-Multiplier Tube (PMT) and electronics for acquiring and digitizing voltage pulses from the PMT. Figure 4.6 shows the detector schematically.

To deploy DOMs in the ice, a pressurized hot water drill melts a hole in the polar ice pack. DOMs are deployed on a long cable which both supports the weight of the DOMs in the water until they freeze into place, and carries communication to the surface.

The detector design consists of 80 vertical strings of 60 DOMs each. The DOMs are spaced 17 meters apart on the cable that leads to the surface. DOMs are deployed starting at a depth of 1450 meters and continuing to a depth of 2450 meters. In addition to screening down-going muons, the detector must be this deep in order to get past particularly bad ice that resides in the first kilometer of ice. Strings are arranged (from above) in a hexagonal pattern with each string spaced 125 meters from the 6 strings nearest it (See Figure 4.7). The spacing of the strings was the result of optimizing the geometry for the detection of TeV neutrinos. All told, the geometrical instrumented volume will be about 1 km^3 .

In addition to the instrumentation below the ice, DOMs are deployed at the surface in tanks of frozen water to form an air-shower array called IceTop. Four DOMs are deployed at the surface at each string location. The DOMs are deployed in two tanks, each containing two DOMs. The two tanks (four DOMs) at the top of each string are collectively known as an IceTop Station.

The IceCube array is being constructed surrounding the Antarctic Muon And Neutrino Detector Array (AMANDA) which served as a prototype for the IceCube ex-

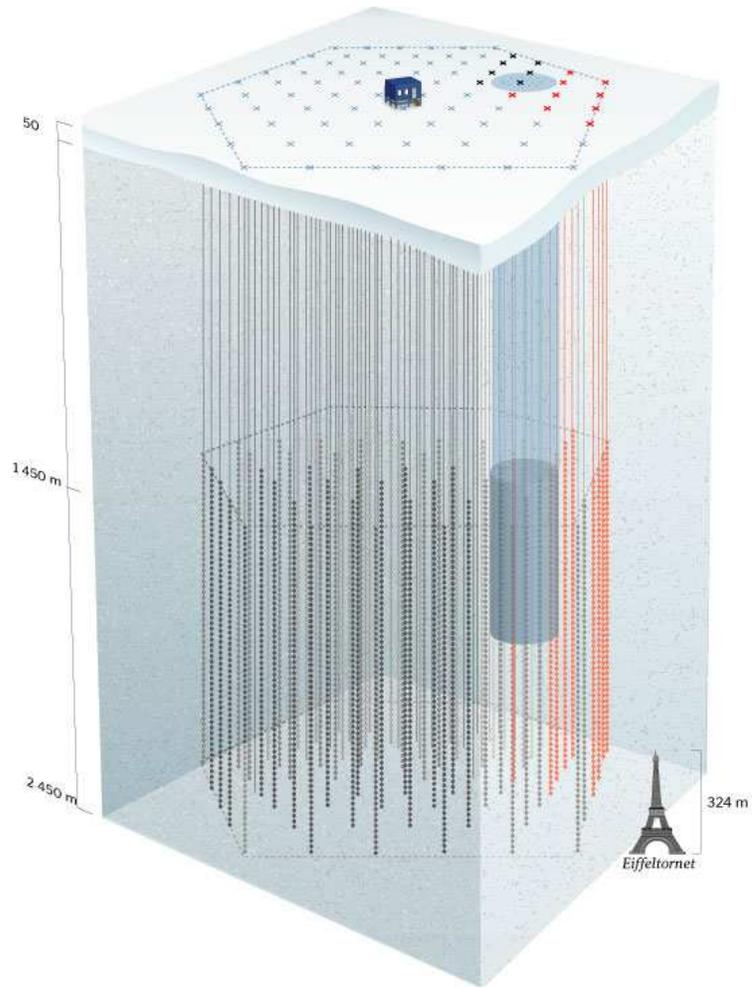


Figure 4.6: A schematic of the IceCube detector. Identified in red are strings which were deployed as of the writing of this thesis (2006). The dark black X's at the surface are the string locations which have IceTop stations deployed but with no corresponding InIce string yet deployed. Also visible is the footprint of the AMANDA experiment which is partially contained by the IC-9 array.

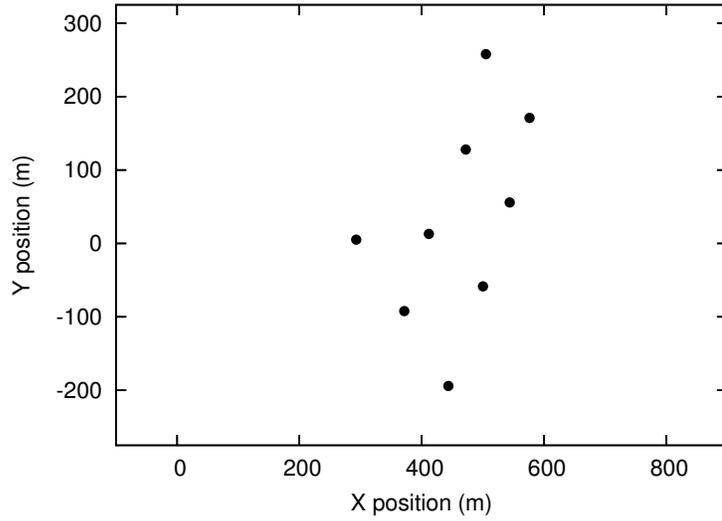
periment. Though the completed IceCube will be significantly larger than AMANDA, the density of strings is lower, affecting IceCube's sensitivity to low-energy (100 GeV and lower) muons. Using AMANDA as a dense sub-array is expected to extend the low-energy reach of IceCube. In the 2007 season, the DAQ and processing systems of the two detectors will be combined to form a single detector.

The InIce DOMs are connected to the surface via a twisted-pair cable which simultaneously powers the DOM and carries digital communications to and from the DOM. These cables are gathered together with the IceTop DOMs at the surface and run to the counting house located in the center of the array. The counting house holds computers which communicate with the DOMs and run the Data Acquisition System (DAQ).

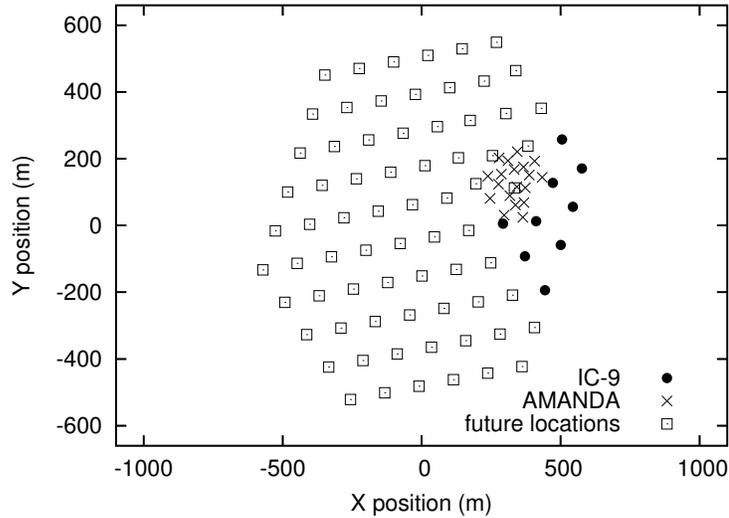
The DOMs are identified by a unique pair of numbers: the string number and the OM number. String numbers run from -19 to 80 (AMANDA strings are identified with a negative string number) and OM numbers run from 1 to 64. OM 1 is closest to the surface and OM 60 is at the bottom of the string. DOMs 61 - 64 are the four IceTop DOMs for a particular string.

Construction of the detector began in the austral summer 2004-2005 with the deployment of the first IceCube string and four IceTop stations [38]. Eight more strings were deployed in the 2005-2006 season with twelve IceTop stations deployed. The plan calls for 12 strings to be deployed in the 2006-2007 season and then 14 strings each year after that until the detector is completed in 2010-2011.

The data used in this thesis is acquired from the detector in its state following the 2005-2006 deployment season, referred to as IC-9. Shown in Figure 4.7a are the locations of the currently deployed IceCube strings and the geometry of the final planned detector is in Figure 4.7b.



(a)



(b)

Figure 4.7: String locations for the completed IceCube detector. Shown are the deployed strings which comprise the IC-9 detector, and the locations of future strings. Also indicated are the positions of AMANDA strings. Panel (a) shows the IC-9 strings used in this analysis. These strings were in place after the 2005-2006 deployment season. Panel (b) shows the projected completed detector along with AMANDA. The origin for the coordinate system is the center of the future array. IceTop stations have been omitted.

4.2.1 IceCube DOM

The Digital Optical Module (DOM) (see Figure 4.8) is the building block of the IceCube detector. The DOM is a complicated piece of hardware and a full description is beyond the scope of this discussion. For more details the reader is directed to [38].

The DOM is a stand-alone data acquisition computer. A DOM consists of a Photo-Multiplier Tube (PMT) encased in a pressure sphere. Included in the DOM main electronics board are components to acquire, digitize and transmit digital signals to the surface, as well as components to assist in the calibration of the detector. Figure 4.9 is a block diagram of the DOM internals. Two Analog Transient Waveform Digitizers (ATWDs), custom application-specific integrated circuits (ASIC), digitize the PMT waveform with fine-binned timing resolution, and a commercial Fast Analog to Digital Converter (fADC) digitizes the waveform with coarse timing resolution. An on-board Field-Programmable Gate Array (FPGA) is used to control mainboard components, and a 40 MHz free-running clock provides local timing. Much of the DOM functionality is implemented as FPGA firmware. This firmware as well as the software running on the DOM can be changed remotely after the DOMs are deployed, allowing for future upgrades. As power at the Pole is somewhat limited, the DOMs maintain a low power consumption of about 3.5 W.

PMTs

A photo-multiplier tube (PMT) is a piece of hardware that is used in many experiments that require low-intensity light measurements because of their ability to precisely measure the arrival time of individual photons.

Figure 4.10 shows schematically how a PMT works. The photo-cathode is a thin (to the point of transparency) alkali layer covering the inside of the glass hemisphere. The photo-cathode has a very low work function so that visible light photons can liberate an electron into the vacuum of the PMT. In the base of the

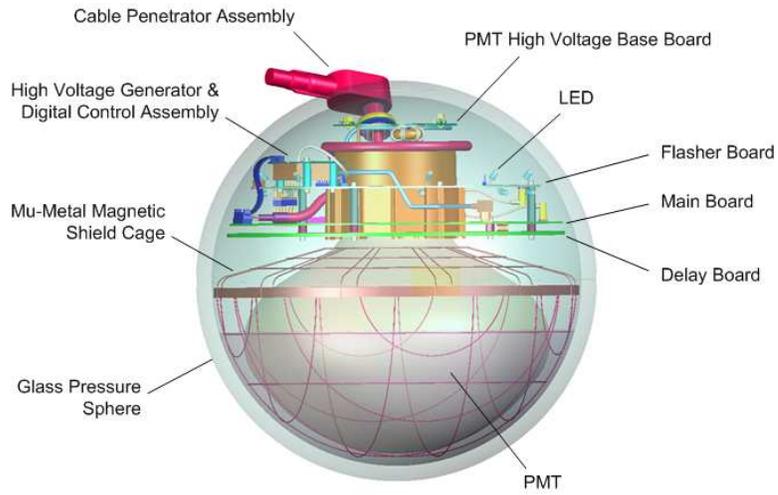


Figure 4.8: A schematic of an IceCube DOM

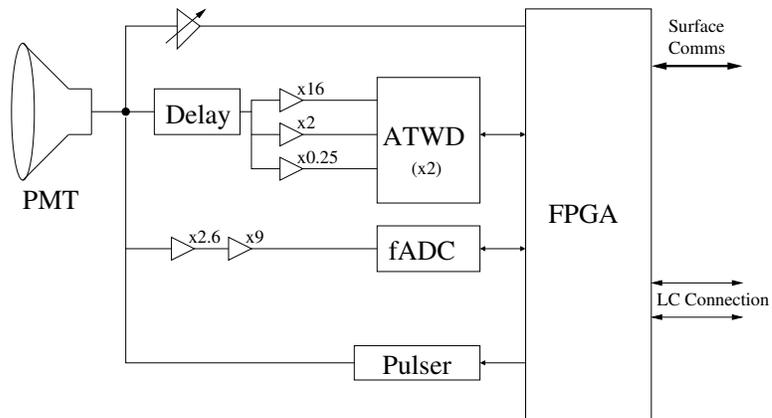


Figure 4.9: Block diagram of the IceCube DOM. Note the topology of the PMT, the discriminator, the ATWDs, fADC and the FPGA. Some components have been omitted for simplicity.

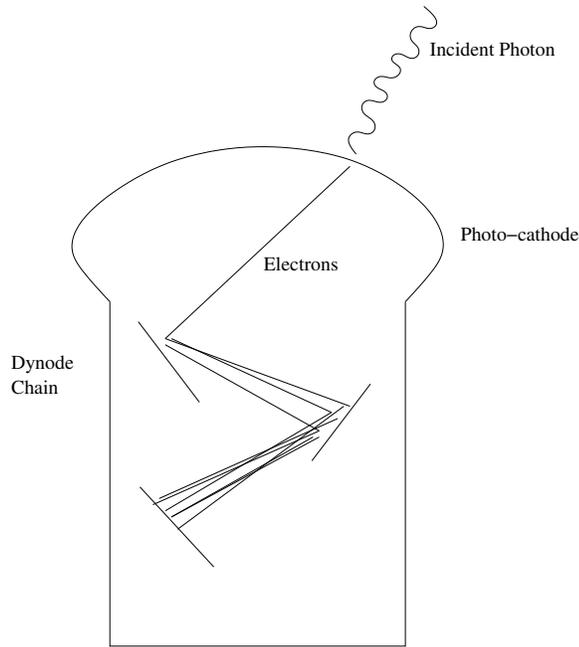


Figure 4.10: A schematic of a PMT.

PMT, there is a series of metal plates, the dynodes. These plates are given very high potential (typically ~ 1200 Volts) with respect to the photocathode which is at ground.

When a visible light photon is incident on the photocathode, there is a chance (of order 20%) that it will knock a single electron (a 'photoelectron' or PE) out of the photocathode and into the vacuum of the PMT. The high potential on the first dynode causes the electron to be accelerated toward the dynode chain. It strikes the first dynode with enough energy to liberate a few more electrons. These electrons in turn are accelerated toward the second dynode and the result is a large amplification in the number of electrons by the time the last dynode is reached so that the final current can be measured. At low PE amplitudes, the final current is proportional to the number of PEs emitted. The average number of electrons to come out of the final dynode per incident PE is the gain of the PMT.

The PMT for IceCube DOMs is a HAMAMATSU R7081-02 25 cm diameter PMT, with 10 dynodes. During 2006 data-taking, the PMTs were operated at a

gain of 10^7 and a high voltage of between 1200 - 1400 Volts. The typical dark noise rate for a deployed IceCube DOM is 300 Hz.

Triggering a Readout

The PMT is connected to a pair of discriminators. For the 2006 data set, only one of these discriminators is used. When the signal on the PMT goes over a defined threshold (typically 0.3 PEs), the discriminator fires and a waveform capture can occur. The capture is initiated by the on-board FPGA which causes the ATWDs and the fADC to capture the PMT waveform. Once the waveform is captured, the waveform is digitized only if a local coincidence occurs. Local coincidence will be described later.

The time required to processing triggering signals through the various mainboard components requires that the ATWD waveform capture must be delayed. This delay allows the full waveform to be captured starting before the waveform begins to rise. This delay is achieved by including 75 ns delay line between the PMT and the ATWD. Since the waveform capture process only begins on a clock heartbeat, the arrival of the pulse within the capture window of the ATWD varies uniformly over 25 ns depending on the phase of the clock when the PMT pulse arrived. Figure 4.11 shows the effect of this clock heartbeat on the DOM readout times.

ATWD and fADC

The Analog Transient Waveform Digitizer (ATWD) is a custom switched-capacitor array containing four channels each capable of sampling and digitizing a voltage waveform [39]. The capture and digitization of a waveform by the ATWD occurs in two steps. On a 'start' signal from the FPGA, each channel of the ATWD captures and stores the waveform. After the waveform is captured, the captured waveforms from the ATWD channels are sequentially digitized. 128 samples each of 10-bit depth

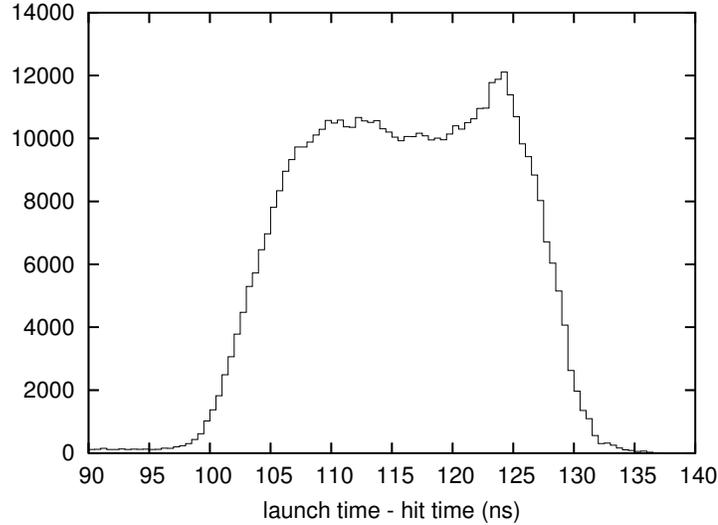


Figure 4.11: Time between the time stamp of the ATWD readout and the feature-extracted first hit. The hits arrive systematically earlier than the capture time due to a PMT transit time correction by the feature extraction algorithms. The time of the hit varies roughly uniformly over 25 ns from the time of the DOM capture because a capture can only be initiated on a 25 ns clock heartbeat.

are acquired from each channel. The ATWD sampling begins on one of the beats of the 40 MHz clock counter, and this time is latched as a time stamp of the readout event. The first three ATWD channels are reserved for acquiring the waveform from the PMT. The last channel is used for calibration and does not participate in regular data taking. Each of the three PMT-acquiring ATWD channels is run at a different gain, nominally 16x, 2x, and 0.25x. Typically only the high-gain ATWD channel is used for reconstruction unless it saturates. If the high-gain ATWD channel saturates, then the highest gain non-saturated channel is used to measure the waveform in the saturation region. The front end gains were arranged so that even if the PMT is saturated, thousands of PEs, the lowest gain ATWD channel will not saturate.

The ATWD channels on a particular ATWD all sample the waveform at the same rate which can be set between 1.6 and 5 ns/sample. For all data considered here, the ATWDs are sampled at a nominal rate of 3.3 ns/sample. Furthermore, each DOM contains two ATWDs. The second ATWD can potentially acquire and digitize

a signal 2 clock cycles after the end of the fADC trace, but all readouts from the second ATWD are ignored in this initial processing.

Note that though the timing on the DOM is limited to the rather coarse 25 ns clock period, by fitting the ATWD waveform, measurement of PE arrival times can be made to a fraction of the ATWD sample width. The onboard clock is used for coarse timing, and the ATWD is used for fine timing. All fine timing measurements are done offline.

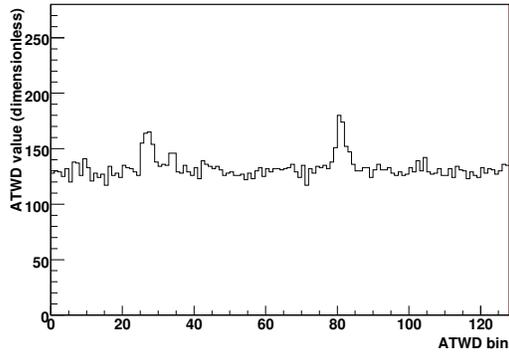
In order to further extend the time over which waveforms can be acquired, the mainboard also contains a commercial Fast Analog to Digital Converter (fADC). The fADC has the advantage that it can acquire samples as long as $6.4 \mu s$, but it is limited to 256 samples. The ATWDs are used for fine-binned recording of the early part of the waveform where the timing is most important. The fADC also digitizes the early part of the waveform seen by the ATWD (though at a coarse timing resolution), but is primarily intended to catch late light arriving after the ATWD window.

For the 2006 dataset, the second ATWD begins reading out after the fADC stops recording data, though the data is not typically used. Furthermore, in order to reduce data bandwidth of the DOMs to the surface, for ATWD channels 1 and 2 only 32 bins are read out instead of the full 128 bins. We acquire the full 128 bins of the highest gain ATWD channel, channel 0.

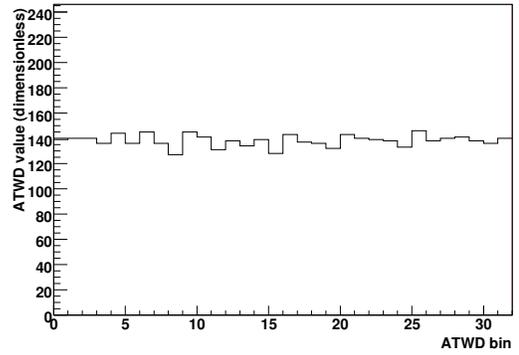
Figure 4.12 shows an example of the three ATWD records as well as the final calibrated waveform from the DOM. Calibration of the ATWDs is described in Section 4.4.2.

Beacon Hits

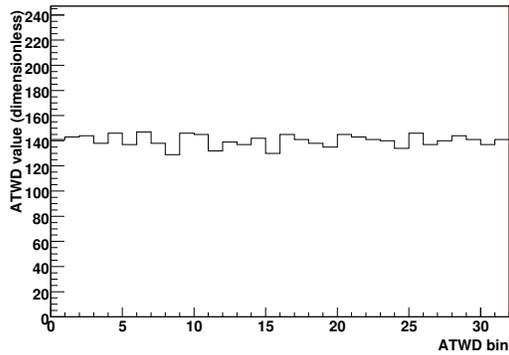
Periodically (about 1 Hz) the IceCube DOMs will trigger a readout without requiring a local coincidence and without requiring the discriminator to fire. During these readouts, there may or may not be any activity in the DOM's PMT, but the signal



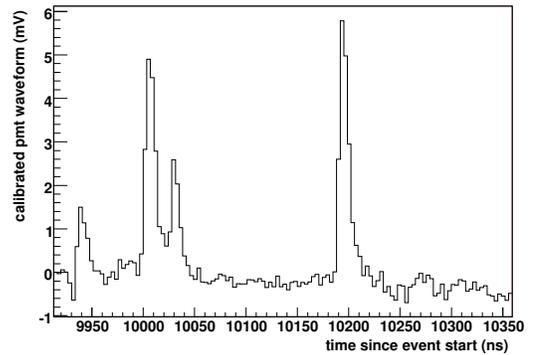
(a)



(b)



(c)



(d)

Figure 4.12: Example uncalibrated and calibrated ATWD waveform readouts. Panel (a), (b) and (c) are the raw digital ATWD readouts for channels 0 1 and 2 respectively. Panel (d) is the final calibrated waveform. The fADC trace is omitted because it is not important for this analysis.

from the ATWD is digitized and transmitted anyway. These readouts are termed 'beacon hits' and are used by the DAQ system to maintain continuous contact with DOMs. These DOM readouts are not used in the analysis of physics events.

Local Coincidence

The DOM waveform is only digitized and sent to the surface if the readout event satisfies a Local Coincidence (LC) condition. This condition says that a DOM is readout only if either the DOM above it or the DOM below it on the string also had a readout event within ± 1000 ns (See Figure 4.13). There are a few rare identified cases where a DOM can experience a readout without an LC. If the hit difference time is very close to the LC window, then one DOM can be read out without its partner because the LC windows up and down are not necessarily perfectly symmetric. Furthermore the neighboring DOM can have a beacon hit or be hit while it is recovering from an earlier readout.

This LC condition suppresses almost all of the isolated noise hits in the detector. Because each readout actually contains quite a bit of data (~ 300 bytes) from the waveforms, the bandwidth to the surface must be suppressed and the LC condition suppresses the readout rate for a particular DOM from ~ 300 Hz to around 6 Hz as seen in Figure 4.14.

The implementation of the LC condition is done in the DOM firmware and can be changed when as the firmware and acquisition software matures. Further improvements to the DOM firmware and DAQ software are anticipated. These improvements will implement a Soft Local Coincidence (SLC) condition which will recover the DOM hits lost to LC by omitting the full waveform and sending summary information when LC is not satisfied.

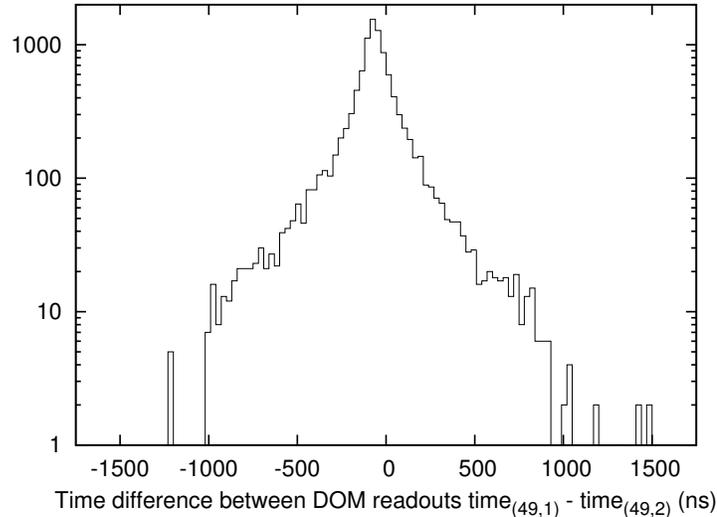


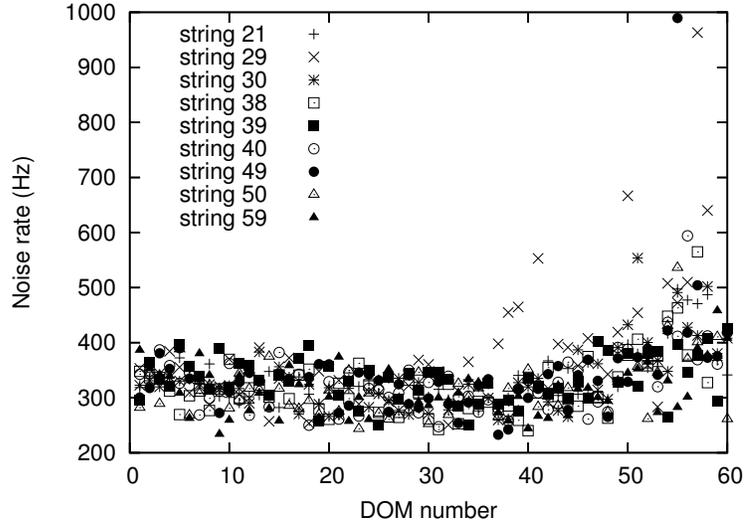
Figure 4.13: The time difference between readouts on the top DOM on a string, DOM (49,1) and the DOM below it, DOM (49,2). The distribution cuts off sharply at ± 1000 ns due to the Local Coincidence condition. The peak is offset from 0 because the top DOM is overwhelmingly the earliest hit.

4.2.2 Data Acquisition

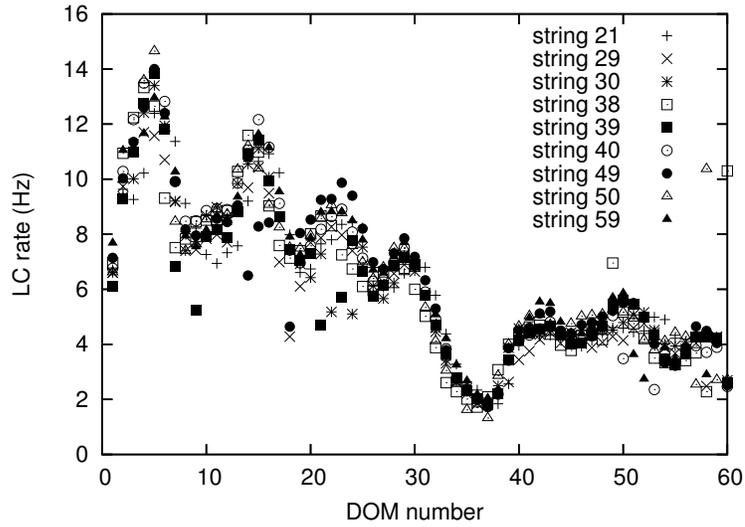
The Data Acquisition system (DAQ) is a collection of hardware, firmware, and software components for acquiring data from the DOMs and assembling the DOM readouts into events based on flexible trigger criteria. See Figure 4.15 for a schematic overview of the system.

For each string, a dedicated computer, the DOMHub, houses several custom PCI cards (the DOM Readout or 'DOR' cards) which perform the communication with, and supply power to, the DOMs. Several other computers called the String Processors are providing the rest of the DAQ with a well-ordered stream of data from the DOMs attached to it. The String Processor maintains the time calibration records for each DOM and performs the calibration for all readout and monitoring events.

Lightweight summary information from each readout event (essentially just the coarse time stamp of when the ATWD was read out) is forwarded from the String Processors to a trigger process which applies several criteria for selecting the events



(a)



(b)

Figure 4.14: Noise rates for IceCube DOMs. Panel (a) shows the raw noise rate and panel (b) shows the rate of local coincidences. The local coincidence readouts are due primarily to real muon events and so the structure is due to the local ice properties.

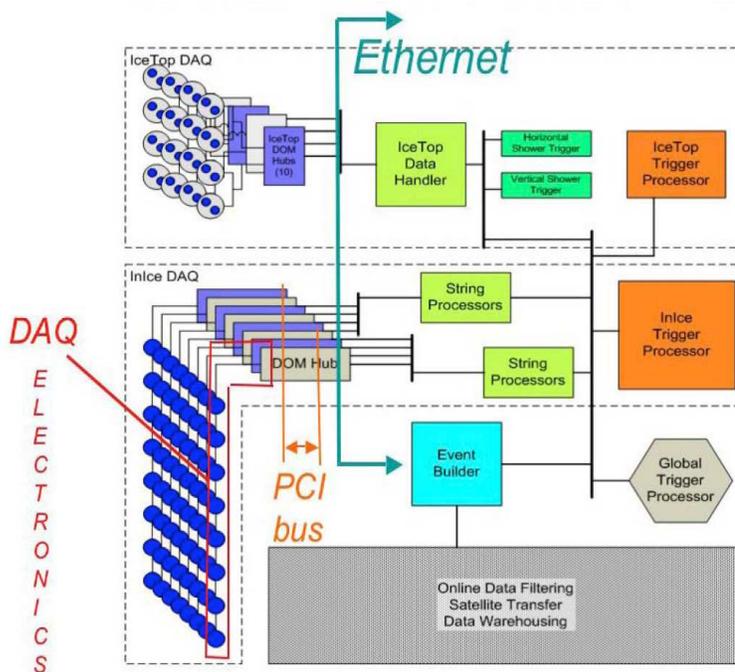


Figure 4.15: Overview of the DAQ components.

and forwards trigger decisions on to the Event Builder.

The Event Builder collects the various triggers and groups them together into events, then retrieves from the String Processors the full readouts of those DOMs which participate in the event.

4.2.3 Trigger

There are a variety of trigger conditions that can be used to identify events in the detector. For the purposes of this analysis, only one is important: the InIce Simple Majority Trigger (SMT). This trigger requires that 8 of the deployed DOMs were read out and transmitted to the surface, with their hits occurring in a $5 \mu\text{s}$ time window. This implies that all DOMs participating in the trigger had a Local Coincidence (see Figure 4.16) with another DOM. If the trigger is satisfied, then a trigger window is formed which is the union of all time windows in which the trigger condition was satisfied. All DOM readouts in a $16 \mu\text{s}$ window around this trigger window are

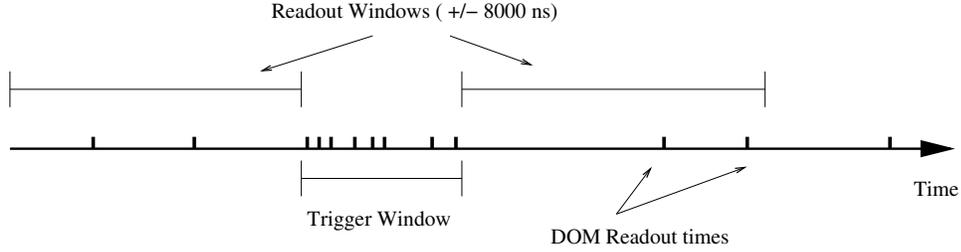


Figure 4.16: A hypothetical InIce Simple Majority Trigger. Shown are the individual DOM readout times in an event. A sliding time window of 5000 ns is applied. If 8 or more DOMs are read out in that window, then a trigger is formed. The trigger window consists of the time over which the 8 DOM threshold is satisfied. In addition to the trigger window, all DOMs in a 16000 ns readout window (8000 before the trigger and 8000 ns after the trigger) contribute to the event.

included in the constructed event. The trigger window is set this wide in order that any activity in IceTop (or activity in the deep detector for IceTop triggers) associated with the event is contained in the event.

Recall that the triggering is done in software, so there is no dead time resulting from a trigger and it is possible that two or more readout windows overlap. In this case, the final event contains all readouts from the union of the readout windows from all the satisfied triggers. This is important because there is no single trigger window, and in the analysis of these events we need to be conscious of the fact that several triggers may be present in an event and that interesting events can fall at any time during the readout window.

At this point, it is worth noting the scale of things. The rate of events satisfying this trigger is a manageable 143 Hz for the 9-string detector. For comparison, the expected rate of muons from atmospheric neutrinos is 10^6 times smaller at $\sim 6 \cdot 10^{-4}$ Hz. The rates in the full detector are expected to be roughly 10x larger.

Figure 4.17 shows a nice example of an event in the detector. The colors indicate the timing of the hit DOMs with red DOMs hit early and blue DOMs hit late. The results of the linefit, dipolefit and likelihood reconstructions (described in Chapter 5) can be seen.

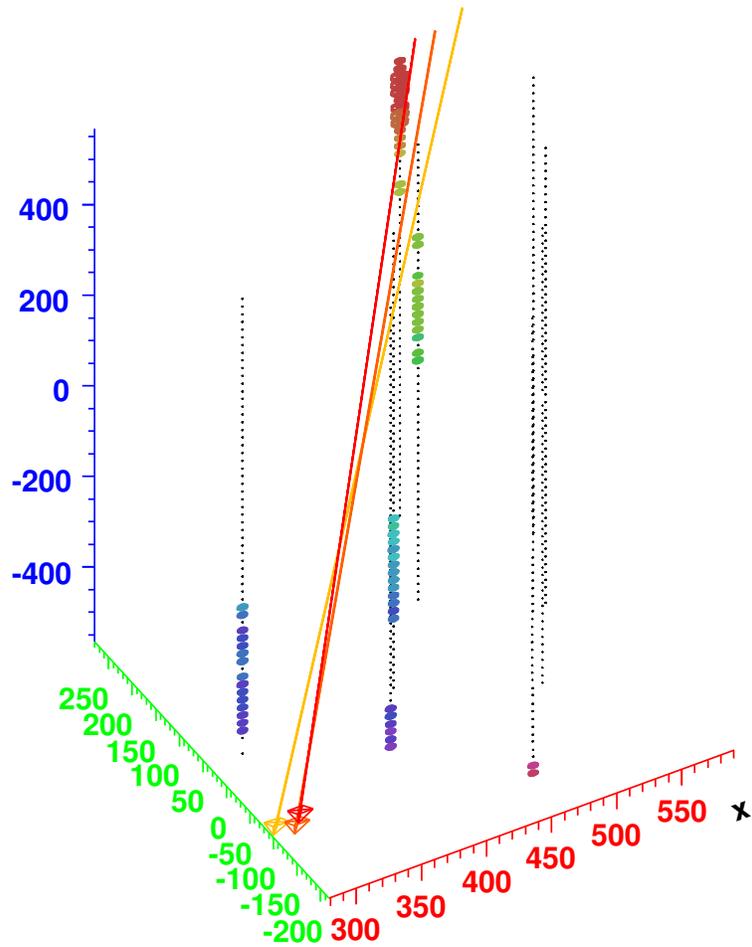


Figure 4.17: A clear down-going muon event in IceCube. 65 DOMs participated in this event and are represented by the colored dots. The colors show the relative time that the DOMs were hit, with red being early hits and blue being late hits. The small black dots represent DOMs which were not hit. The axes indicate the position of DOMs relative to the center of the future array. Also shown are the results of the linefit, dipolefit and likelihood reconstructions, described in Chapter 5.

4.3 Online PNF System and Filtering

The DAQ outputs approximately 120 GB/day (50 GB/day after compression is applied) for all the triggers currently implemented. Though this data rate is quite modest on the scale of other high-energy physics experiments, it is much larger than the bandwidth allowed for IceCube from the Pole. All data acquired from the Pole must fit over an approximately 4 GB/day satellite bandwidth allocation allowed by the South Pole TDRSS Relay (SPTR). Communications satellites are only visible from the Pole for less than half of the day. The limited bandwidth for IceCube to the Pole means that on order of 90% of the data volume must be rejected at the Pole and only the physically interesting events sent to the North. In order to facilitate this filtering, a Processing and Filtering (PNF) cluster of computers running at the Pole reconstructs the events in real-time as they come out of the Event Builder. The system selects interesting events based on a number of criteria (for different analyses) as well as sending a pre-scaled unbiased sample for study.

Data from the DAQ is buffered in DAQ-dispatch. DAQ-dispatch is a simple component which serves both to provide a clean interface between the DAQ and PNF as well as allowing each system to run independently.

The PNF system (Figure 4.18) is a server-client system with one master process, the pfserver, and several client processes, the pfclients. The pfclients run reconstruction software on the events as they come through, deciding which events to keep and which events to reject. The pfserver receives events from DAQ-dispatch and assembles chunks of events for transmission to the clients. When the clients have finished processing events, they send filter summary information back to the server. The server then applies the filter decisions and outputs both a data stream including all events and a data stream including just the filtered events. The filtered data stream is written to disk and is queued for transmission to the North. The unfiltered data stream is available for processes which want access to all events, like detector

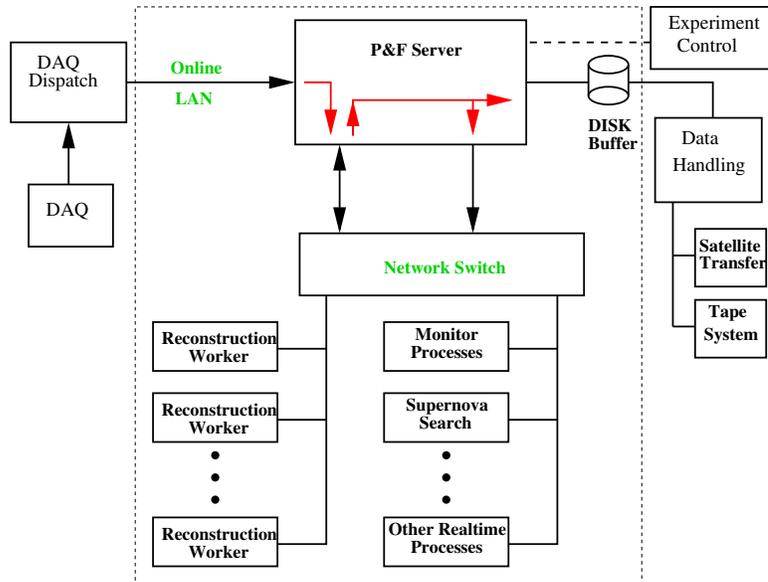


Figure 4.18: PNF System

and performance monitoring, and is eventually written to tape at the Pole. Thus all events are available on tape, though it is hoped that they are not needed.

Data that is preserved by the filter is queued up for transmission to the North. The South Pole Archival and Data Exchange (SPADE) system handles the transmission of data to the North by negotiating with the SPTR system.

The different filtering streams are useful for different analyses. For example, there are filtering streams for InIce-IceTop coincidences, and for high-energy events which have no bearing on this work. For this work, the two filtering streams of interest are the InIce SMT filtering scheme and the Up-going Muon filtering scheme.

The InIce SMT filtering scheme prescales events that pass the InIce SMT trigger for transmission to the North. The prescale rate has been set both to 1/40 and 1/400 during the the data of interest here. This scheme is useful for providing an unbiased set of down-going muons from cosmic rays while meeting bandwidth requirements.

The second filtering scheme attempts to select events which look up-going. Because there is limited processing available at the Pole, the fast first-guess reconstructions linefit and dipolefit are used to determine whether an event should be kept or

rejected. See Section 5.2 for a description of these reconstructions. This up-going muon filter requires either the linefit or the dipolefit to reconstruct the event with a zenith angle greater than θ_{cut} , and requires that the events have $nHit_{cut}$ or more DOMs hit.

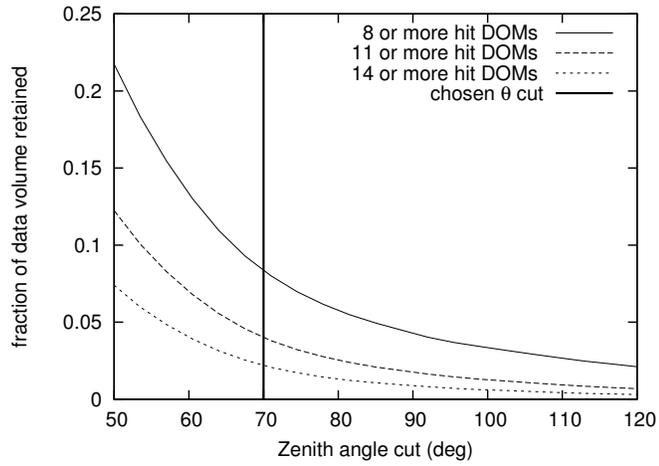
In selecting a filter configuration, we want to reject uninteresting events while keeping potentially interesting events. To do this, we set the filter settings as loose as possible while satisfying the bandwidth requirements. The results of a study on the effectiveness of the filter are presented in Figure 4.19. Figure 4.19a shows the data reduction acquired for varying settings of $nHit_{cut}$ and θ_{cut} . We must reject 95% of the data volume at the Pole in order to meet bandwidth requirements for the up-going muon stream. From Figure 4.19a we can determine settings of $nHit_{cut}$ and θ_{cut} that meets this requirement. Figure 4.19b shows the efficiency for preserving atmospheric neutrinos for a given $nHit_{cut}$ and θ_{cut} . The efficiency is calculated relative to events which pass the given $nHit_{cut}$. For the chosen filter settings of $\theta_{cut} = 70$ degrees and $nHit_{cut} = 11$, we reject 95% of the data volume at the pole, while keeping 95% of the atmospheric neutrinos which have at least 11 hit DOMs.

4.4 Calibration

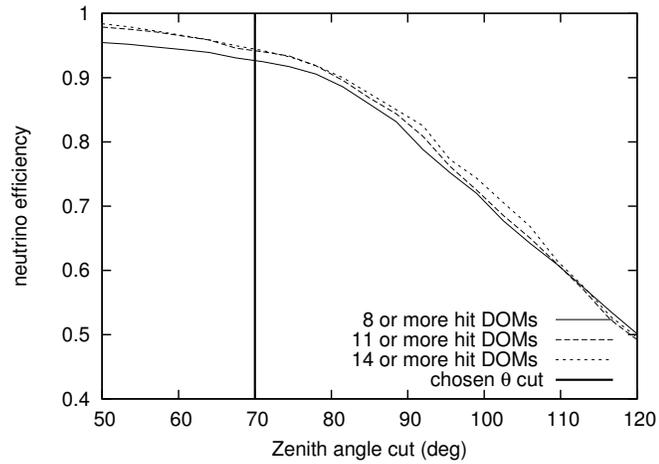
4.4.1 RAPCal Time Calibration

The IceCube detector uses a novel procedure to calibrate the time at which readouts occur. This Reciprocal Active Pulsing (RAPCal) procedure periodically synchronizes the surface time (Universal Time) to the local DOM time. Since the DOM oscillator is stable, this procedure is done rarely (once every two seconds).

The DOR card sends out a short bipolar pulse to the DOM. The DOM uses the same hardware that is typically used for digital communication to digitize the pulse when it is received. By the time the waveform is received at the DOM, it



(a)



(b)

Figure 4.19: Study of possible online filter settings. Panel (a) shows the fractional data volume kept by a particular zenith angle cut and for different $nHit_{cut}$ thresholds. Panel (b) shows the efficiency for keeping atmospheric neutrino events above the $nHit_{cut}$. On both plots the chosen θ_{cut} of 70 degrees is shown.

has dispersed to μs length and the waveform can be sampled at the relatively slow 20 MHz communications rate. The DOM then waits a known amount of time and transmits a pulse identical to the one it was sent by the DOR card. The DOM also transmits the time stamp and waveform that it recorded in the ice. The DOR card at the surface digitizes the DOM-sent waveform the same way the DOM did.

By comparing global and DOM-local timestamps for the transmission and reception events, we can get a rough scale of the timing offset between the surface and in-ice clocks, but by isolating the final waveform features we can get precision beyond the intrinsic 20 MHz timing. The precision arises from the reciprocity of the system. Since the pulse is sent down the same path both to and from the DOM, the final dispersed pulse should be identical in shape at the DOM and the DOR. It does not matter what fiducial mark on the dispersed pulse is used to time the pulse as long as we choose the same mark on both the sending and receiving pulse.

Figure 4.20 shows example waveforms acquired by the DOR and by the DOM.

Several timing calibration events are used in performing the timing calibration so as to measure both the offset between a DOM clock and the surface clock and to measure precisely the frequency of the DOM clock.

4.4.2 DOM Calibration

The DOM includes some components to allow the DOM to calibrate itself. This section contains information on the specific calibrations that are applied and sketches how the calibration is done. The chief references for this section are [40], [41], and [42].

DOMs are calibrated periodically (nominally once a month). The calibration is performed by a program *domcal* run on the DOM.

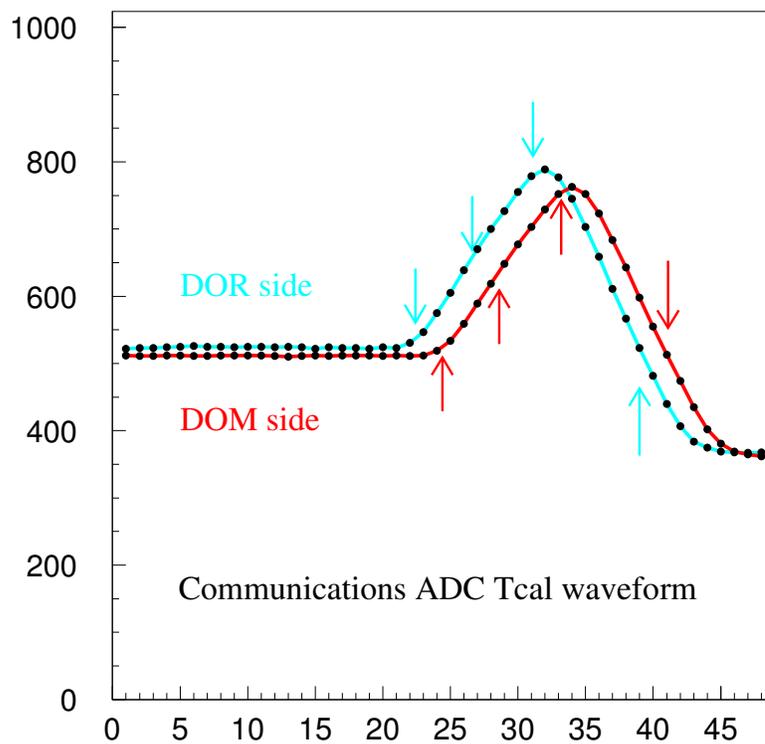


Figure 4.20: Example waveforms used in the timing calibration. Shown are the waveform recorded by the DOR and the waveform recorded by the DOM. The x-axis is units of the 20 MHz communications cycle. The arrows show waveform features as extracted by different algorithms.

Pulser Calibration

On the mainboard is an electronic pulser which produces fast PMT-like pulses. The mainboard is laid out so that the entire DOM readout can be driven by the pulser rather than the PMT. The first step of DOM calibration is to calibrate the pulser so that the height of the pulser output is known. This calibration is done by using the discriminator on the mainboard. The discriminator is set to a fixed value and the pulser is ramped up until its waveforms trigger the discriminator. In this way it is possible to set the pulser to a known voltage.

ATWD Bin Calibration

Each ATWD produces a set of 128 10-bit numbers that constitute the measurement of the waveform at the ATWD. Each of these bins has a unique linear response to the true voltage across it when the waveform was sampled. Figure 4.12 shows the waveform before calibration. It is ragged, varies bin-to-bin, and has a constant offset from 0. This offset is chosen sufficiently high so that the pedestal pattern does not cause zero values to be delivered by the ATWD digitizer. The response of each ATWD bin is measured by stepping through different values for the DOM front-end voltage and mapping out the ATWD response for each voltage. In this way we can take the discrete 10-bit response of the ATWD and map it back to the true voltage that the ATWD measured.

Amplifier Calibration

Each ATWD channel is subject to a different amplification. Nominally these amplifications are 16x, 2x, 0.25x but calibration is needed to measure these values directly. This measurement is done by feeding pulses from the pulser (now of known voltage) into the readout electronics. From the ATWD bin calibration, the voltages can be measured and the gain of each channel can be calculated.

ATWD Sampling Speed Calibration

Each ATWD is set to a nominal sampling speed of 3.3 ns/sample. Nevertheless, this value can vary from ATWD to ATWD. To measure the ATWD sampling speed, the waveform from the DOM clock is fed into the ATWD, resulting in a sinusoidal pattern in the ATWD readout. Counting the zero-crossings in the readout pattern gives the relationship between the ATWD sampling speed and the clock frequency.

PMT Calibration

After the electronics are calibrated, we must calibrate the PMT's response to single photoelectrons. To do this, real PE events are needed, so the DOM high-voltage is turned on and the DOM acquires individual photoelectron events (noise hits). Using the timing and the voltage calibrations from the previous steps, we can calculate an integrated charge in pC for each event and build up a charge histogram (Figure 4.21). This charge histogram is characterized by two components. The exponential tail of low integrated charge is the noise pedestal and results from low-amplitude noise in the PMT. The second component is the gaussian SPE peak which results when SPEs trigger the readout. The quality of the PMT at these settings can be quantified as a measure of the ratio between the height of the SPE peak and the 'valley' between the noise pedestal and the SPE peak. A higher peak-to-valley ratio means that distinguishing noise hits from true SPE hits is easier. The charge histogram is measured at a variety of high-voltage settings for the PMT. For each high-voltage setting, the peak-to-valley ratio and the position of the SPE peak can be identified and the gain measured. In this way we have a complete measurement of the PMT's response at different high voltages. Using the relationship between the high-voltage setting and the gain, the DOMs are set to the high-voltage corresponding to a gain of 10^7 . The 10^7 gain is the 2006 setting but the value can be chosen and will likely be changed reduced in the future.

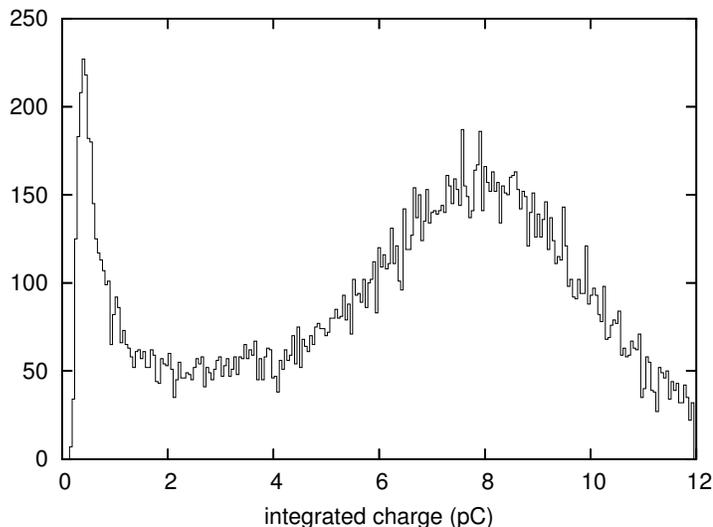


Figure 4.21: Typical charge histogram acquired during the calibration of an IceCube PMT with a voltage of 1340 V. Two features are prominent. First is the noise pedestal, the exponentially decaying low amplitude component. Also clear is the SPE peak which allows us to make the conversion from integrated charge to PEs.

Applying the Calibration

The calibration of a DOM consists of the described parameters. These parameters are acquired for each DOM periodically throughout the life of the detector (nominally once a month). During physics data-taking, the DOM is set to a high-voltage and the characteristics of the DOM at this high-voltage are interpolated from the values measured during calibration.

One point to note is that the ATWD waveforms that come from the detector do not have any calibration applied. They are the raw 10-bit numbers acquired by the DOM. The calibration is applied offline during reconstruction of the event.

It should be noted that significant compression of the waveform can be achieved in the DOM, saving bandwidth to the surface, if the ATWD pedestal is subtracted in the DOM. After the pedestal subtraction the ATWD signal is mostly 0, facilitating the compression. This is anticipated in the future.

4.4.3 Geometry Calibration

IceCube strings are deployed with two high-resolution pressure sensors spaced 1000 meters apart at the top of the string and at the bottom of the string [38]. From pressure readings at the string's final location, and incorporating the depth of the waterline in the hole after deployment, the absolute depth of the DOMs can be determined. Horizontal positions of the DOMs are measured by the combination of a survey of hole positions and data of the drill position acquired during drilling. The total accuracy in the position measurement is about 50 cm, comparable to a 1.5 ns error in the timing of individual readouts. The geometry measured during deployment can be cross-checked and potentially improved by using the flashers or down-going muons information.

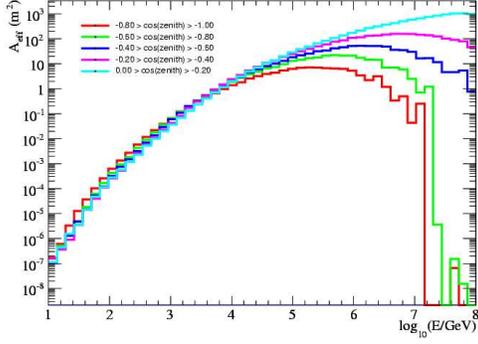
4.5 Effective Area

It is useful to characterize the response of the detector in terms of an effective area A_{eff} . For a predicted number of neutrino events N_{events} associated with a diffuse neutrino flux (differential in area, time, energy and solid angle) $\Phi(E, \theta)$, the effective area A_{eff} is defined as the function that satisfies Equation 4.2.

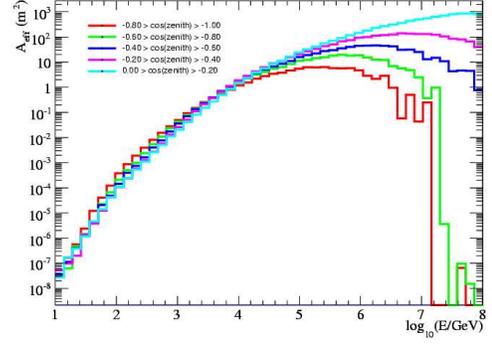
$$N_{events} = \int dt \int d\Omega \int dE \cdot \Phi(E, \theta) \cdot A_{eff}(E, \theta) \quad (4.2)$$

The effective area is defined so that it is flux-independent. Note that in order to quote only one effective area (to be precise we should quote different effective areas for ν_μ and $\bar{\nu}_\mu$) we assume that ν_μ and $\bar{\nu}_\mu$ arrive in their atmospheric neutrino ratios. With the effective area, we can calculate the expected number of events in the detector to any diffuse neutrino signal without re-running simulation. All of the effects in the simulation can be tied up in this one function.

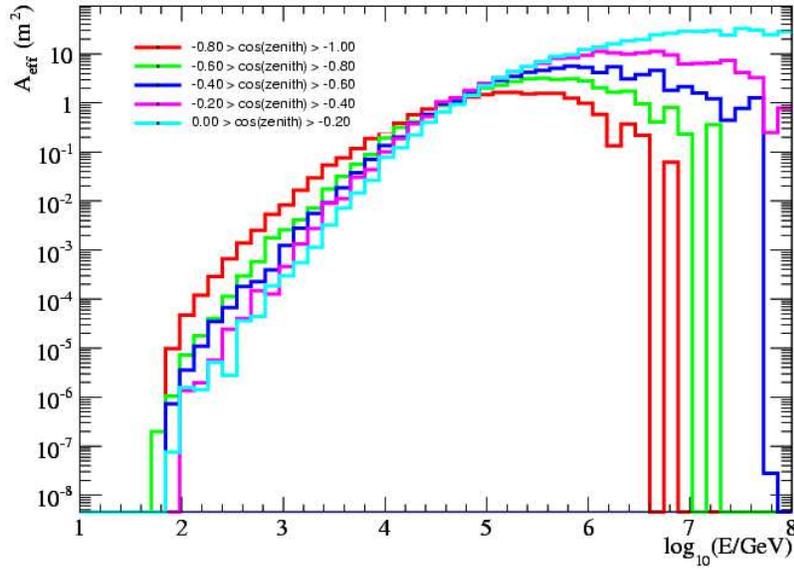
The effective area is different for different selection criteria since they result in



(a)



(b)



(c)

Figure 4.22: Effective area of the IC-9 detector to diffuse sources of ν_μ and $\bar{\nu}_\mu$ as defined in Equation 4.2. Notice the loss of effective area for high energies for vertical neutrinos. Panel (a) is the effective area at trigger level. Panel (b) is the effective are after the online filter is applied, and panel (c) is the effective area after the final event selection.

different numbers of events recorded in the detector. Figure 4.22 contains the calculated effective area for IC-9 to neutrinos. Shown are the effective area for neutrinos to trigger the detector, to pass the filter and then to make it through the final selection cuts. These selections are described in detail in Section 7. Note the loss of effective area for high energy neutrinos for nearly vertical neutrinos. Detection of high-energy neutrinos must occur at the horizon. Also note that this is the effective area to neutrinos, so the effective area is significantly smaller than the geometrical area.

Chapter 5

Reconstruction Techniques

The reconstruction of IceCube muon events takes place in several steps with different reconstruction techniques employed at each step. The reconstructions gradually improve the quality of the fit with successive algorithms taking more computing resources. With the full likelihood reconstruction, we can achieve an angular resolution of the muon's direction approaching 1 degree, depending on the strength of the cuts applied. This is on the same scale as the irreducible error from the difference between the muon and neutrino directions and is acceptable for doing neutrino astronomy.

5.1 Hit Preparation

This section examines the process of preparing hits for reconstruction. By this I mean the process of taking the calibrated DOM waveform and extracting a list of PE arrival times and positions that will be used by the reconstruction algorithms. In this process, we remove all hits from malfunctioning DOMs and all hits that are obviously noise or obviously not part of the single muon event being reconstructed.

5.1.1 DOM Readout Cleaning

DOM readout cleaning is the first step in preparing hits. In this step, we clean out all readouts that occur on a pre-defined list of 13 bad DOMs described below. Some of the bad DOMs have bad calibrations so it is helpful to avoid applying later algorithms to these DOMs. Also, recall that it is possible for both ATWDs to participate in an event, and at this point we completely remove any second ATWD waveforms because they are obviously outside the event, since they occur a minimum of $6.4 \mu s$ after the first ATWD readout.

DOMs were excluded for a variety of reasons: they were known to be dead, gave only beacon hits, had high current or a broken LC connection. Two DOMs were excluded because the calibration record was not complete, but were otherwise good. Because bad DOMs have not been included during the simulation, we manually remove data from DOMs which are known not to give good data in order to use the same processing for data and simulation.

The bad DOM list is somewhat aggressive and several DOMs were excluded that might be included when the problems with them are better understood.

5.1.2 Feature Extraction

The semi-continuous calibrated waveform is de-convolved into the sequence of photo-electron arrival events by the process of feature extraction [38], [43]. The feature extraction process is an iterative fit to the calibrated waveform.

The assumed form of the waveform is given by the sum of n single SPE-like waveforms:

$$F^n(t) = b_0 + \sum_{k=1}^n A^k f\left(\frac{t - t^k}{\delta^k}\right) \quad (5.1)$$

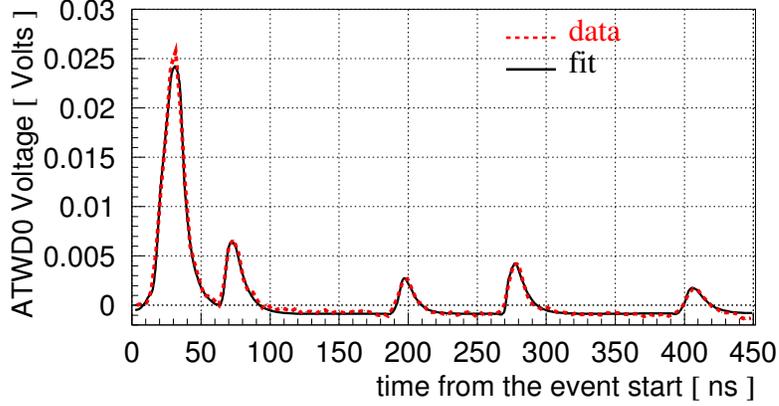


Figure 5.1: An example calibrated ATWDO waveform and the feature-extracted best fit.

Here b_0 is the baseline estimate. A^k , t^k and δ^k are respectively the amplitude, time and width of the k th pulse. The function $f(\zeta)$ is the SPE waveform.

The fit begins with $n = 1$ and proceeds with successively more SPE waveforms in the hypothesis function. At each iteration, the parameters are modified to minimize χ^2 for the fit. The difference between the fitted waveform and the predicted waveform is then used to seed the time for the $n + 1$ st pulse in the next iteration. The fit proceeds with successively more SPEs in the hypothesis until the χ^2 is small, stops improving, or we reach a pre-defined maximum number of pulses. The final result is then a linear combination of a number N of SPE-like waveforms as seen in Figure 5.1.

5.1.3 Time Window Hit Cleaning

The triggering of IceCube events allows DOM readouts to be included in events for a $16 \mu s$ window. Furthermore, since the global trigger combines overlapping triggers, it is not uncommon to have events where the times between the first and last hits in the event can be up to $40 \mu s$. The time that the muon spends within the detector volume is considerably smaller at a maximum of $3 \mu s$, and noiseless simulation shows that

most of the light is detected within a window of $4 \mu s$. Any hits that arrive outside of a window of $4 \mu s$ cannot be light from a single muon.

In order to reduce contamination from other sources (random noise or other muons), the first step in the hit cleaning procedure considers a sliding time window of $4 \mu s$. Over all of the hits in the event, the $4 \mu s$ time window in which the maximum number of hits occur is chosen. All hits which fall outside of this time window are cleaned out and not used in reconstruction.

This cleaning procedure has the additional benefit that it facilitates comparison between data and simulation since the detector simulation currently only models single triggers.

5.1.4 Isolated Hit Cleaning

Isolated hit cleaning is another algorithm used to reduce contamination from hits that are not believed to be physically related to the muon passing through the detector. In this step, a hit is cleaned out if there is no other hit within 100 meters and within 500 ns. These values are standard for AMANDA analyses and have not been tuned for IceCube and are largely redundant after the Local Coincident condition.

5.1.5 First Hit Cleaning

The final step in the hit cleaning procedure is to consider for each DOM only the earliest hit on the DOM. This step was typically used in AMANDA muon analyses, and it is not strictly necessary for IceCube analyses since in IceCube the entire waveform is available. This step may be removed in later analyses and has been chosen in order to reduce the complexity of this first analysis. We know the approach of taking only first hit times is satisfactory for AMANDA. In general, we do not expect the extra waveform information to be really useful until we get to higher energies.

5.1.6 Re-triggering

The simulation does not currently involve simulation of bad DOMs. For this reason, the triggering condition is subtly different for data than for simulation. As a final check to make sure that simulation faithfully represents data, we reconsider the triggering threshold of 8 DOMs and remove any events which – as a result of the hit cleaning – now fall below this threshold.

5.2 First-guess Reconstruction Methods

The first-guess reconstructions serve two purposes in the reconstruction of IceCube events. First, they are used as seeds to the full likelihood reconstruction (Section 5.3) and second they are used at the Pole as a fast way to determine the direction of an event for filtering. Two first-guess reconstructions are presented here: the linefit and the dipolefit.

Additionally, we can use the sobol seed algorithm to generate a list of seeds of directions distributed uniformly on the unit sphere. This is a brute force way to search for the best solution and is used on a very restricted dataset.

5.2.1 LineFit

The linefit algorithm is a simple χ^2 fit based on the false assumption that light from the muon is a simple plane moving at a fitted velocity. Assuming the position of the hit i of a total of N_{hit} hits is given by \mathbf{r}_i and the time given by t_i , the equation for χ^2 is

$$\chi^2 = \sum_{i=1}^{N_{hit}} (\mathbf{r}_i - \mathbf{r} - \mathbf{v} \cdot t_i)^2 \quad (5.2)$$

where \mathbf{r} is the reconstructed position and \mathbf{v} is the reconstructed velocity vector. This equation can be minimized analytically and yields a solution

$$\mathbf{r} = \langle \mathbf{r}_i \rangle - \mathbf{v} \cdot \langle t_i \rangle \quad (5.3)$$

$$\mathbf{v} = \frac{\langle \mathbf{r}_i \cdot t_i \rangle - \langle \mathbf{r}_i \rangle \cdot \langle t_i \rangle}{\langle t_i^2 \rangle - \langle t_i \rangle^2} \quad (5.4)$$

where the $\langle X_i \rangle$ denotes an average over all hits of the parameter X_i . Though it isn't used in this analysis, the reconstructed velocity \mathbf{v} can be used to distinguish good fits when $|\mathbf{v}|$ is close to c .

5.2.2 DipoleFit

The dipolefit is another simple, fast algorithm. In this algorithm, we average the unit vectors from a hit DOM to the subsequently hit DOM. The reconstructed direction is then the direction of the dipole moment \mathbf{M}

$$\mathbf{M} \equiv \frac{1}{N_{ch} - 1} \sum_{i=2}^{N_{ch}} \frac{\mathbf{r}_i - \mathbf{r}_{i-1}}{|\mathbf{r}_i - \mathbf{r}_{i-1}|} \quad (5.5)$$

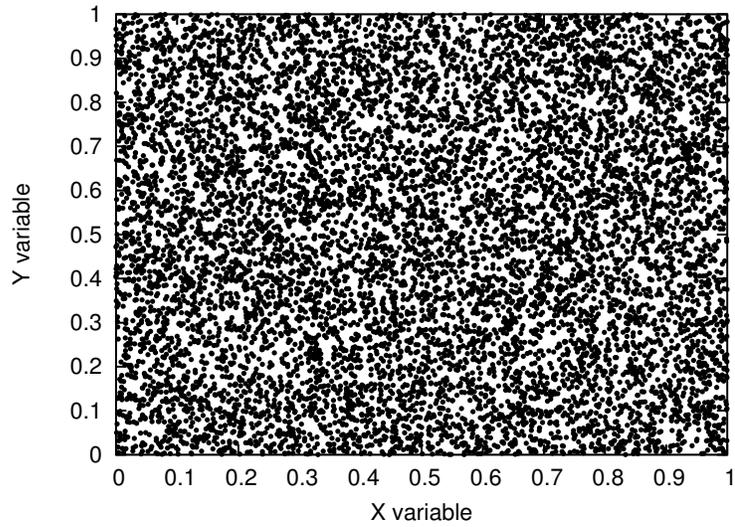
where \mathbf{r}_i are the positions of each of the N_{ch} hit channels, understood to be taken in the time order they occur.

5.2.3 The Sobol Seed

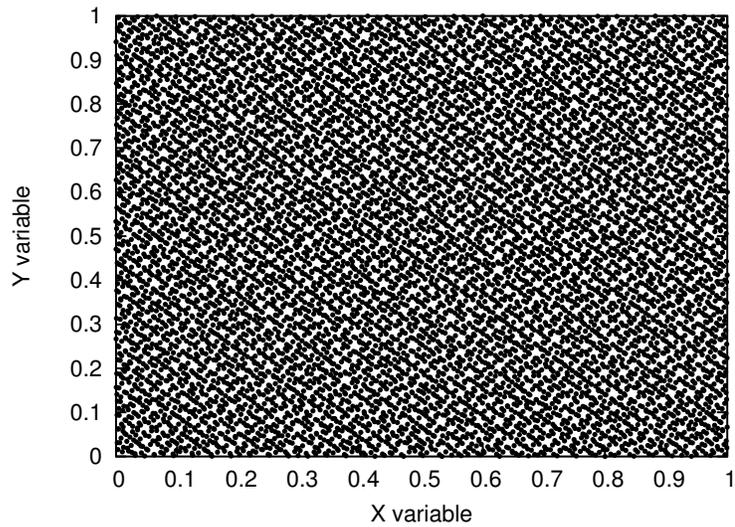
The sobol sequence [44] is an algorithm for generating a pseudo-random sequence of numbers. The sobol sequence differs from a random sequence of numbers in that it uniformly samples a space rather than randomly sampling it (see Figure 5.2).

This uniform sampling occurs no matter how many samples are chosen because the algorithm fills in empty space uniformly as more and more samples are chosen.

The 2-D sobol sequence can be used to generate a list of directions sampled uniformly across the unit sphere by choosing the first number x to be on the interval



(a)



(b)

Figure 5.2: Panel (a) is 10^5 points drawn from a uniform random distribution. Since the distribution is random, there are areas of the space that are oversampled and areas which are under-sampled. Panel (b) is 10^5 points drawn from the 2D Sobol sequence. Here there is no clustering and the space is uniformly sampled.

$x \in (-1, 1)$ and the second number y on the interval $y \in (0, 2\pi)$. The spherical variables θ and ϕ are then distributed uniformly on the unit sphere with the assumption

$$\theta = \arccos(x)$$

$$\phi = y$$

5.3 Muon Likelihood Fitter

5.3.1 The Approach

The method of maximum likelihood is a well-known technique for estimating a set of unknown parameters \mathbf{a} from a set of independent measured values $\mathbf{x} = \{x_i\}$. The technique proceeds by forming a likelihood function L

$$L(\mathbf{x}|\mathbf{a}) = \prod_i p(x_i|\mathbf{a}) \quad (5.6)$$

which is the likelihood of obtaining the measured values \mathbf{x} under the assumption that the parameters are \mathbf{a} . Once we have this function, the parameters \mathbf{a} that correspond to the maximum value of L are taken to be the measured parameters.

This technique can be used – with some modifications – in the context of reconstructing muon events in IceCube. Here the unknown parameters \mathbf{a} are parameters that determine an infinite muon track in the detector. Though it is possible to parametrize the muon track in a way that is unambiguous, it is typically assumed that the parameters that specify a muon track are some position x, y, z and time t along the track as well as the zenith angle θ and azimuth angle ϕ of the direction the muon is coming from. There is one too many degrees of freedom in this set because the vertex position specified is arbitrary. To keep the likelihood function a function of the five independent parameters, the time t of the vertex is taken to be a constant value.

In the present definition we do not consider DOMs which are not hit as part of

the measured signal, and the only DOMs that contribute to L are the DOMs which actually have a hit. The measured parameters x_i are then the position and times of the measured hits in the detector.

5.3.2 The Likelihood Function

All the physics is contained in our choice of the single-hit likelihood function $p(x_i|\mathbf{a})$. It is useful, in considering the requirements for this function, to define the number t_{geo} which is the time that the PE is expected to arrive at a particular position, under the simple hypothesis of unscattered Cherenkov emission. Given the variables as defined in Figure 5.3, it can be shown that

$$t_{geo} = t_0 + \frac{\hat{\mathbf{p}} \cdot (\mathbf{r}_i - \mathbf{r}_0) + d \cdot \tan(\theta_c)}{c_{vac}} \quad (5.7)$$

where c_{vac} is the speed of light in a vacuum. This is the simplest form of the equation under the assumption that the phase and group velocity of light in ice are the same. In truth, the calculation of t_{geo} uses the phase velocity to determine the angle of the Cherenkov cone, and the group velocity to determine the amount of time taken by the light to travel from the emission point to the detection point [45].

With this definition of t_{geo} we can define for a hit at time t_{hit} and for a given hypothesis track the residual time t_{res} as

$$t_{res} \equiv t_{hit} - t_{geo} \quad (5.8)$$

which is the amount of time by which the hit time differs from a pure Cherenkov hypothesis. A positive t_{res} would be a hit that arrives later than expected, and a negative residual would be a hit that arrives earlier than expected, for a given hypothesis.

For each hit in the detector (that is for each x_i), if we have a hypothesis \mathbf{a} we can assign a residual for that hit $t_{res,i}$ and the function $p(x_i|\mathbf{a})$ becomes a probability

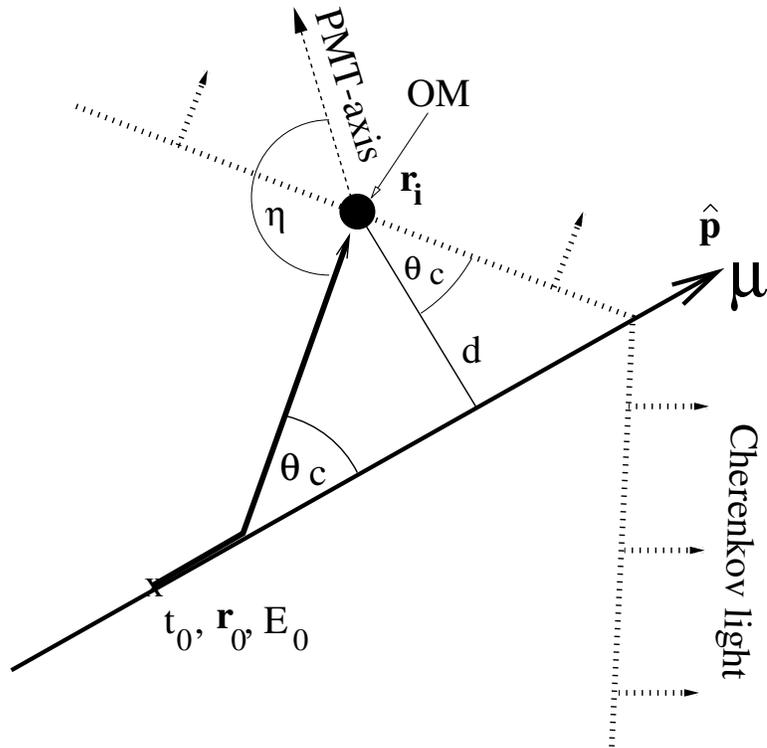


Figure 5.3: The coordinate system used for defining the single-hit likelihood function. θ_c is the Cherenkov angle, and $\hat{\mathbf{p}}$ is the momentum unit vector. The angle η is the angle of light arriving at the DOM. The vector \mathbf{r}_i is the position of the DOM. t_0 and \mathbf{r}_0 are an arbitrary time, and the corresponding position of the muon, and d is the distance of closest approach between the muon and the DOM.

density function (PDF) in t_{res} for observing the hit at the given position with a residual t_{res} . Furthermore, under the assumption that the distribution of residuals depends only on the distance d_i of the muon from the observation point and the angle η_i of the Cherenkov cone on the DOM, the single hit probability becomes a simple function of three variables, $t_{res,i}$, d_i , and η_i

$$p(x_i|\mathbf{a}) \rightarrow p(t_{res,i}, d_i, \eta_i) \quad (5.9)$$

Pandel Function

The Pandel function [46] [47] is an analytic estimate for the arrival time distribution of light emitted from a monochromatic, isotropic point light source in a medium where scattering is a dominant effect. The Pandel function is expressed relative to the residual time t_{res} and the distance d from the emission point to the detection point:

$$p_{pandel}(t_{res}, d) \equiv \frac{1}{N(d)} \frac{\tau^{(-d/\lambda)} \cdot t_{res}^{(d/\lambda-1)}}{\Gamma(d/\lambda)} \cdot e^{-\left(t_{res} \cdot \left(\frac{1}{\tau} + \frac{c_{medium}}{\lambda_a}\right) + \frac{d}{\lambda_a}\right)} \quad (5.10)$$

and the normalization is given by:

$$N(d) = e^{-d/\lambda_a} \left(1 + \frac{\tau \cdot c_{medium}}{\lambda_a}\right)^{-d/\lambda} \quad (5.11)$$

The free parameters are λ , an effective scattering length, λ_a , an effective absorption length and a time scale τ .

Forming the PDF

The pandel function has some limitations. It is undefined for $t_{res} < 0$ and has a pole at $t_{res} = 0$. Furthermore, in its current form it does not allow for realistic PMTs because of PMT jitter or the fact that PMTs have an orientation relative to the source of the light.

To account for orientation of PMTs, the procedure of patching according to [47] has been followed. The distance d is taken to be an effective distance $d_{eff}(\eta)$ which accounts for the fact that the light is more scattered when the PMT is oriented away from the light source. To account for the problems with the pandel function near $t_{res} = 0$, we use the pandel function only for large enough t_{res} . For negative t_{res} we use a Gaussian with a width corresponding to the uncertainty in the timing (from PMT jitter or other timing uncertainty), and for small positive t_{res} we use a polynomial which matches the value and first derivative at the points where it overlaps with the pandel function and the Gaussian.

In addition to the patching, a small constant offset is added to the PDF to account for the fact that a noise hit can occur at any time.

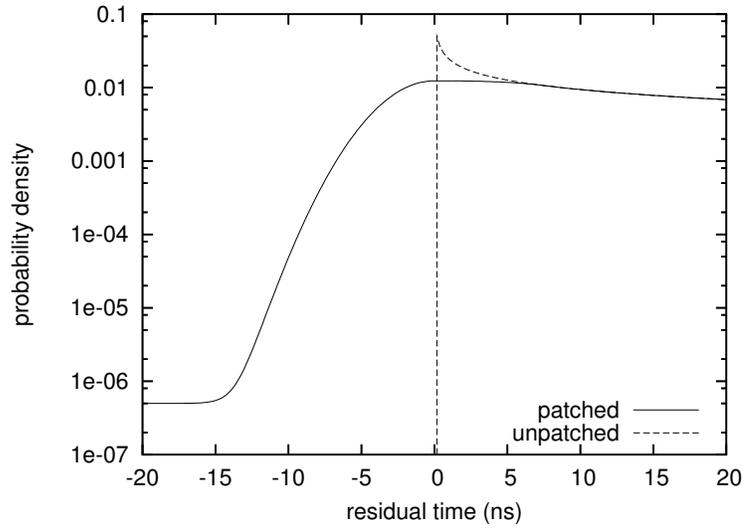
The values for the free parameters of the pandel function are taken from [46].

5.3.3 Implementation

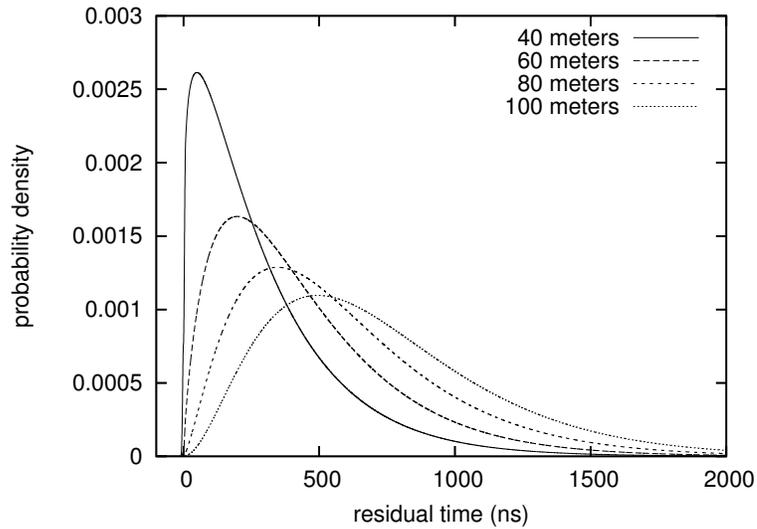
In broad terms, the task of the likelihood fitter is to identify the parameters \mathbf{a} that maximize the value of $L(\mathbf{x}|\mathbf{a})$. In practice, what is actually done is finding the minimum of $-\ln(L)$, which is mathematically equivalent, but it is easier to work with logarithms since we are dealing with numbers of order 10^{-7} .

Practically speaking, what is done is to first construct a list of hypotheses for \mathbf{a} using the first guess methods from Section 5.2 or by sampling the parameter space for \mathbf{a} systematically using the sobol seed from Section 5.2.3. For each of the selected hypotheses, a numerical minimizer algorithm is used to identify the minimum. In the present analysis, the minimization algorithm used is the *simplex* algorithm as implemented by the TMinuit package in ROOT [48].

The likelihood space for any given event can contain several local minima. For instance, Figure 5.5 shows one projection of the likelihood space for a particular event. To obtain this plot, the likelihood reconstruction was run for a variety of



(a)



(b)

Figure 5.4: Panel (a) shows the unpatched and patched panel PDF at 20 meters. The effect of the patching is clear as it gives the PDF a finite value for negative residual times and accounts for timing uncertainty. The noise level is visible. Panel (b) shows the PDF at several distances from the track. Further hits tend to result in more delayed light.

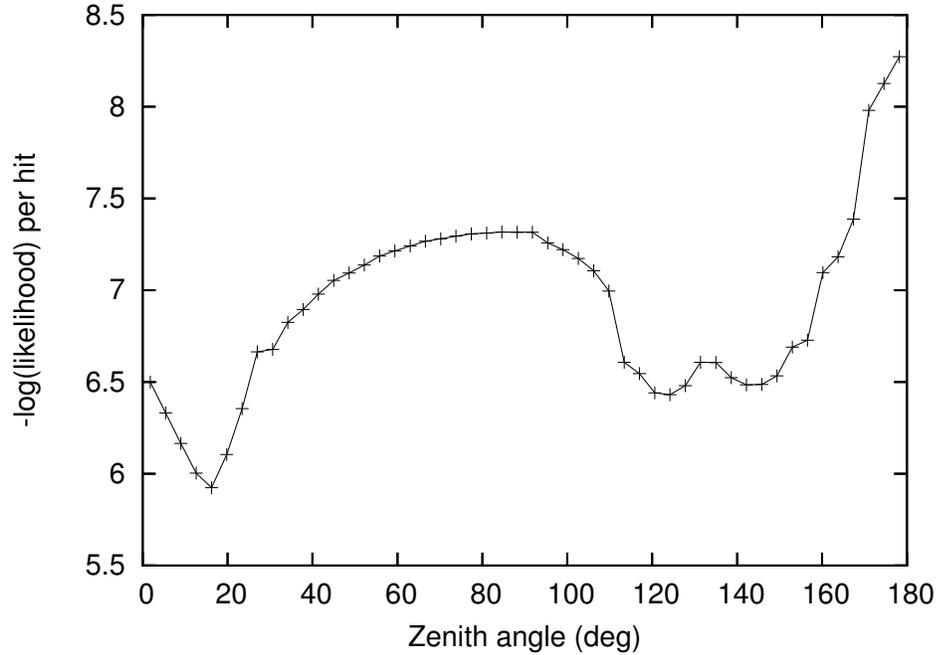


Figure 5.5: A cross-section of the likelihood space for one particular event. For each zenith angle, the optimal fit was found by varying the vertex and azimuth of the track while holding the given zenith angle fixed. The likelihood space shows three strong local minima and it is possible for the numerical minimizer to get 'stuck' in the wrong minimum if the seed is poor.

zenith angles with the zenith angle fixed during the minimization of the likelihood. The plot shows the best likelihood value corresponding to each zenith angle.

In order to ensure that the reconstructed 'maximum likelihood' fit is indeed the global maximum likelihood, the fitter can be run with the sobol seed which will sample the track parameter space and identify any minima that would otherwise be missed.

One detail that bears mentioning is that since the first guess methods and the sobol seed algorithms do not understand the Cherenkov cone and do not reconstruct a realistic vertex (only the direction), the vertex for each seed is therefore chosen such that the residuals are all positive.

5.4 Quality Parameters

There are a number of quality parameters that can be used to determine whether or not a fit is 'good' or not. Events that have good quality parameters are more likely to be reconstructed correctly.

These quality parameters are based on the idea of *direct hits*. A DOM hit is considered to be direct if it has a residual time t_{res} such that:

$$-15\text{ns} > t_{res} > 75\text{ns} \tag{5.12}$$

This condition suggests that the hit arrives without much scattering and therefore the hit is trusted to convey more reliable information about the track geometry. The first quality parameter is the number of direct hits N_{dir} . An event with a large number of direct hits has a higher quality than an event with only a few direct hits.

The second quality parameter is the direct length L_{dir} . This parameter is calculated by projecting the position of each of the direct hits onto the reconstructed muon (Figure 5.6). The distance between the furthest direct hit projections is the direct length. A long direct length represents a high-quality event because there is a long lever arm of direct hits so that we trust the reconstruction result.

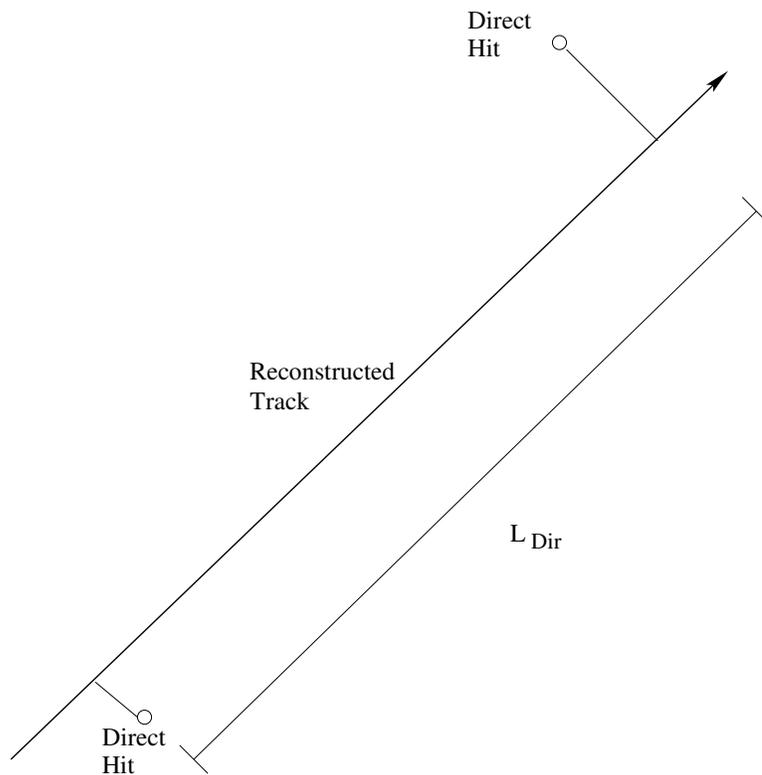


Figure 5.6: The definition of L_{dir} . Direct hits are projected onto the reconstructed muon track. The distance between the furthest separated projections is the direct length.

Chapter 6

Simulation

The chief difficulty in tagging up-going muons in IceCube is the 10^6 larger background of muons traveling down through the detector. These down-going muons are produced in cosmic-ray air showers in the atmosphere and constitute the dominant source of events in the detector. In order to study the effect of any cuts on the data, simulation is needed both of the down-going muon sample from cosmic rays (the background) and of the up-going muon sample from atmospheric neutrinos (the signal).

Simulation of events in IceCube takes place in several steps and with several different independent components responsible for pieces of the simulation:

1. *Generation.* In this step, the interaction of primary particles (either cosmic rays or atmospheric neutrinos) are simulated up until the point that they produce muons that might be seen in IceCube. The result of this stage of simulation is a set of muons that could potentially trigger the detector.
2. *Muon Propagation.* In this stage, the muon is simulated from the point of its creation until it loses all of its energy. The energy loss of the muon is simulated and points of stochastic energy loss are identified.
3. *Photon Propagation.* Photon propagation is performed using tables pre-tabulated

to contain the photoelectron density and photoelectron arrival time PDFs. The result of this step is a list of photo-electron arrival times for each DOM.

4. *PMT Simulation.* Knowing how many photo-electrons arrive at each DOM, we can simulate the PMT, the result of which is a continuous voltage waveform that is incident on the DOM mainboard.
5. *DOM Simulation.* In this step, the capture of the PMT voltage pulse by the ATWD and fADC is simulated, accounting for the DOM hardware. The LC condition between DOMs is then applied.
6. *Trigger Simulation.* Since the DAQ is lossless, the only component of the DAQ that must be simulated is the logic of the trigger.

6.1 Generators

6.1.1 CORSIKA Air Shower Generator

The CORSIKA [49] program for simulating extensive air-showers in the atmosphere was used to model the down-going cosmic-ray muon flux. The QGSJET01 hadronic interaction model from [50] implemented in CORSIKA is used to model the high-energy hadron interactions, and the GHEISHA model from [51] in CORSIKA was used to model the low-energy hadronic interactions. The cosmic-ray spectrum and composition is based on the model of [4].

The standard CORSIKA program has been modified to correct for the curvature of the atmosphere which becomes important for penetrating muons at high zenith angles [52].

There is an approximately 20% uncertainty in the overall uncertainty in the cosmic-ray flux. In order to get agreement on the event trigger rate between the data

and the simulation, we end up having to scale down the trigger rate in simulation by 15%, consistent with this overall uncertainty.

6.1.2 Neutrino Generation with NeutrinoGenerator

Simulating neutrinos is tricky because of their low cross-section. The first naive simulation would throw neutrinos and see how many interacted near the detector. The problem with this approach is that, because of the minuscule interaction cross section, we would have to throw a huge number of neutrinos to see a single interaction. This is not practical.

Instead, the neutrino generator begins with the assumption that a particular neutrino interaction occurred and then works back to see the probability that such an interaction would occur, considering the amount of matter crossed in order to have this interaction. The probability that the interaction would occur assigns a weight to the neutrino. Furthermore, the spectrum simulated is typically a very hard spectrum E^{-1} , and the events can be re-weighted to whatever spectrum is desired.

By re-weighting the events, we can consider any particular neutrino flux model without re-running the simulation. For this study, the neutrino flux model of [30] was chosen to estimate the neutrino rates based on a full three-dimensional simulation of the neutrino flux. This is the 'bartol' line in Figure 3.2. No prompt neutrino flux from charmed meson production was assumed.

The parton distributions assumed in the calculation of the neutrino-nucleon cross section come from the CTEQ5 calculation [35].

6.2 Muon Propagation with MMC

In order to simulate a muon's passage through matter we first make a rather arbitrary distinction between continuous and stochastic energy loss by the muon. The

distinction is arbitrary because each loss of energy by a muon is inherently a stochastic event. To the extent that we can make this distinction, however, it is the job of the muon propagation code to simulate when the stochastic energy loss events occur and what kind of interaction occurs.

Muon Monte-Carlo (MMC) [36] is a program written for AMANDA for simulating TeV-PeV muon passage through ice. MMC divides the muon's energy loss into a stochastic and a continuous component. For each stochastic loss event, MMC reports what kind of event it was (a pair-production event, a bremsstrahlung event or a photo-nuclear interaction) and the energy lost by the muon at that point. The range of the muon is also calculated.

The formulas used to calculate the various energy loss mechanisms are accurate to within about 1% for muons less than 10 TeV. Above 10 TeV, the theoretical are somewhat higher.

6.3 Photon Propagation with Photonics

From MMC, we have a record of the muon, where it lost energy continuously and at what points along its length it lost energy due to either an electromagnetic or hadronic shower event. Any hadronic or electronic secondaries quickly result in a hadronic or electromagnetic shower. It is important to note that the scale of the development of this shower is much smaller than the typical separation between DOMs, so the approximation is made that these showers occur at a single point.

It is the job of the photon propagation simulation to take this information and simulate where visible light photons will be detected. Because it is impractical to track every photon from these interactions over a km^3 , the photon simulation (photonics) is performed once and a table is constructed with photon densities and arrival time distributions are pre-tabulated. This simulation is done separately for the dis-

crete point-source shower-like events along the muon track and for the muon itself. The resulting photonics tables are parametrized by the direction of the muon (or shower), the relative position of the muon (or shower) and the DOM, the angle of detection on the DOM, and the absolute depth of the muon (or shower) and DOM in the ice. The absolute depth of the muon (or shower) is needed because the ice has different scattering and absorption properties at different depths.

In a second pre-tabulating step, the assumption is made that the angular efficiency of each DOM is the same, and the angle of detection on the DOM is folded in with the measured DOM angular efficiency. The final tables actually contain PE densities, not photon densities.

When the simulation is finally run for a particular event, the photon propagation step is simple. For each DOM, we look up the average number of PEs, as specified by the photonics tables. The mean photons detected are summed over the muon and shower tables. We then draw a random number from a Poisson distribution with that mean. This gives us, for each DOM, the number of PEs that that DOM will see. For each PE seen at the DOM, we look again at the photonics tables and draw a random time from the arrival time PDF. In this way we have all the PE arrival times for each DOM.

6.4 Hardware Simulation

From photonics we get individual PEs and their arrival times. In the next step, the Root-based Optical Module EmulatOr (ROMEEO), from a parametrized model of the PMT, simulates the charge measured by the PMT and superimposes a sequence of Gaussian PMT voltage pulses for each PE arrival. The resulting voltage waveform (for example Figure 6.1) is the input to the DOM mainboard simulation.

The mainboard simulation consists primarily of determining the measured ATWD

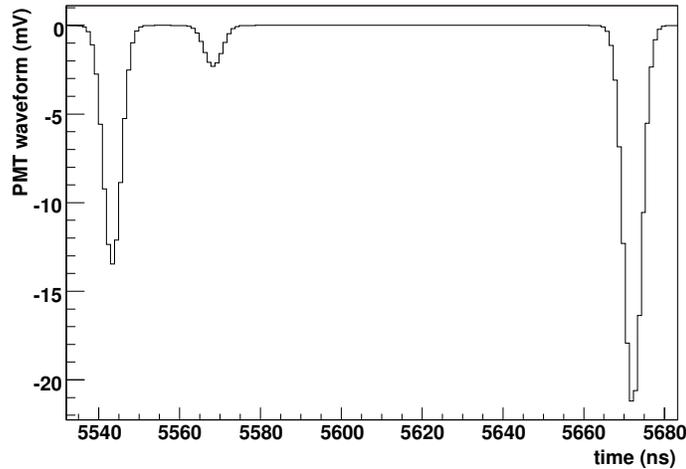


Figure 6.1: Example true PMT waveform simulated by ROMEO.

and fADC waveforms based on the true waveform ROMEO predicts. The delay board and clock phase (cf Section 4.2.1) are simulated, varying the ATWD readout window uniformly by 25 ns with respect to the arrival of the recorded waveform, before the ATWD and fADC simulation. The ATWD and fADC simulation is then essentially an 'uncalibration' taking the true waveform and assigning a discrete binary value for each ATWD and fADC bin using the ATWD calibration constants, with a small amount of bin-to-bin noise added, at the level of 1 count. At this point, Local Coincidence is simulated and only those DOM readouts which satisfy the Local Coincidence condition participate in the event.

Finally, since the DAQ is lossless, the only component of the DAQ that must be simulated is the trigger logic. Events are considered to have triggered the detector if they have more than 8 readouts in a $5 \mu s$ time window. All readouts in a $16 \mu s$ window are assembled into the final event.

6.5 Coincident Air Shower Simulation

Events occur in IC-9 at a rate of about 143 Hz in IceCube, with each event consisting of all fired DOMs in least a $16 \mu s$ window. The bulk of this is due to muons from single air showers at the surface. However, uncorrelated air showers can occur at the surface within this event window. Uncorrelated muons from coincident air showers form a background to the observation of atmospheric neutrinos and can fake up-going muons if the first event hits near the bottom of the detector and the second event hits near the top. Furthermore, in IC-9 all the strings are 'outer' strings meaning that no strings are completely contained by other strings. A typical event hits only two or three strings. This means that it is particularly easy for just the right uncorrelated down-going muons to fake an up-going muon. This situation will improve when IceCube is larger and a veto region around the detector can be formed. For all these reasons, it is important to carefully handle coincident air-showers in the detector.

The coincident air-shower simulation occurs with the same simulation chain that produces the regular single air-shower simulation. When the simulation is run, the trigger threshold simulated is set to 1 DOM rather than the normal 8 DOM trigger simulation. This forms a set of events which produce some signal in the detector. These events occur at a nominal rate of 548 Hz (469 Hz after the same rate correction applied to the single air-shower simulation). To form the coincident air-shower simulation, pairs of these events are overlaid into single events, varying the time of the second event by ± 6000 ns with respect to the first event. These events are then run through the same processing as the single events. This processing includes removing all events with fewer than 8 DOMs hit in $4 \mu s$ (See Section 5.1.3), which is more restrictive than the trigger condition. In this way we simulate events which would make it through the first level of hit cleaning and processing due to uncorrelated coincident muons.

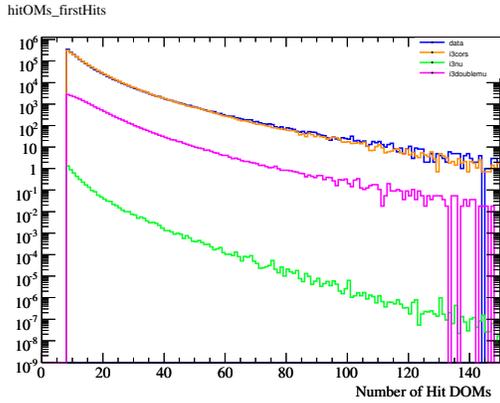
The calculation of the coincident muon rate is trivial. The product $469 \text{ Hz} * 12$

μs is the probability that a 'single hit' event will occur within a given $12 \mu s$ window. The rate of coincident 'single hit' events is then $469 \text{ Hz} * 12 \mu s * 469 \text{ Hz} = 2.64 \text{ Hz}$. This is the rate of 'coincident air-shower events' with at least 2 DOMs hit. The rate passing the trigger is lower when the 8-DOM threshold is applied.

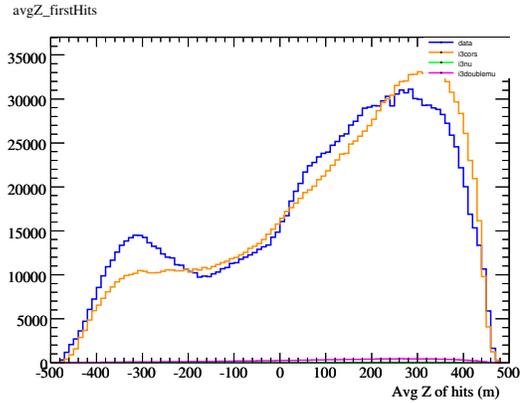
6.6 Comparison to Data

The simulation as described agrees rather well with data after the mentioned 15% rate adjustment. A full treatment of the agreement between data and simulation must look at all cut levels in analysis. This will follow, but a natural place to start is to compare the rates and shapes of a few important distributions at trigger level, as events roll off the detector.

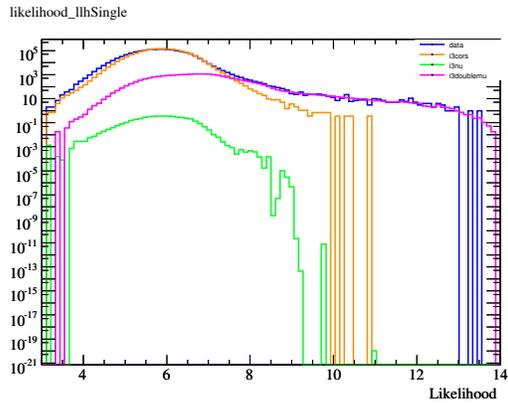
In Figure 6.2 we see these first comparisons. For these plots, the data was prepared as detailed in Chapter 5. Only the time of the earliest feature-extracted hit is used, and the reconstruction shown is the result of the likelihood fitter using both the dipolefit result and the linefit result as seeds. One important area of disagreement between data and simulation is the disagreement between the rate of events reconstructed as up-going. This disagreement will be discussed in Chapter 8.



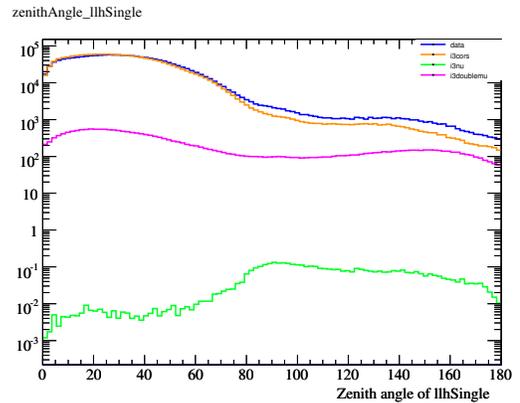
(a)



(b)



(c)



(d)

Figure 6.2: A comparison of the agreement of data to simulation. Shown are three simulated lines and data. The data is shown in blue, the single air-shower simulation is shown in orange, the coincident air-shower simulation is shown in purple, and the atmospheric neutrino simulation is shown in green. Panel (a) shows the distribution of number of hit DOMs. Panel (b) shows average position of hits in each event. Panel (c) shows the likelihood from the likelihood reconstruction. Panel (d) shows the reconstructed zenith angle.

Chapter 7

Analysis Procedure and Optimization

7.1 Overview of the Analysis

The analysis procedure is designed to generate a set of events which are likely to be caused by up-going neutrinos, with a minimum contamination from other sources. Up-going events are identified by using the muon likelihood reconstruction to determine the direction of the muon. If the muon is moving up through the detector, the event is kept. If the muon is moving down through the detector, the event is cut. Most down-going events are relatively easy to reconstruct and identify, but the number of down-going events is roughly 10^6 times larger than the number of up-going muon events. The overwhelming quantity of down-going events results in a fairly high rate (a few Hz) of mis-reconstructed events. These mis-reconstructed down-going events are of relatively poor quality and can be removed with quality cuts. Furthermore, it is computationally prohibitive to run the best reconstructions on all the events, so the analysis proceeds by a set of levels with successively better reconstructions and tighter cuts. After these reconstructions, we have a set of events that is still dominated by down-going events which are mis-reconstructed as up-going events. A final set of quality cuts is applied to remove these mis-reconstructed events and give us the final neutrino candidates.

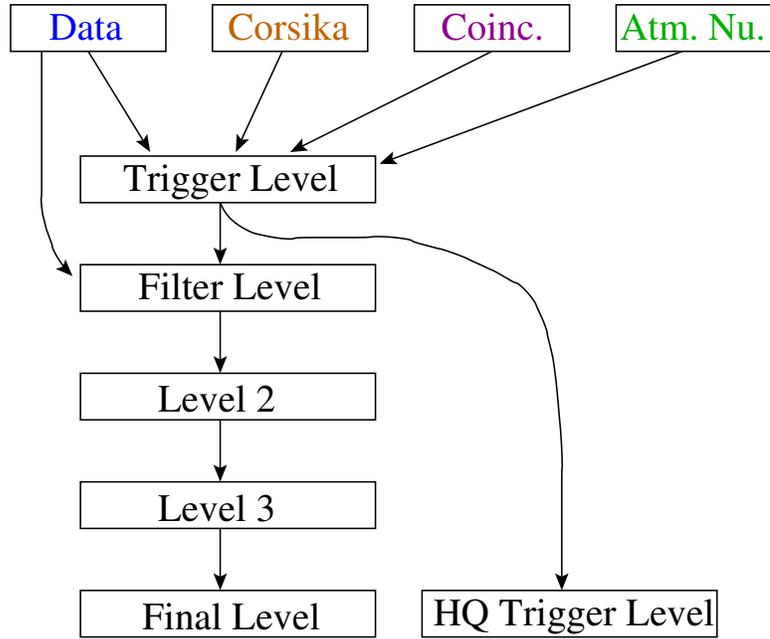


Figure 7.1: The processing levels used in the analysis. The sequential levels are cuts to remove down-going events with gradually better reconstructions. Level 3 is comprised of all up-going events. From level 3, quality cuts are applied to obtain the final dataset. These same quality cuts are applied to the down-going events to generate a high-quality down-going sample for study as well.

The same processing is applied identically to the experimental data and the simulated events so that we can gauge the response of the different components of the simulation. Using the simulation, we can place the cuts appropriately to optimally expose neutrino events and remove background.

7.2 Processing Levels

The following are the levels used in the extraction of neutrinos and for studying the down-going muon background. Figure 7.1 shows schematically the levels and the datasets.

7.2.1 Trigger Level

'Trigger level' is a small unbiased set of data that is saved and transmitted to the North regardless of the filter result. This data is identical to data as it rolls off the detector, except that it survives at a reduced rate. At this level, the hit preparation described in Section 5.1 is applied, and the events are re-triggered insisting that at least 8 hit DOMs survive the hit cleaning. At this stage the linefit and dipole fit are applied to the data. At the Pole, the likelihood reconstruction is not run, but for the small sample preserved through the filter the likelihood reconstruction is run using the results of the linefit and dipolefit as seeds.

7.2.2 Filter Level

The data in the 'filter level' were preserved through the filter running as the data was acquired. This implies that they reconstructed with a linefit or dipolefit zenith angle above 70 degrees and had at least 11 hit DOMs after hit cleaning. For these events the likelihood fitter was run using the linefit and dipolefit as seeds.

7.2.3 Level 2

After the filter, events for which the likelihood fitter (seeded by the linefit and dipolefit) returns a zenith angle less than 80 degrees are cut. The remaining events survive to 'level 2'. At level 2, the sobol seed is run with 32 seeds, and the likelihood reconstruction is run again, this time with all 32 seeds from the sobol seed, the linefit and the dipolefit used as seeds. This likelihood reconstruction is the best reconstruction run and constitutes the best estimate, in this analysis, of the muon's direction.

7.2.4 Level 3

Using the 34-seed likelihood reconstruction run at level 2, we reject any events which reconstruct with a zenith angle less than 80 degrees. All other events make up the 'level 3' dataset. This dataset is still dominated by events which are actually down-going but which reconstruct as up-going and by coincident, uncorrelated air showers.

7.2.5 Final Event Selection

The background events that survive to level 3 are not typically high quality events. We can remove these events and preserve neutrino events by cutting out low quality events. A study of several cut parameters typically used by AMANDA analyses suggests that cuts on the direct length L_{dir} and the number of direct hits N_{dir} can be used to reject these background events while retaining enough signal events to perform a measurement of the atmospheric neutrino rate.

The cut values are chosen to maximize the number of detected events while maintaining a purity of true neutrino events of at least 95%. The determination of the the cut values used are discussed in section 7.4. Events survive to the final sample if they reconstruct with $L_{dir} > 250$, $N_{dir} \geq 10$, and a 34-seeded likelihood reconstructed $\theta > 92$.

7.2.6 High Quality Down-going Events

In addition to the high-quality up-going event sample, the same processing and event selection criteria are applied to down-going events from the 'trigger level' dataset to generate a high-quality down-going event sample. These events are not caused by neutrinos, but can nevertheless be used to study the effects of the high-quality cuts on a much more abundant and well-understood source.

7.2.7 Passing Rates

Table 7.1 shows the passing rates of events through the different levels of the analysis. As noted in Section 6.6, the rate of the Single Shower and Double Shower simulation have been corrected by 15% to force them to agree at trigger level, consistent with the overall uncertainty in the cosmic ray air-shower rate.

Table 7.1: Event Passing Rates

	Data	Single Shower	Double Shower	Atmospheric ν
Trigger Level	117.7	117.7	1.51	6.6×10^{-4}
Filter Level	6.09	4.69	0.45	3.7×10^{-4}
Level 2	1.70	1.13	0.32	3.5×10^{-4}
Level 3	0.77	0.46	0.21	3.2×10^{-4}
HQ (down)	11.47	10.05	0.07	–
HQ (up)	2.01×10^{-5}	9.0×10^{-7}	3.22×10^{-7}	1.77×10^{-5}

7.3 Detector Stability

The data acquired is organized into runs, with individual runs lasting anywhere from a few minutes to 4 hours. The software that acquires the data is being developed and is subject to some instabilities which are not fully understood.

In order to investigate the stability of the detector, runs were omitted if they were particularly short or if they had an obviously extreme rate for events passing the Pole filter. Only runs longer than 2000 seconds and with a rate passing the filter between 5.8 and 6.4 Hz were included here. This minimal quality cut results in only a 3% loss of data.

The easiest measure of the detector stability is simply the rate at which events trigger the detector and pass the various processing levels. Figure 7.2 shows these rates. Clearly there are some systematic effects going on. In the small amount of data we have acquired, we do not see the whole annual cycle, but in AMANDA a

periodicity of 1 year has been observed [53] in the muon rate coincident with changes in the temperature and atmospheric overburden at the Pole. The results of IC-9 are consistent with that observation.

Looking at the absolute rate reveals that the rate changes during the year, but in order to quantify the stability it is useful to consider the significance of the deviation from the average. In Figure 7.3a, we see the deviation of the rate for a particular run from the expected average rate. The distribution is quite wide at 3.36σ . If the detector were perfectly stable we would expect a width of this distribution of 1σ . Figure 7.3b shows the time-dependence of this behavior, and it is clear that the width of this distribution is due primarily to the overall change in the rate of events. Because of the large number of recorded events we are sensitive to small variations in the rate.

Because the overall rate of events in the detector is changing, it is useful to consider deviations in the ratio of events passing through the different processing levels. Figure 7.4 shows the deviation in σ for the ratios level0/level2 and level0/level3. Since these values factor out the overall rate change, the histograms show the expected 1σ width. There is some slight residual broadening but overall things look stable.

7.4 Determining the Cuts

The final step of the analysis procedure is to cut out low-quality up-going events to remove any contamination by down-going muons. A survey of several potential cut variables typically used in AMANDA analyses and the variables of L_{dir} and N_{dir} where able to sufficiently reject the background. We also cut on the reconstructed zenith of the event to remove events above the horizon.

Shown in Figure 7.5 are the distributions of the three cut variables chosen for

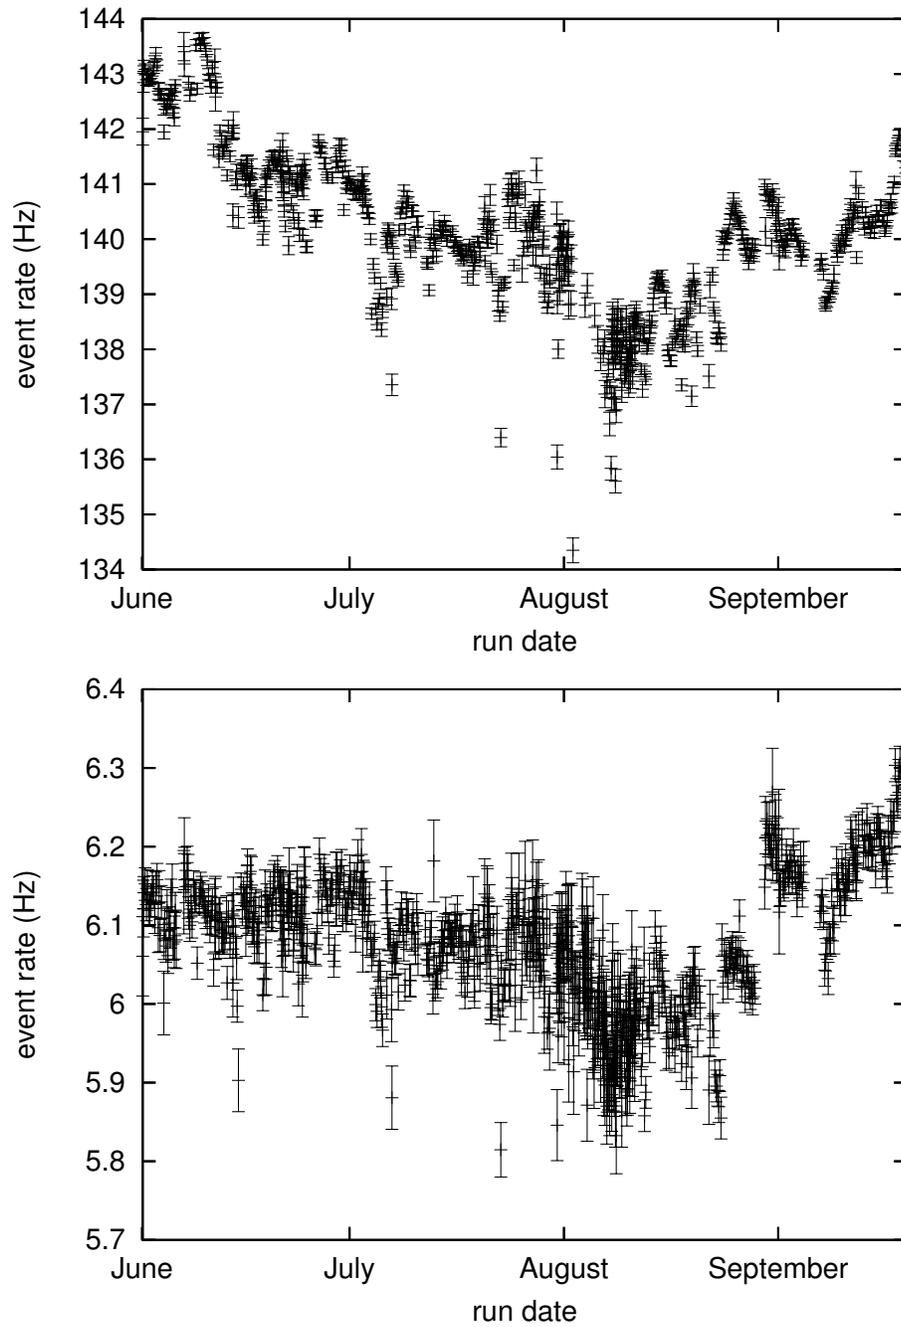
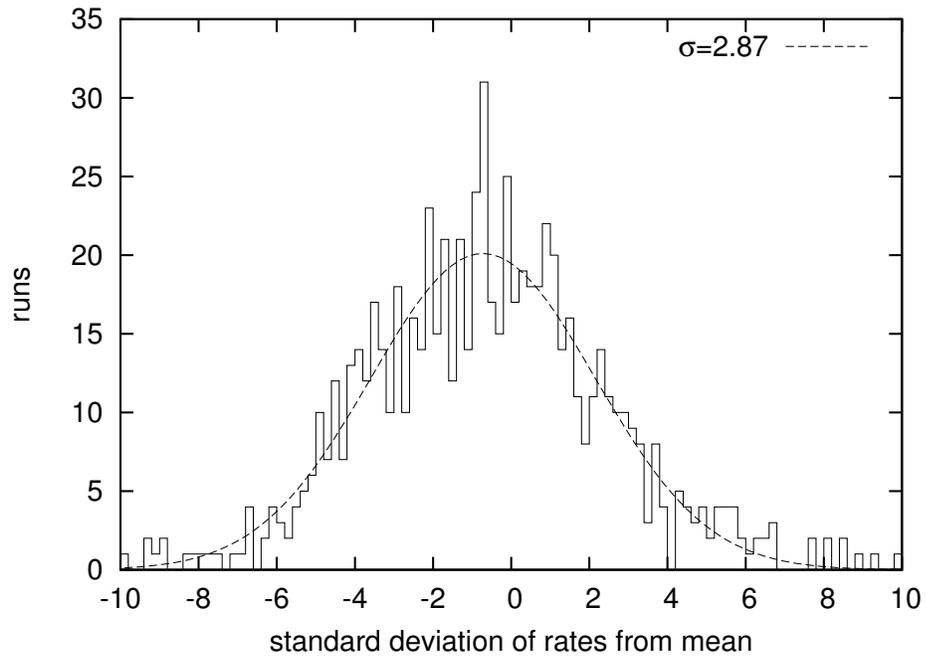
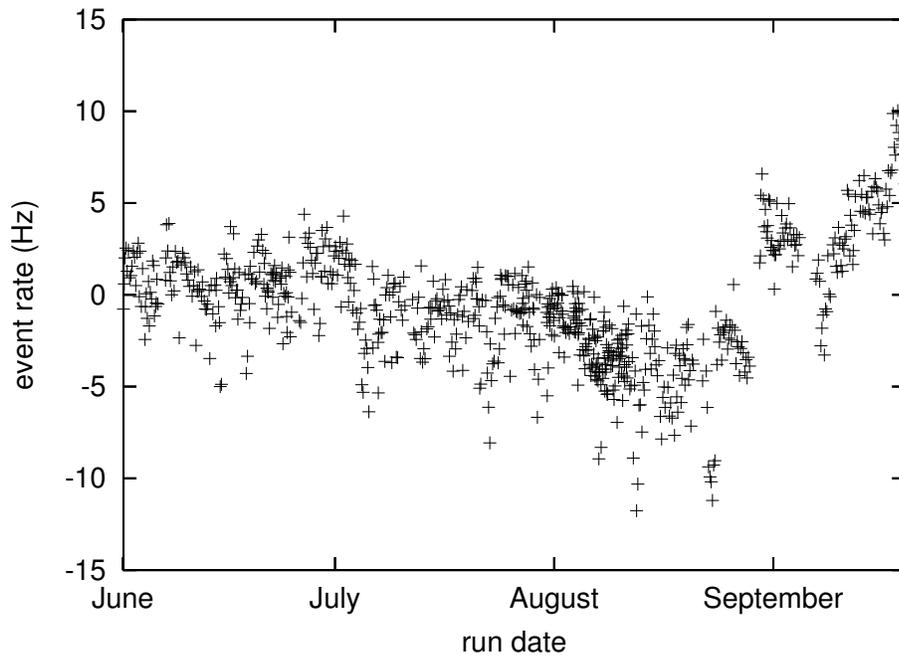


Figure 7.2: The rate of events triggering the detector is shown in panel (a) and the rate passing the Pole filter is shown in panel (b). The rates are given for each run as a function of time in seconds since the start of the year. The error bars shown are the statistical error due to the number of events used to determine the rate.

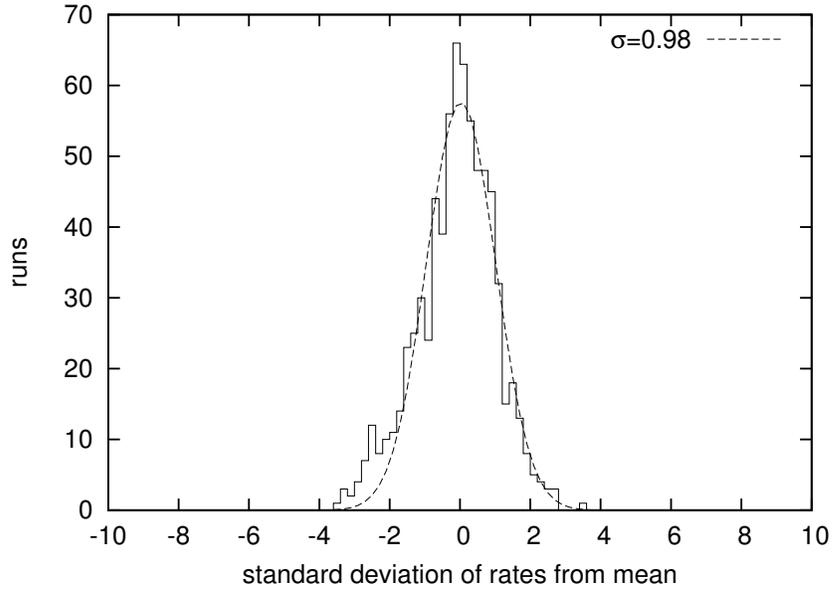


(a)

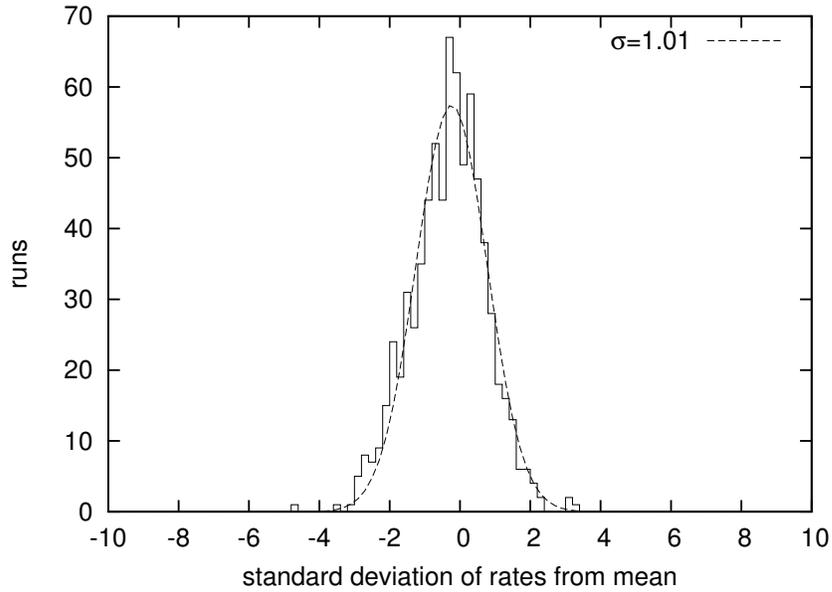


(b)

Figure 7.3: Panel (a) is a histogram of the deviations in the rate of individual runs from the average. Panel (b) shows the time-dependence of these deviations.



(a)



(b)

Figure 7.4: These plots show the deviations from the mean of the ratio of events passing through the level 2 and level 3 processing levels. Panel (a) shows the deviations in the ratio level0/level2 and panel (b) shows the deviations in the ratio level0/level3.

the two background categories (muons from single and coincident air-shower events) and for the signal of up-going neutrino-induced muons.

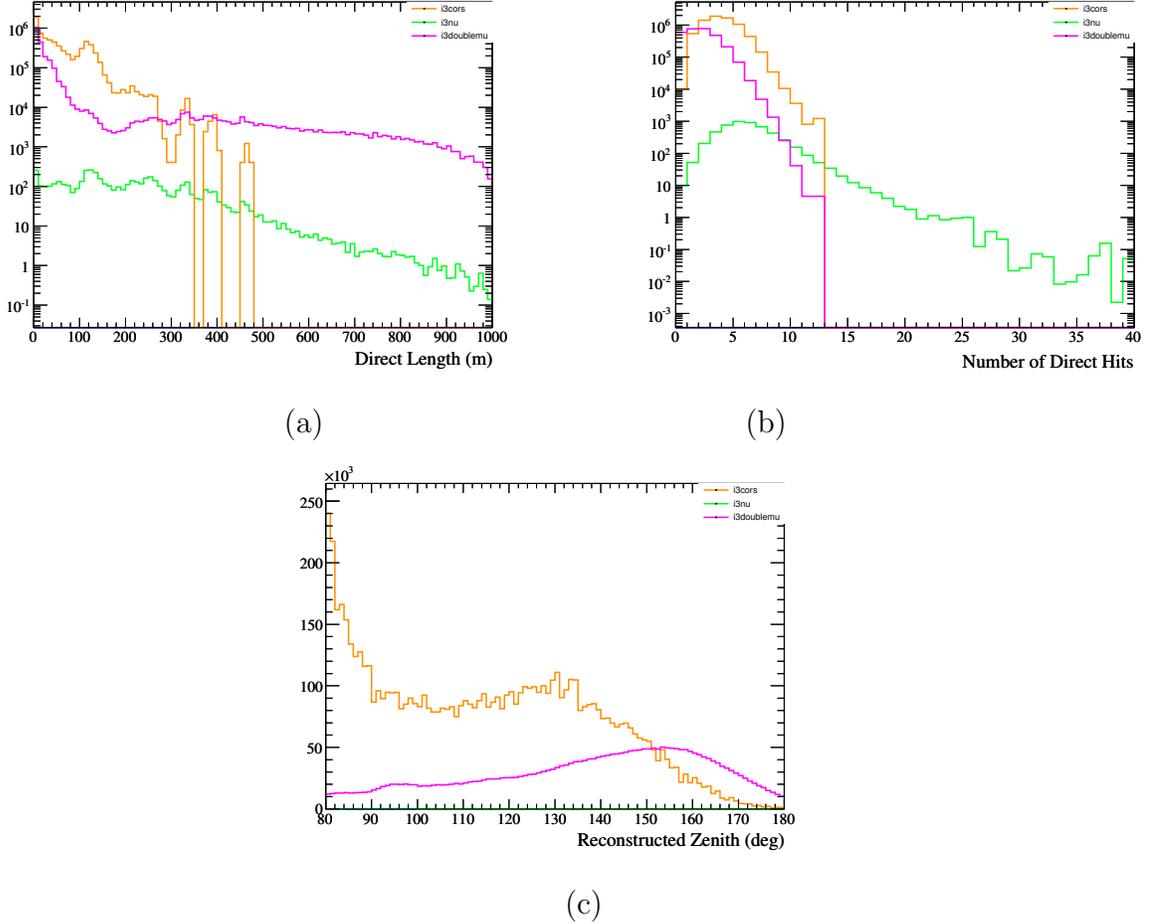


Figure 7.5: The three variables used to separate signal from background at level 3. Panel (a) is the direct length. Panel (b) is the number of direct hits, and panel (c) is the reconstructed zenith angle. The distributions shown are for 180 days of livetime.

The L_{dir} and N_{dir} parameters are clearly able to distinguish signal from background. Also clear is the contamination from events from just above the horizon, so events are cut at a zenith of 92 degrees. It is not as clear where to put the cuts on L_{dir} and N_{dir} , so the approach is to sample the potential L_{dir} and N_{dir} cut space and determine which cuts are optimal. The cut values are chosen so that we maximize the number of signal events while rejecting enough background events to have less

than 5% contamination of our sample by either 'single shower' or 'double shower' CORSIKA events.

One cause for concern is that the statistics of the single air-shower CORSIKA sample are low and so we do not sufficiently probe the tail of that sample. We cannot use the simulation to reliably estimate, for instance, how many background events have at least 15 direct hits because we do not generate enough simulation to sample these events at all.

The number of single-shower background events remaining after a cut at L_{dir} and N_{dir} forms a function of L_{dir} and N_{dir} which was fit. The fit allows us to extrapolate the single-shower background out past where the statistics expire.

The actual measurement from the coincident air-shower background simulation is used without any re-fitting. Since the coincident muons are relatively rare, it is easier to get acceptable statistics in the simulation.

With this procedure we can make the plots shown in Figure 7.6. Figure 7.6a shows the number of signal and background events remaining after cuts at L_{dir} and N_{dir} are applied. Figure 7.6b shows a rather complicated cross section of the plot in Figure 7.6a. Figure 7.6b shows the number of signal events remaining for L_{dir} cuts. The N_{dir} cuts chosen for each point maximize the number of signal events while maintaining a purity of 95%.

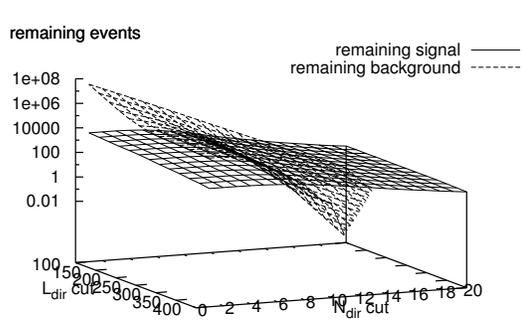
The cuts which overall retain the most events while maintaining a purity of at least 95% are as follows:

$$L_{dir} > 250$$

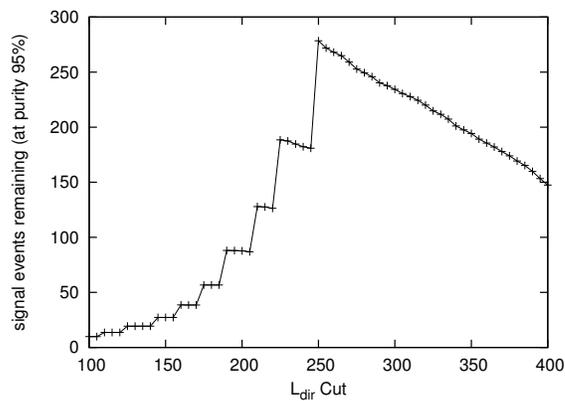
$$N_{dir} \geq 10$$

$$\theta > 92^\circ$$

With these cuts, we expect for 6 months of livetime 278.4 signal events with 8.8 single-shower and 4.6 double-shower background events, for a purity of the final sample of 95.4%.



(a)



(b)

Figure 7.6: These plots demonstrate the selection of the final cuts for the analysis. Panel (a) shows the remaining signal and background as a function of cuts on L_{dir} and N_{dir} . A signal-dominated region is evident. Panel (b) shows the remaining signal events as a function of cuts on L_{dir} where we have applied the optimal cut on N_{dir} .

Chapter 8

Results

8.1 Surviving Events

From the processing in Chapter 7, we have a set of measured events N_{data} surviving cuts during the 90.0 days of recorded livetime. We also have the prediction from three kinds of simulation: single air-shower N_{single} , coincident air-shower N_{double} and atmospheric neutrinos N_ν .

Table 8.1 summarizes the number of events expected or measured in each of these categories in the recorded 90.0 days of livetime. The 156 measured data events are recorded arriving constantly during the exposure time as seen in Figure 8.1. The expected number of events for the double air-shower simulation and atmospheric neutrino simulation have been estimated directly by scaling the simulated livetime to match the 90.0 days of livetime. No events from the single air-shower simulation survive and so the estimate of N_{single} has to be extrapolated from looser cuts as described in Section 8.1.2.

What we want to calculate is the normalization \mathcal{R} of the neutrino simulation implied by our measurement of the atmospheric neutrino rate as given by:

$$\mathcal{R} = \frac{N_{data} - N_{single} - N_{double}}{N_\nu} \quad (8.1)$$

Table 8.1: Number of events passing cuts in 90.0 days.

Category	Passing Events
Single-shower simulation N_{single}	7
Double-shower simulation N_{double}	2.5
Atmospheric ν simulation N_ν	139.1
Data N_{data}	156

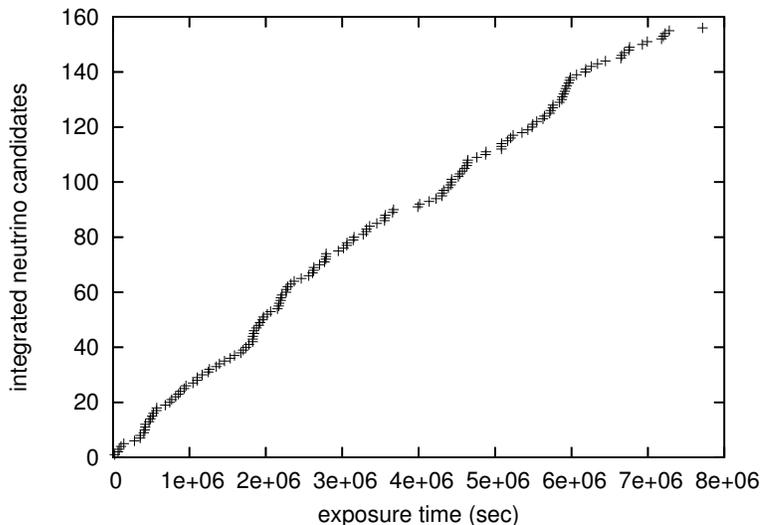


Figure 8.1: The integrated number of neutrino candidates as a function of the exposure time of the detector. Also shown is the expectation if events were arriving at a perfectly constant rate. The deviations are small and consistent with a constant rate of arrival.

For the numbers given in Table 8.1 we arrive at a measurement of \mathcal{R} . The uncertainty in this measurement of \mathcal{R} is summarized in Section 8.2.

8.1.1 Cut Strength

It is useful to define the quantity *cut strength*. The cut strength is a number that quantifies how hard we are cutting on the data. The final cuts are $L_{dir} > 250$, $N_{dir} \geq 10$, and $\theta > 92$. We will define the cut strength S_{cut} such that the cuts at a cut strength S_{cut} are given by

$$L_{dir} > 25 \cdot S_{cut}$$

$$N_{dir} \geq S_{cut}$$

$$\theta > 92$$

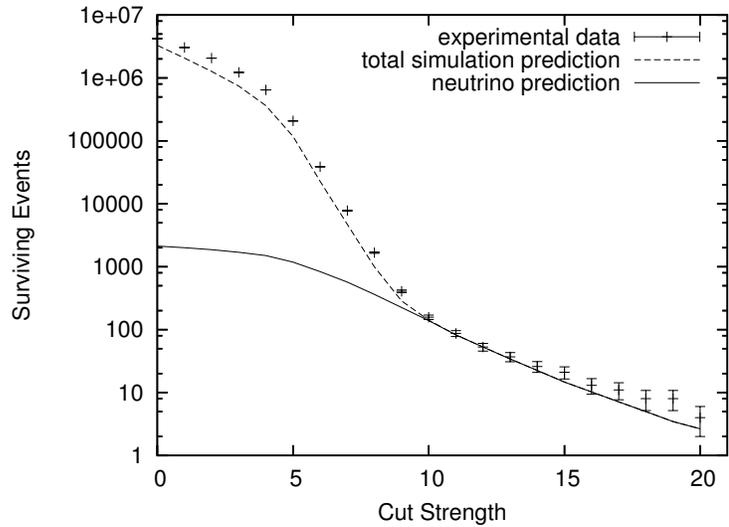
In this way the events only cut on θ when $S_{cut} = 0$ and the final analysis cuts are placed at $S_{cut} = 10$.

It is useful to look at the populations of data and simulation remaining as the cuts are turned up. Figure 8.2 shows how many data and simulated events survive as a function of S_{cut} . Shown are three things: the experimental data, the total simulation prediction (single showers, double showers and neutrinos) and the neutrino prediction by itself. When we get to a cut strength of about 10, we transition from background-dominated to signal-dominated samples. In the background-dominated region, the disagreement between data and simulation is of the order 60% and peaks at 80%. When we have removed all the background and only high quality up-going events remain, the agreement improves. This disagreement between data and simulation for mis-reconstructed events is significant. Note that the relative agreement between data and simulation for high-quality down-going events (as seen in Table 7.1) gives us confidence that the simulation models high-quality events satisfactorily. The disagreement is in how the ice properties give rise to mis-reconstructed events. Since we have relatively few mis-reconstructed events in the final sample, this systematic error due to disagreement does not impact the final results much.

8.1.2 Single-Shower Contamination Estimate

No events from the single-shower simulation survive the cuts. Because the neutrino events are so rare and the single-shower events are so plentiful, the final analysis cuts are made at the very tails of the single-shower distributions. We must run a prohibitive amount of single-shower simulation in order to have events survive the cuts.

In order to estimate the single-shower contamination after analysis cuts, we make



(a)

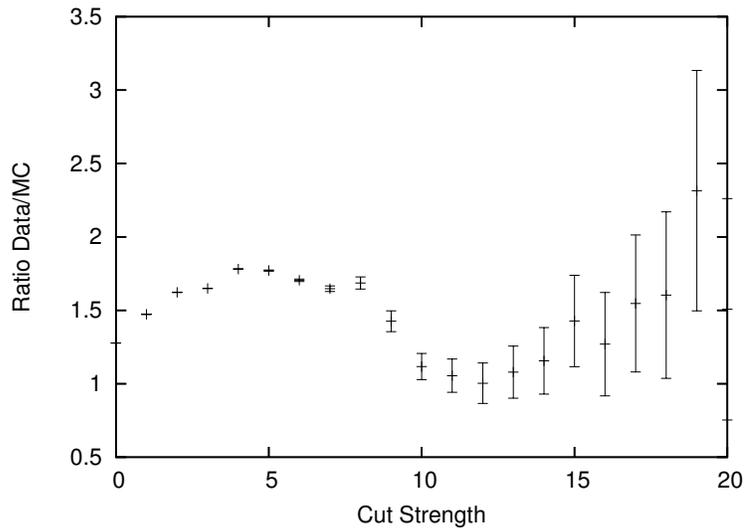


Figure 8.2: Shown in panel (a) is the remaining number of events as the strength of the cuts applied are varied. There are points shown both for data and simulation. Panel (b) shows the ratio of data/MC as a function of cut strength. The error bars on the data are purely statistical and no errors are shown for simulation. The analysis cuts are at a cut strength of 10.

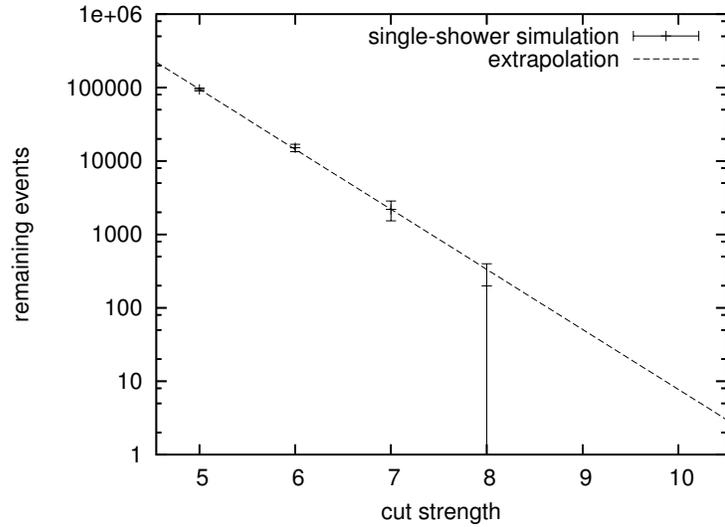
the plots shown in Figure 8.3. Figure 8.3a shows the remaining number of single-shower events as a function of cut strength. We run out of simulation at a cut strength of 8 where there is one simulated event left, which scales to 200 events for 90.0 days. In order to estimate the single-shower contamination at a cut strength of 10 where the analysis cuts are, we fit the curve in Figure 8.3a and extrapolate out to a cut strength of 10. The fit gives some nominal errors for the parameters and varying these parameters within their reported errors gives a distribution of single-shower contamination estimates shown in Figure 8.3b, from which we can estimate the error in this contamination.

For the estimate of the single-shower contamination in our event sample, we take 7 ± 13 single-shower events in the dataset. The error comes by considering the point at which 95% of the independent parameter trials fall. This estimate is conservative in order to capture the fact that we are in fact estimating the single-shower contamination by extrapolating beyond what has been estimated from Monte Carlo.

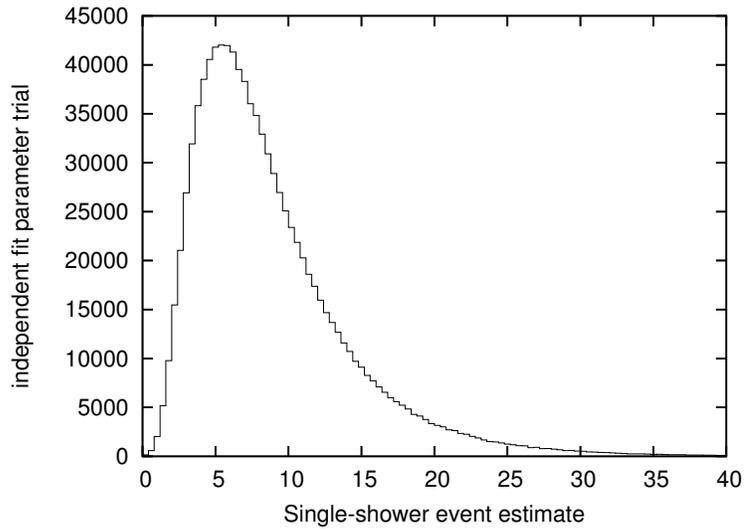
8.1.3 Comparison to Simulation

Figure 8.4 exhibits the agreement between data and simulation for the analysis neutrino-candidate sample. Shown are the data, the atmospheric neutrino simulation and the one surviving coincident air-shower event. Recall that, though it is not shown on the plots, there are expected to be about 7 ± 13 events from single air-showers (as predicted from the extrapolation in Section 8.1.2). In general, the data agree well with some slight excesses near the horizon and at low L_{dir} . It is likely that these excesses come from the single-shower background.

One exception to the generally good agreement between data and simulation is in quantities that measure how much light is deposited at different depths. The two quantities typically considered are the occupancy and the average z position of hits, as seen in Figure 8.5. The discrepancy between data and Monte-Carlo is likely



(a)



(b)

Figure 8.3: Panel (a) shows the single-shower prediction as a function of cut strength with error bars corresponding to the statistical error due to low simulation statistics. When the identified extrapolation is allowed to have its parameters drift by the errors from the fit, we get panel (b) which shows how many single-shower events are estimated to remain at a cut strength of 10.

due to poor modeling of the ice below 2200 meters (-250 meters in detector-centered coordinates). AMANDA is not deployed below 2200 meters and so IceCube is probing the ice in an area unexplored by AMANDA.

8.2 Systematic Errors

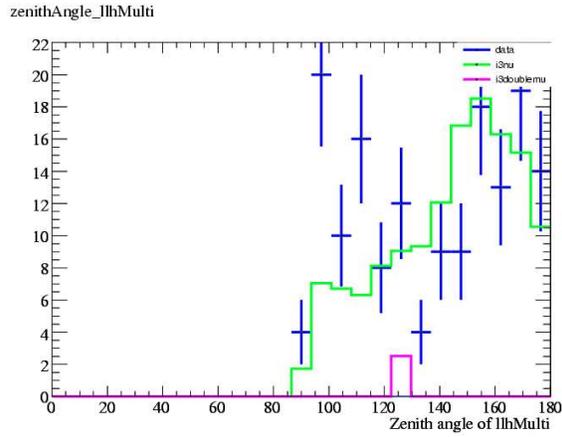
The final result is the ratio \mathcal{R} of the experimentally measured number of events to the number expected from simulation. The effect of the systematic error sources considered are quoted as a percent error in our measurement of \mathcal{R} .

8.2.1 Timing Uncertainty

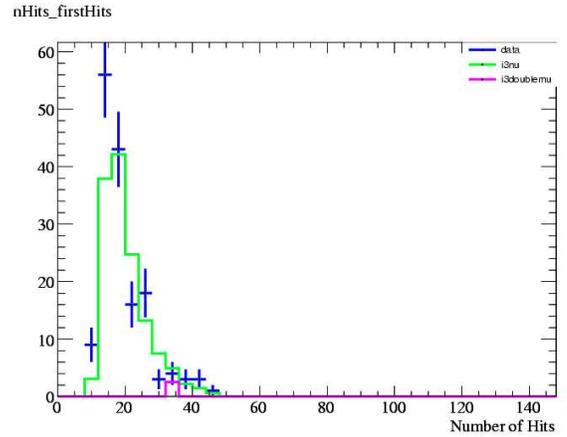
There are several factors that contribute to uncertainty in the hit times. The RAP-Cal procedure and Feature Extraction algorithms both introduce uncertainty in the timing in addition to intrinsic uncertainty due to variations in the photo-electron transit time from the photocathode to the dynode chain. Furthermore, we can represent uncertainties in the DOM positions as timing uncertainties of order 1.5 ns. In order to assess the overall uncertainty in the timing, the timing residuals of high-quality down going muons have been studied in addition to studies using the flashers. The conclusion of these studies is that the absolute time of the PE arrival is known to about 2-3 ns with some outliers at 5 or 6 ns [38].

In order to study the effect of this timing uncertainty, the full analysis processing has been run on the neutrino simulation after introducing an artificial jitter in the timing of individual hits. This timing difference causes some events to pass the cuts which might otherwise not have passed and also causes us to lose some events. By performing this processing a number of times, we can get an idea of the uncertainty in the measured number of neutrino events due to timing uncertainty.

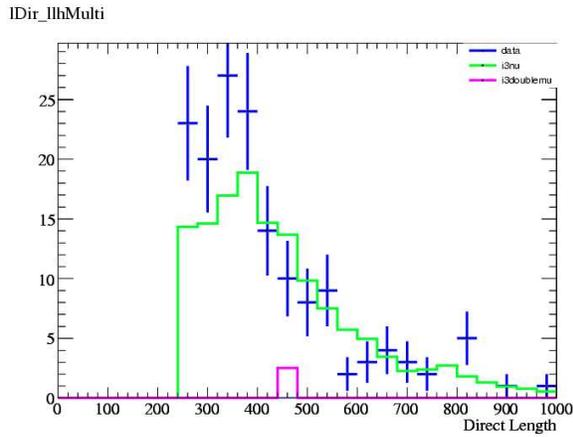
Figure 8.6 shows the results of 25 independent processings of the neutrino simu-



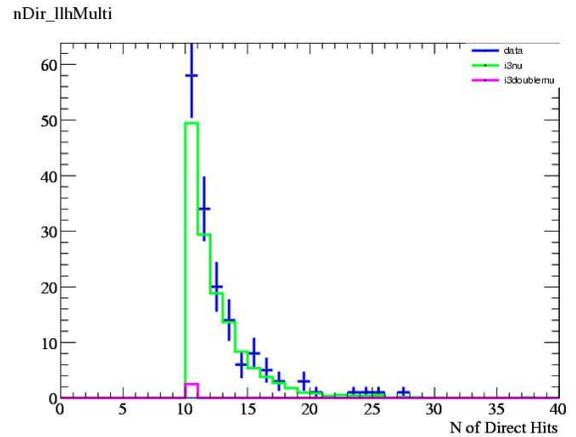
(a)



(b)

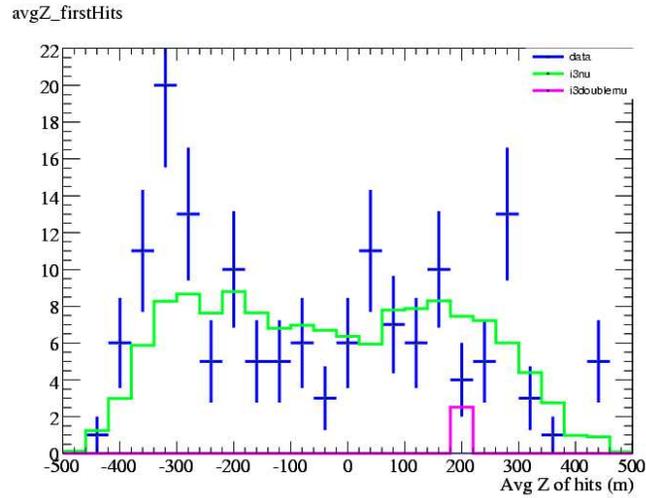


(c)

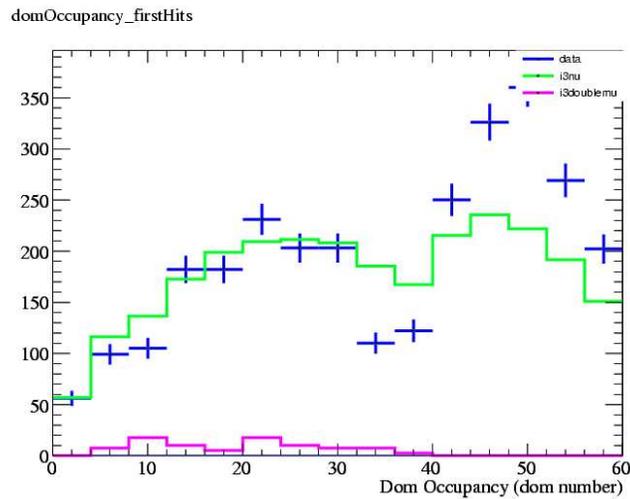


(d)

Figure 8.4: Comparing experimental neutrino candidates to simulation. Panel (a) shows the reconstructed zenith angle. Panel (b) shows the number of hit DOMs in the event. Panel (c) shows the L_{dir} and panel (d) shows N_{dir} .



(a)



(b)

Figure 8.5: Panel (a) shows the averaged position of hits in the events from the final event selection. Panel (b) shows the position of all hit DOMs in events in the final event selection.

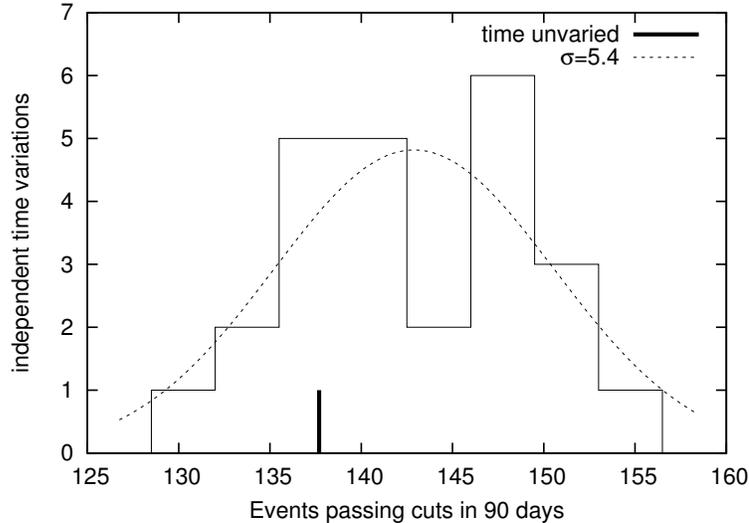


Figure 8.6: Histogrammed is the number of events predicted in 90.0 days in independent processings of the neutrino simulation, allowing the timing of hits to vary by 4 ns. The best fit is shown with a width of 5.4 events. Also shown is the prediction without varying the hit times at all.

lation. In each processing of the simulation, each hit of every event was modified by some random time drawn from a Gaussian distribution with a width of 4 ns. Shown is the number of events which are predicted to pass all the cuts for 90.0 days of livetime for each processing. The width of the distribution is 5.4 events, suggesting an error in \mathcal{R} due to timing of less than 5%.

8.2.2 Uncertainties in Muon and Neutrino Interaction and Propagation

The simulation of muon and neutrino propagation and interaction includes several sources of error which affect the neutrino prediction. Principal among these are uncertainties in the density of the rock below the detector and uncertainties in the neutrino-nucleon cross-section. These errors have been estimated for a search of point sources of neutrinos with AMANDA [25]. Since IceCube and AMANDA use the same software for muon and neutrino interaction and propagation, it is safe to take the AMANDA estimates. These estimates place the errors in \mathcal{R} from the

neutrino interaction modeling at $\pm 3\%$ and the estimated error in \mathcal{R} due to muon propagation at $\pm 1\%$

8.2.3 Error in Background Estimates

From the low statistics in the background simulation we get a systematic error in the number of events predicted to survive to the final level. In the case of single-showers, we predicted in Section 8.1.2 a contamination of 7 ± 13 events in the final sample. This gives us a error of $\pm 3.5\%$ in \mathcal{R} .

In the case of double-shower simulation, the event rate is predicted from one double air-shower event surviving the cuts. This one event scales to 2.5 events when scaled to 90.0 days livetime. We can take an error of 100% on the predicted number of double air-shower events to survive the cuts, giving a error of about $\pm 2\%$ in \mathcal{R} .

8.2.4 Uncertainty in Light Propagation Simulation

The light propagation simulation is a significant source of error. The simulation of light propagation includes simulation of the depth-dependent scattering and absorption due to dust which has been deposited over the years. Furthermore, though the ice at IceCube depths is remarkably clear, the effect of bubbles that may have formed around the DOMs during deployment is uncertain. Also included in this uncertainty is any error due to uncertainty in the DOM sensitivity.

In order to consider how variations in the light propagation simulation will change estimates of our neutrino signal, we consider the final analysis cuts. We chose cuts of $N_{dir} \geq 10$ and $L_{dir} > 250$. The plots in Figure 8.7 show how much signal remains if we varied the cuts slightly. All of our cuts are based on direct hits, hits which arrive at the DOM unscattered. If we modify the ice properties, we modify the behavior of direct hits, either increasing or decreasing the occurrence of direct hits. We can see that if we were to remove 1 direct hit from each event, the number of signal events

passing the cuts could change by more than 30%. If we were to modify the L_{dir} of events by 20 meters we can change the number of expected signal events by 15%.

In order to gauge the uncertainty due to the modeling of light propagation this way, we need to gauge what constitutes a reasonable variation in the N_{dir} and L_{dir} of events. This is done by looking at the quality of the background simulation. The discrepancy between data and simulation at level 3 is expected to be primarily due to incorrect modeling of the light propagation.

Figure 8.8 helps quantify what constitute reasonable variations in N_{dir} and L_{dir} . The plot shows the amount of remaining data as a function of cut strength. In this plot the S_{cut} of the MC events has been shifted by $S_{cut}^{new} = S_{cut}^{old} + 0.5$. This corresponds to a change in N_{dir} of 0.5 and a change in L_{dir} of 12.5 meters. With this modification, the agreement between data and simulation is much better. Figure 8.8 should be compared to Figure 8.2 to see the improvement. This suggests that our modeling of the direct light is inaccurate at the 0.5 direct hits level.

Looking again at Figure 8.7, if we were to consider a change of N_{dir} of 0.5 we would get a 20% variation in the number of signal events passing. This dominates any change in L_{dir} and suggests that we have a 20% error in \mathcal{R} due to the simulation of the light propagation. This estimate may be conservative because we have estimated this uncertainty based on low-quality down-going events rather than the high-quality events that constitute our signal and we expect the high-quality sample to be better simulated.

Note that this approach is insensitive to the actual number of measured signal events because we have calculated this error based on the data when it is still background-dominated. We are using the down-going muon events to help estimate the errors in the up-going flux.

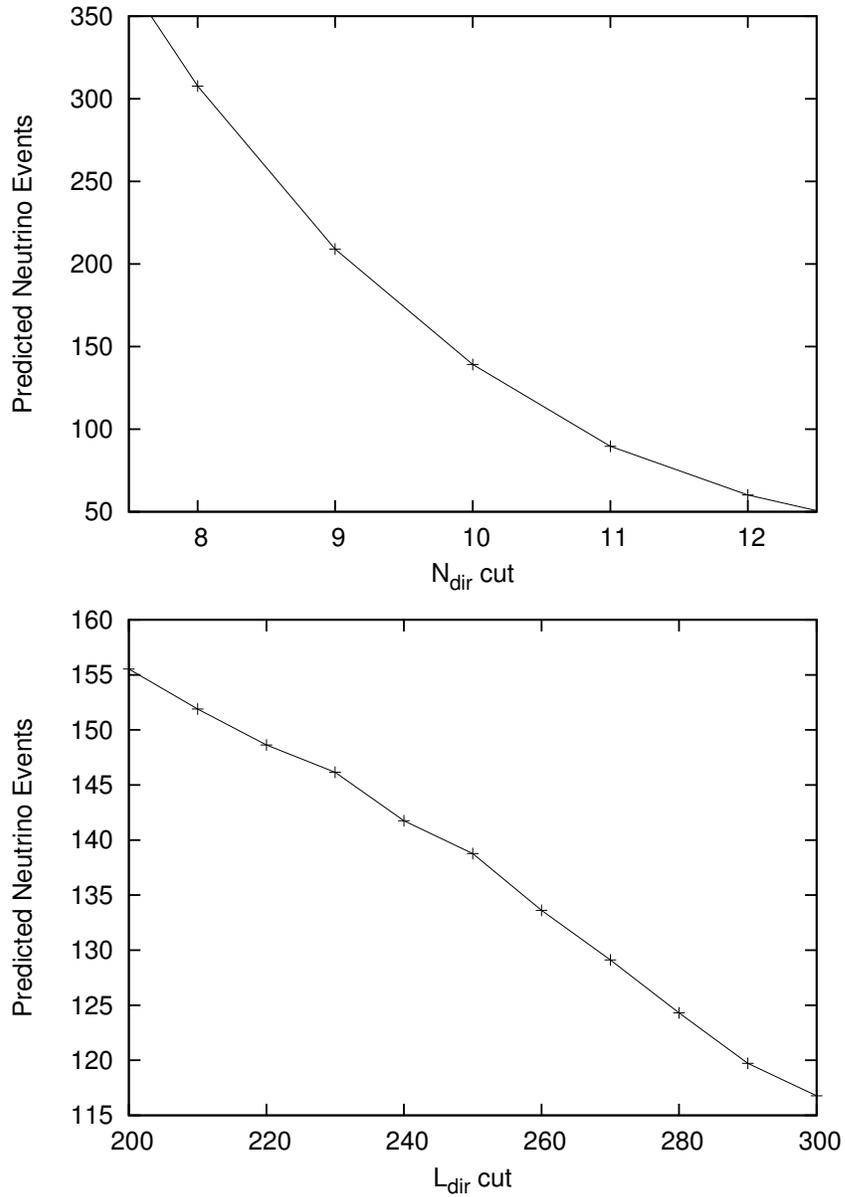


Figure 8.7: Panel (a) shows the predicted number of remaining neutrino events as a function of the cut on N_{dir} after the nominal L_{dir} and θ cuts have been applied. Panel (b) shows the predicted number of remaining neutrino events as a function of the cut on L_{dir} after the nominal cuts on N_{dir} and θ have been applied. The figures give us an idea of how stable the predicted number of neutrino events are to our modeling of direct light.

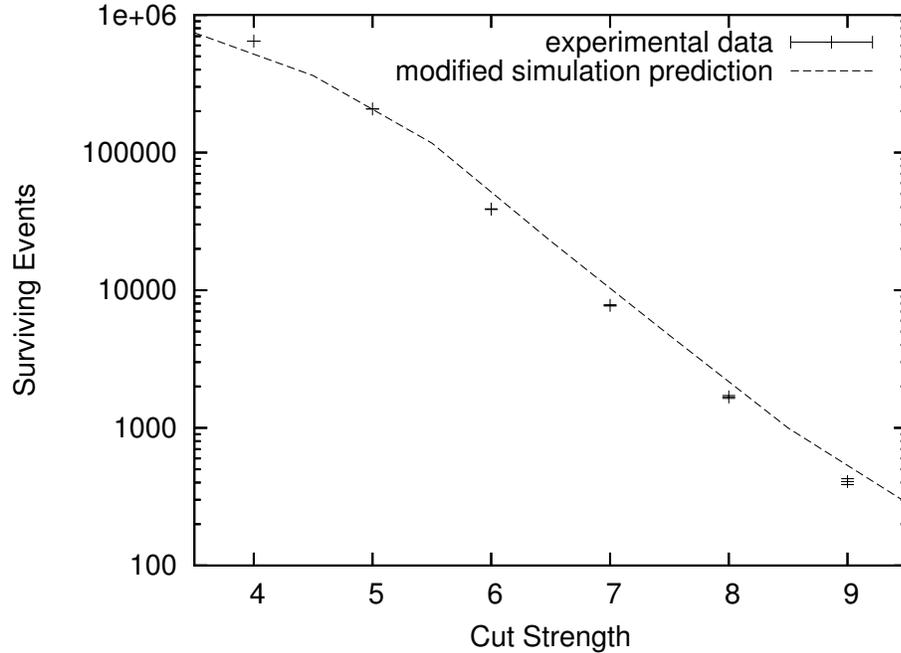


Figure 8.8: Shown is the remaining amount of data as a function of cut strength. The simulation prediction has been shifted by a cut strength of 0.5 to show that our modeling of direct hits is uncertain.

8.2.5 Error Summary

Table 8.2 shows the systematic uncertainties used in the determination of \mathcal{R} . In considering the error in N_ν , we do not use the theoretical $\sim 30\%$ error in the overall normalization of the atmospheric neutrino spectrum since that is what we are measuring. The largest source of error is due to the light propagation simulation.

8.3 Final Results

Totaling the errors we arrive at an answer for \mathcal{R} as

$$\mathcal{R} = 1.05 \pm 0.24_{syst} \pm 0.09_{stat} = 1.05 \pm 0.26_{tot}$$

Figure 8.9 compares this measurement to the previous measurements of Frejus [34] and AMANDA [17]. Since, in this work, the energy spectrum was not measured, we show the energy spectrum from the *bartol* model assumed in the simulation, with

Table 8.2: Error Summary

Source of Error	Error
Muon Propagation	$\pm 1\%$
Timing Uncertainty	$\pm 5\%$
Neutrino Cross-Section	$\pm 3\%$
Light Propagation	$\pm 20\%$
Single-Shower Contamination	$\pm 9\%$
Double-Shower Contamination	$\pm 2\%$
Data/Sim Background Disagreement	$\pm 6\%$
Total Systematic Error	$\pm 24\%$
Total Statistical Error	$\pm 9\%$
Total Error in \mathcal{R}	$\pm 26\%$

the normalization given by our measurement of $\mathcal{R} = 1.05 \pm 0.26$. Interpreting these results in light of previous measurements is difficult because other measurements actually estimate the energy spectrum and we have only reported an overall normalization. However, within the limits of such a comparison, the agreement is good. The theoretical uncertainty in the predicted neutrino flux is about 30% and the errors in our results prevent constraining atmospheric neutrino models. Nevertheless, it is clear that IceCube is seeing atmospheric neutrinos at the expected level.

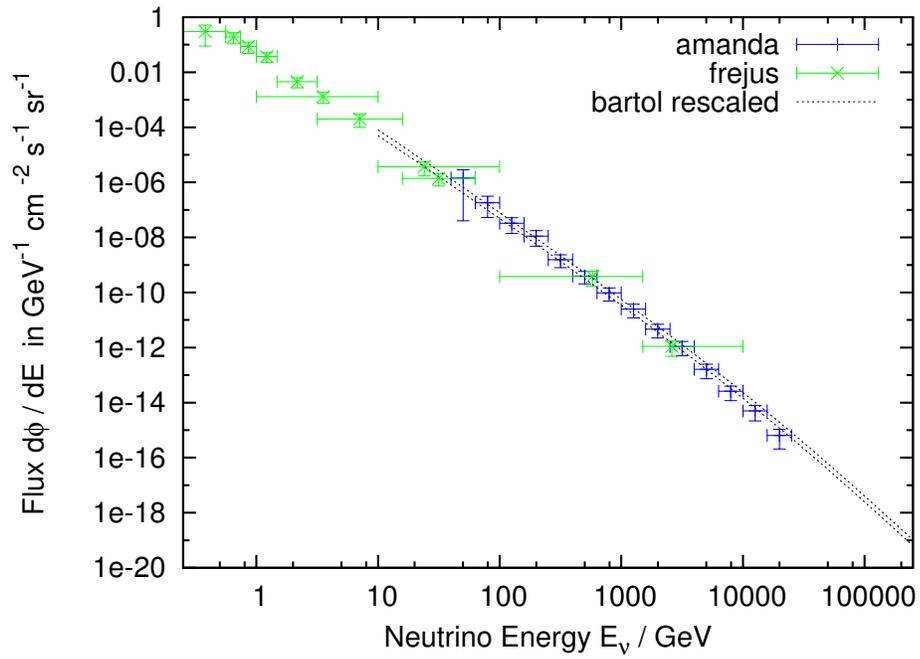


Figure 8.9: Experimental measurements of the atmospheric neutrino spectrum from the Frejus experiment and the AMANDA experiment. Overlaid is the bounds of the *bartol* flux, with the normalization coming from this work.

Chapter 9

Summary and Outlook

9.1 First Steps with IceCube

Neutrino astronomy offers a new window on the universe and will explore phenomena inaccessible by conventional astronomy. The natural scale for neutrino astronomy has long been known to be a cubic kilometer [9]. With the deployment of IceCube, high-energy neutrino astronomy begins in earnest.

The calibration beam for potential astrophysical sources of neutrinos is the relatively well-understood atmospheric neutrino population. It is with atmospheric neutrinos that IceCube will understand its energy response. In addition to providing calibration, atmospheric neutrinos constitute a background to astrophysical observations. In order to accomplish the goal of neutrino astronomy we must understand the atmospheric neutrino population and how IceCube responds to it.

This thesis has taken the first few steps toward understanding this new instrument. We have established that IceCube can reconstruct muons and reject the down-going cosmic-ray muon background while preserving atmospheric muon neutrinos for study. Atmospheric neutrinos in IceCube are seen at the expected level. The ratio between atmospheric neutrino predictions and data was measured to be $\mathcal{R} = 1.05 \pm 0.24_{\text{sys}} \pm 0.09_{\text{stat}}$. Additionally, the simulation and reconstruction al-

gorithms have been investigated. Work is required to perfect the simulation and the reconstruction algorithms require re-tuning for the new IceCube geometry and hardware, but the initial performance of the AMANDA-era tools is suitable.

One notable characteristic of this IceCube analysis is how remarkably stable and clean the data has been. Significant engineering effort has been expended in the design of the DOM and DOR cards and it has paid off. Typical AMANDA analyses have had to work hard to remove unsimulated cross-talk between the electrical channels as well as to remove poorly-understood non-physical flares in the event rate due to electronic noise at the surface. Furthermore, AMANDA analyses have typically been delayed for more than year while the data is hand-carried to the North and reprocessed. IceCube data has none of these problems and, consequently, this first analysis has been conducted exclusively on data transmitted over the satellite with less than six months turn-around. This is a significant accomplishment for the IceCube collaboration.

9.2 Future Analyses

This work is very forward-looking and one of the main focuses has been to understand and identify what areas the IceCube collaboration must address in its future work. This includes priorities for future atmospheric neutrino analyses as well as analyses looking for other potential signals.

9.2.1 Ice Simulation

The simulation of light propagation in the presence of dusty ice is a significant hurdle that must be addressed. The disagreement between data and simulation at high cut levels is believed to be due to incorrect modeling of the ice because the occupancy

and average hit z disagreement tracks the ice properties. Systematics studies of the ice are essential.

9.2.2 Multi-Muon Reconstruction Algorithms

The rate of coincident air-shower events triggering IC-9 is of order 1%, but as the detector grows this effect will become more important. A back-of-the-envelope calculation suggests it could be as high as 10% for the full array at trigger level. The approach in this analysis has been simply to reject these events, but in the full IceCube detector it will be possible to resolve these events and reconstruct them separately. Development of these separation algorithms is important because the strategy of simply rejecting coincident events will result in a large loss of signal when signal events are coincident with a cosmic-ray event.

9.2.3 Waveforms

This analysis has largely ignored the waveforms that come from the DOMs. The chosen hit cleaning leaves us with only one hit time per hit DOM meaning that we are largely insensitive to the details the waveforms provide. This means that the fine details of the waveform simulation have not been explored and the full power of having the waveforms available has not been realized. Future analyses should begin to take full advantage of the waveform information. This means, among other things, verification that the simulation of the DOM response and the waveform capture is satisfactory. The performance of the feature extraction algorithms on simulation will also be necessary.

9.2.4 Online Filtering

The online filtering used in this analysis is an initial pass aimed at obtaining first results. The filtering requires an artificial threshold of 11 hit DOMs while we trigger on 8 hit DOMs. This requirement is imposed simply to meet the bandwidth requirement. We want to continue this precedent of analyzing data transmitted over the satellite. The filter must be improved in order to extend the low-energy reach of IceCube. Improvements will be absolutely necessary considering that the size of the detector will be constantly increasing until 2011. There are several ways to increase the effectiveness of the online filter.

- *Waveform Compression at the Pole.* In the current setup, the waveform acquired at the Pole is transmitted to the North as it is acquired. The waveforms are large and constitute the bulk of the data sent north. Waveforms could be compressed further at the Pole either by feature-extracting SPE waveforms and discarding the waveform for these simple DOM readouts or by employing a smart compression algorithm which makes use of the fact that the digital waveform is mostly 0 after the ATWD pedestal is subtracted. Either of these approaches significantly reduces the amount of data that needs to be sent to the North.
- *Online Likelihood Reconstruction.* The likelihood reconstruction is quite fast when using 1 or 2 seeds. The original fast-guess methods employed by AMANDA were necessary when AMANDA started taking data in 1997, but since then computers have increased speed more than an order of magnitude. It may be possible to run the full likelihood reconstruction on all events online. Perhaps the likelihood reconstruction could be run only on events which are up-going according to the first-guess methods.
- *Smart First-Guess Methods.* AMANDA has two reconstructions which were

employed for background rejection at the Pole: Direct Walk and Just Another Muon Search (JAMS). The reconstructions were better at rejecting background than either the linefit or the dipolefit and achieved that performance by being carefully tuned to reject background. JAMS is particularly good at rejecting background but is built using a neural-net to decide the final answer. The optimal parameters for AMANDA are not the same as for IceCube. In fact, the performance of these highly-tuned AMANDA fitters on IceCube events is quite poor. When these or other smart fitters become available in IceCube, they may assist in background reduction and can replace the linefit and dipolefit.

9.3 The Future of IceCube

The future of IceCube is bright. The detector is partially deployed and working well. Already the instrumented volume is larger than AMANDA, and will increase steadily. Table 9.1 summarizes the deployment plan until 2011. IceCube plans to deploy 12 strings during 2006-2007 season and then 14 strings each year after that. In the last year, a minimum of 7 strings are planned to be installed bringing the total to at least 70 strings. Plans include the possibility of deploying up to 80 total strings if things go well during the remaining construction years. Also shown is the integrated exposure at the end of each season, and we reach a $\text{km}^3 \cdot \text{year}$ sometime in 2008, assuming 100% uptime. If we scale the rate of neutrinos measured in this analysis, we will have some 7800 neutrino candidates at that point. However, it is likely that our efficiency will be higher by then and it is reasonable to expect some 10^4 neutrino candidates before 2009, double the current AMANDA dataset.

The plan also calls for the integration of the AMANDA DAQ with the IceCube online system during the 2006-2007 season, giving the unified detector a large instrumented volume combined with a high-density sub-detector.

Table 9.1: The plan for IceCube deployment.

Season	Strings	Instrumented Volume (km^3)	Instrumented Exposure ($\text{km}^3 \cdot \text{year}$)	Exposure including AMANDA ($\text{km}^3 \cdot \text{year}$)
2004-2005	1	–	0.01	0.11
2005-2006	9	0.11	0.12	0.24
2006-2007	21	0.26	0.39	0.51
2007-2008	35	0.43	0.83	0.95
2008-2009	49	0.61	1.44	
2009-2010	63	0.78	2.23	
2010-2011	70+	0.96	3.10	

With more than an order of magnitude improvement over AMANDA, the IceCube experiment is poised to search for extraterrestrial sources of high-energy neutrinos from gamma-ray bursts, active galactic nuclei, microquasars and other potential sources. Somewhere in the universe, cosmic-rays are accelerated to macroscopic energies and we expect neutrino production in those sources. Neutrinos are the only practical messenger particle above 100 TeV and IceCube is – for the time being – the only detector that can look for this expected neutrino emission.

Bibliography

- [1] J. Alcaraz et al. The AMS Collaboration. *Physics Letters B*, 490:27–35, 2000.
- [2] T. Sanuki et al. *The Astrophysical Journal*, 545:1135–1142, 2000.
- [3] M. Boezio et al. *Astroparticle Physics*, 19:583–604, 2003.
- [4] J. Hoerandel. *Astroparticle Physics*, 19:193–220, 2003.
- [5] J. Cronin et al. *Scientific American*, 276:44, 1997.
- [6] T.K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990.
- [7] E. Fermi. *Physical Review*, 75:1169–1174, 1949.
- [8] A.M Hillas. *Annual Review of Astronomy and Astrophysics*, 22:425–444, 1984.
- [9] T.K. Gaisser, F. Halzen, and T. Stanev. *Physics Reports*, 258:173–236, 1995.
- [10] J. Ahrens. *Astroparticle Physics*, 20:507–532, 2004.
- [11] F. Halzen. *astro-ph/0506248*, 2005.
- [12] F.W. Stecker and C. Done. *Physical Review Letters*, 66:2697–2700, 1991.
- [13] K. Mannheim. *astro-ph/9703184 v1*, 1997.
- [14] F. Halzen and D. Hooper. *Astroparticle Physics*, 23:537–542, 2005.
- [15] A. Levinson and E. Waxman. *Physical Review Letters*, 87(17), 2001. 171101.

- [16] P. Mesaros and E. Waxman. *Physical Review Letters*, 87(17), 2001. 171101.
- [17] The AMANDA Collaboration. *Proceedings of the 28th International Cosmic Ray Conference*, HE 2.3:1313–1316, 2003.
- [18] W. Rhode et al. *Astroparticle Physics*, 4:217, 1996.
- [19] M. Ambrosio et al. (MACRO Collaboration). *Astroparticle Physics*, 19:1, 2003.
- [20] J. Ahrens et al. (The AMANDA Collaboration). *Physical Review Letters*, 90:251101, 2003.
- [21] M. Ackermann et al. (The AMANDA Collaboration). *Astroparticle Physics*, 22:339, 2004.
- [22] M. Ackermann et al. (The AMANDA Collaboration). *Astroparticle Physics*, 22:339, 2004.
- [23] V. Aynutdinov et al. (Baikal Collaboration). *Astroparticle Physics*, 25:140, 2006.
- [24] E. Waxman and J. Bahcall. *Physical Review D*, 59:023002, 1998.
- [25] The IceCube Collaboration: A. Achterberg et al. Five Years of Searches for Point Sources of Astrophysical Neutrinos with the AMANDA-II Neutrino Telescope. *Physical Review D.*, 2006. In preparation.
- [26] M. Ambrosio et al. (MACRO Collaboration). *Astrophysical Journal*, 546:1038, 2001.
- [27] T. Montaruli et al. for the ANTARES Collaboration. *Acta Phys. Polon.*, B36:509, 2005. arXiv:hep-ex/0410079.
- [28] L.V. Volkova. *Soviet Journal of Nuclear Physics*, 31:784, 1980.

- [29] G. Battistoni et al. *Astroparticle Physics*, 19:269–290, 2003. Erratum-ibid. 291-294.
- [30] G.D. Barr et al. *Physical Review D*, 70(023006), 2004.
- [31] M. Honda et al. *astro-ph/0404457 v4*, 2004.
- [32] J.G. Learned and K. Mannheim. *Annual Review of Nuclear and Particle Science*, pages 679–749, 2000.
- [33] T.K. Gaisser and M. Honda. *Annual Review of Nuclear and Particle Science*, 52:153–199, 2002.
- [34] K. Daum et al. *Zeitschrift Fur Physik C*, 66:417–428, 1995.
- [35] H.L. Lai et al. *European Physical Journal*, C12:375–392, 2000.
- [36] D. Chirkin and W. Rhode. Muon Monte Carlo: A High-Precision Tool For Muon Propagation Through Matter. *AMANDA Technical Report*, 2004.
- [37] The IceCube Collaboration: M. Ackerman et al. *Journal of Geophysical Research*, 111(D13203), 2006.
- [38] IceCube Collaboration: A. Achterberg et al. *Astroparticle Physics*, 26:155–173, 2006.
- [39] S. Kleinfelder. *A Multi-Gigahertz Analog Transient Waveform Recorder Integrated Circuit*. PhD thesis, University of California, Berkeley, 1992.
- [40] K. Hanson. IceCube DOM Users’ Guide. 2003. IceCube Internal Documentation, docushare document 4071.
- [41] J. Kelly and J. Braun. *DOM-Cal Users’ Guide*. IceCube Internal Documentation.

- [42] J. Kelly. DOM-Cal. Presentation, Aug 2004. docushare document 9127.
- [43] D. Chirkin. Feature Extraction of IceCube Waveforms. IceCube internal Documentation.
- [44] I. M. Sobol. *USSR Comput. Math. Math. Phys.*, 7:86–112, 1967.
- [45] P. B. Price and K. Woschnagg. *Astroparticle Physics*, 15(97), 2001.
- [46] J. Ahrens et al. The AMANDA Collaboration. *Nuclear Instruments and Methods*, A524:169–194, May 2004.
- [47] C.H. Wiebusch. In *Proceedings of Workshop on the Simulation and Analysis Methods for Large Neutrino Telescopes*, pages 302–316, 1999.
- [48] <http://root.cern.ch>.
- [49] D. Heck et al. Hadronic Interaction Models and the Air Shower Simulation Program CORSIKA. *Proceedings of the 27th International Cosmic Ray Conference*, pages 233–236, 2001.
- [50] N.N. Kalmykov et al. *Nuclear Physics B (Proc Suppl.)*, 52B:17–28, 1997.
- [51] H. Fesefeldt. *RWTH, Aachen*, (PITHA-85/02), 1985.
- [52] D. Chirkin. *hep-ph/0407078*, 2004.
- [53] A. Bouchta et al. Seasonal Variation of the Muon Flux Seen by AMANDA. International Cosmic Ray Conference, 1999.