

OBSERVATION OF ATMOSPHERIC MUON NEUTRINOS WITH AMANDA

by

TYCE ROBERT DEYOUNG

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

DOCTOR OF PHILOSOPHY
(PHYSICS)

at the

UNIVERSITY OF WISCONSIN – MADISON

2001

© Copyright by Tyce Robert DeYoung 2001

All Rights Reserved

OBSERVATION OF ATMOSPHERIC MUON NEUTRINOS WITH AMANDA

Tyce Robert DeYoung

Under the supervision of Professor Francis Halzen

At the University of Wisconsin — Madison

The Antarctic Muon and Neutrino Detector Array (AMANDA) is designed to detect high energy neutrinos using the three kilometer thick ice cap covering the South Pole as a target and Cherenkov medium. Neutrinos that undergo charged current interactions with nucleons in the ice will produce ultrarelativistic charged leptons, which are detected through their Cherenkov and stochastic radiation by a three dimensional array of phototubes embedded in the ice cap at depths of 1500 to 2000 meters.

The background to the observation of neutrinos is the flux of penetrating muons produced in cosmic ray showers in the atmosphere. This flux is approximately one million times the neutrino flux. To reject this background, we look downward, using the Earth to filter out all particles except neutrinos. To demonstrate the correct operation of the detector, we observe atmospheric neutrinos, which are produced in cosmic ray showers in the Northern Hemisphere. The flux, energy spectrum, and angular distribution of these neutrinos are relatively well known, making them a convenient calibration source.

This work describes algorithms that have been developed to reconstruct and identify upgoing neutrinos in data recorded during the austral winter of 1997. A total of 204 neutrino candidates are identified, containing less than 10% background from misreconstructed downgoing muons. The neutrinos observed are found to agree with theoretical predictions of the atmospheric flux within the estimated systematic uncertainties. Limits are placed on high energy neutrino emission from known astronomical sources of very high energy gamma rays.

Francis Halzen (Adviser)

In Memoriam

G. Paul DeYoung

Thanks, Dad.

Acknowledgements

This document was made possible by years of support and assistance by many people, and it is a great pleasure to acknowledge their contributions. First, I wish to thank my advisor, Francis Halzen, who has supported me through my years at Wisconsin and who has taught me many things, not least of which is the importance of style. I would also like to thank my undergraduate advisor, Charles Cunningham, who first got me hooked on physics.

I must also thank the members of the AMANDA collaboration, past and present, without whom this work would obviously not have been possible. In particular I am grateful to Alex Biron, Jan Conrad, Kael Hanson, John Jacobsen, Marek Kowalski, Jodi Lamoureux, Matthias Leuthold, Ped Mioćinović, ‘Ignitio’ Taboada, and Christopher Wiebusch for discussions and technical assistance.

Closer to home, I have benefitted greatly from my cohorts in the Wisconsin AMANDA group, whom I thank not only for their help but also for their patience. I am especially grateful to Albrecht Karle, who has given invaluable advice and guidance relating both to physics and to physicists, and without whom it would have imploded long ago. I would also like to thank Katherine Rawlins and Rellen Hardtke for their advice, support, and friendship, and Bob Morse, who keeps the whole thing rolling. Last but definitely not least, this dissertation is largely the result of three years of incredibly fruitful, and usually cordial, discussions with Gary Hill, and I gratefully acknowledge his contributions.

I owe a great deal to my parents, Paul and Nancy DeYoung, who taught me that I could do whatever I put my mind to, and also the proper use of the apostrophe. I wish you both could have seen this finished.

And finally, I want to thank Rebecca, for putting up with it all with good grace.

Contents

1	Introduction	1
2	High Energy Neutrino Physics and Astrophysics	4
2.1	Astronomical Sources of High Energy Neutrinos	4
2.1.1	High Energy Cosmic Rays	5
2.1.2	Active Galactic Nuclei	6
2.1.3	Gamma Ray Bursts	9
2.1.4	Diffuse High Energy Neutrinos	11
2.2	Exotic Neutrino Physics	11
2.3	Atmospheric Neutrinos	12
3	Neutrino Detection	19
3.1	Optical Cherenkov Detectors	19
3.2	Lepton Signatures	21
3.3	Muon Neutrino Cross Section	22
3.4	Muon Energy Loss	25
3.4.1	Cherenkov Radiation	25
3.4.2	Stochastic Energy Deposition	26
3.4.3	Muon Range and Detectability	28
3.5	Background	30
4	The AMANDA Detector	34

4.1	Optical Modules	34
4.2	Geometry	35
4.3	Data Acquisition	36
5	Event Reconstruction	41
5.1	Maximum Likelihood Reconstruction	42
5.2	Bayesian Reconstruction	44
5.3	Minimization	47
6	Data	49
6.1	Triggering and Live Time	49
6.2	Hit Cleaning	51
6.3	Filtering	52
6.4	Reconstruction	54
7	Atmospheric Neutrinos	59
7.1	Cuts	60
7.1.1	Reduced Likelihood	60
7.1.2	Number of Direct Hits	62
7.1.3	Sphericity	62
7.1.4	Track Length	65
7.1.5	Line Fit Zenith Angle	68
7.1.6	Smoothness	68
7.2	Systematic Uncertainties	71
7.3	Event Quality	74
7.4	Efficiencies and Effective Area	82
7.5	Background Estimation	85
8	Flux from Point Sources	91
8.1	Candidate Sources	91

8.2	Optimization and Background Estimation	93
8.3	Results	97
9	Conclusions	102
A	Combining Weighted Monte Carlos	109
B	Cross Talk	113
B.1	Channel Cuts	113
B.2	Hit Cuts	116
C	Reconstruction Chain	118
D	Event Quality Levels	127
E	Events	130
F	On Reconstruction	135
F.1	Calculation of the Likelihood	135
F.1.1	The Math	136
F.1.2	Parametrization	137
F.1.3	Multiple Photoelectrons	138
F.1.4	Hit Probabilities	139
F.2	Expanding the Likelihood Space	140
F.3	Minimization	142

Chapter 1

Introduction

Astronomy is among the oldest of the sciences. Humanity has always been fascinated by the heavens, and over the millennia has invested vast effort in attempting to better understand the nature of the cosmos. From the beginning of history until the last century, however, only one means of observing the heavens was available: photons at optical wavelengths.

During the twentieth century, the number of ways of looking at the universe increased rapidly. Photon astronomy expanded to new wavelengths, and we now look at the sky in every band from radio through microwaves and the infrared to x-rays and gamma rays. As we developed these new ways of seeing, we discovered new objects and new phenomena whose existence was not even suspected: quasars, active galaxies, gamma ray bursts, and much more. These discoveries revolutionized our understanding of the universe.

But photons cannot tell us everything about the universe. As Fig. 1.1 illustrates, the sky is opaque to high energy gamma rays — above a few TeV gamma rays will interact with the cosmic microwave background, producing electron-positron pairs. The distance a gamma ray can travel through the universe falls quickly, until at PeV energies a gamma ray has a mean free path of only 10 kpc [1], less than the distance from earth to the galactic center. If we wish to know what the high energy sky looks like, we must find another way to see.

In the past decades there have been attempts to use neutrinos to probe these high energies. Suggestions that the ocean would be a suitable site for a large neutrino detector date to the 1960's [3, 4, 5, 6], and efforts to actually build such a telescope started with the DUMAND project in 1975 [7]. At the present time, there are two operating neutrino telescopes, Baikal [8] and AMANDA,




Figure 1.1: The gamma-ray horizon as a function of energy. The shaded region shows the areas accessible to observation only with neutrinos; the dominant processes and backgrounds responsible for attenuation are listed on the right. From [2].

and three others which are proposed or in the development and prototyping stages [9, 10, 11]. Beyond their capabilities as telescopes, these large neutrino detectors have many other scientific capabilities, and as they enter operation they promise to provide many interesting new results.

However, these are fundamentally novel instruments. Before their scientific potential can be realized, their behavior must be understood. There are no well-known astronomical sources on which a neutrino telescope can be calibrated; there is not even a body of collective experience or communal wisdom on how to operate such a device. Establishing a neutrino telescope as a working instrument is thus a major undertaking, even after the detector is physically assembled — one must learn how to analyze the data that is collected.

Chapter 2

High Energy Neutrino Physics and Astrophysics

2.1 Astronomical Sources of High Energy Neutrinos

The universe has been explored throughout the spectral range of electromagnetic radiation, from radio waves to gamma rays, but it has only recently become possible to use other messenger particles to observe the universe. The nature, number, luminosity, and spectra of neutrino sources are therefore matters of speculation. It is natural to use the observed energy in high energy cosmic rays and known sources of non-thermal, high energy gamma radiation for guidance as to the possible sources and magnitudes of astronomical neutrino signals. History has shown, however, that the opening of each new astronomical window has led to unexpected discoveries; there could well be hidden particle accelerators from which only the neutrinos escape [3]. In the following discussion of possible neutrino sources, therefore, one should not forget the possibility of serendipitous discovery.

High energy gamma rays may be produced in astrophysical sources either through radiative processes by accelerated electrons (synchrotron emission, inverse Compton scattering, and so forth) or through the decay of neutral pions



Neutrinos, by contrast, must be produced through hadronic processes, primarily pion and kaon decay:

$$\begin{aligned}
 p + X &\rightarrow \pi^\pm + Y \\
 &\quad \hookrightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\
 &\quad \quad \hookrightarrow e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu) \\
 \\
 p + X &\rightarrow K^\pm + Y \\
 &\quad \hookrightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\
 &\quad \quad \hookrightarrow e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu)
 \end{aligned} \tag{2.2}$$

$$\begin{aligned}
 p + X &\rightarrow K_L^0 + Y \\
 &\quad \hookrightarrow \pi^\pm + \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\
 &\quad \hookrightarrow \pi^\pm + e^\pm + \nu_e(\bar{\nu}_e)
 \end{aligned}$$

where depending on the environment of the source the target X may be a photon or another hadron, with Y varying accordingly. High energy neutrino astronomy thus has the potential to discriminate between hadronic and electronic models of gamma emitters such as supernova remnants (SNRs), gamma ray burst sources (GRBs) and active galactic nuclei (AGN).

2.1.1 High Energy Cosmic Rays

The origin of cosmic rays is one of the oldest puzzles in science. The prevalent theory is that most cosmic rays, at least those with energies up to perhaps 100 to 1000 TeV, are accelerated in supernova blast waves [12]. The argument is based largely on the circumstantial evidence that the power available from supernova explosions is about right and that strong shock waves naturally produce a spectrum consistent with what is observed, after accounting for effects of propagation. Confirmation of this theory could come by observing evidence of pion production at the correct level in the gas surrounding supernova remnants. Photons with energies up to several GeV have been detected by the EGRET detector [13], but there are only upper limits on most of these objects in

the TeV range and higher, which has raised uncertainties about the standard picture of the origin of galactic cosmic rays [14, 15]. TeV emission is now established from a few remnants, but it is thought to originate from radiation by high energy electrons [16]. Detection of TeV neutrinos from these sources would confirm their role as accelerators of hadronic cosmic rays as well.

Whether supernovæ produce these TeV cosmic rays or not, it is generally accepted that they are incapable of acceleration to higher energies. But somewhere in the universe, Nature accelerates particles to the astonishing energy of 10^{20} eV and even higher [17, 18]. Although there are plausible models for the origin of these particles in the halo of our own galaxy [19, 20, for example], the predominant opinion is that cosmic rays with energies greater than about $3 \cdot 10^{18}$ eV come from extragalactic cosmic accelerators. The production of the very high energy cosmic rays is thus an open question; some candidate sources are discussed in the following sections. However, whatever the mechanisms and sites of acceleration, some fraction of cosmic rays will interact in their sources to produce pions. These interactions may be hadronic collisions with ambient gas or photoproduction in the intense photon fields near the sources. In either case, the neutral pions decay to photons while charged pions include neutrinos among their decay products, with spectra related to those observed in the gamma rays. A neutrino telescope of sufficient sensitivity should be able to observe these sources and give clues to their nature.

2.1.2 Active Galactic Nuclei

One possible class of high energy neutrino emitters is active galactic nuclei, which are among the brightest gamma ray sources in the universe. AGN emit as much energy as entire bright galaxies, but they are extremely compact; within time periods as short as hours, their luminosities have been observed to flare by more than an order of magnitude [16, 21, and references therein]. AGN emit at all wavelengths of the electromagnetic spectrum, from radio waves to TeV gamma rays, largely through the interactions of accelerated electrons with magnetic fields and ambient photons in the source. The standard AGN model involves accretion onto a supermassive black hole, with high energy emission produced by acceleration in jets beamed perpendicular to the accretion disc [22], as shown in Fig. 2.1.

It is assumed that particles are accelerated by Fermi shocks in clumps of matter travelling




Figure 2.1: A model of high energy emission from an AGN. Electrons and possibly protons, accelerated in shockwaves in the jet, interact with photons radiated from the accretion disk or produced by electrons in the magnetic field of the jet. From [22].

along the jet with a bulk Lorentz factor $\Gamma \sim 10$, possibly larger, relative to the observer. In order to accommodate bursts lasting a day or less in the observer’s frame, the size of the clump in its rest frame must be less than $R' = \Gamma c \Delta t \simeq 10^{-2}$ pc. The clumps may in fact be more like sheets, extending across the jet’s size of roughly one parsec. The observed radiation at all wavelengths is produced by the interaction of the accelerated particles in the clump with the ambient radiation in the AGN. From the photon luminosity L_γ received over a time Δt , the energy density of photons in the rest frame of the clump can be inferred:

$$\rho_{E_\gamma} = \frac{L_\gamma \Delta t}{\Gamma \frac{1}{3} \pi (R')^3} \sim \frac{L_\gamma \Delta t}{\Gamma} \frac{1}{(\Gamma c \Delta t)^3} \sim \frac{L_\gamma}{\Gamma^4 \Delta t^2}. \quad (2.3)$$

With high luminosities L_γ emitted over short Δt , high energy photons will be contained by the high photon density through $\gamma\gamma \rightarrow e^+e^-$; the clump will be opaque to multi-TeV photons unless Γ is very large. A boost factor $\Gamma \gtrsim 10$ is required to dilute the clump to the point that 10 TeV gamma rays fall below the pair-production threshold [22].

If protons are accelerated along with electrons to energies of PeV to EeV, they will produce gamma rays via pion photoproduction, as in Eq. 2.1. Near the central black hole, the ultraviolet thermal background provides the target photons; in the jets non-thermal photons may also act as targets. If this is the case, then neutrinos should be produced in similar numbers following Eq. 2.2, and any protons that escape without interacting would contribute to the flux of high energy cosmic rays. The relative merits of electron and proton acceleration models are a topic of debate, but observation of high energy neutrino emission would settle the issue.

Examples of transparent sources with large boost factors are the nearby blazars¹, Markarian 421 and 501. These will be relatively weak neutrino sources because one expects at most one neutrino per photon. A source with the same morphology, but $\Gamma \simeq 1$, would be opaque to high energy photons and protons. It would be a “hidden” source, with reduced or extinguished emission of high energy particles, but undiminished neutrino production by protons on the high density photon target. Waxman and Bahcall have pointed out [23] that sources such as AGN which contribute to the observed ultra-high energy cosmic rays are limited to an energy flux $< 5 \cdot 10^{-8} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1} \text{ GeV}$ around 10^8 – 10^9 GeV. Hidden sources, in which high energy photons and hadrons are trapped but

¹A blazar is an AGN in which the jet illuminates the observer.

from which neutrinos escape, are of course not subject to this bound. Whether some AGN satisfy the conditions to be hidden is at present an open question.

2.1.3 Gamma Ray Bursts

Gamma-ray bursts (GRBs) are short, intense, and randomly distributed eruptions of high energy photons. The most popular mechanism for achieving these high energies is the conversion to radiation of the kinetic energy of ultrarelativistic electrons and protons that have been accelerated in a relativistically expanding fireball [24]. In the fireball's early stages, radiation is trapped by the very large optical depth, and the fireball cannot emit photons efficiently; cf. Eq. 2.3. The fireball's kinetic energy is therefore dissipated until it becomes optically thin, at which time the gamma rays produced by synchrotron radiation, and perhaps also by inverse Compton scattering off the accelerated electrons, can escape.

Neutrinos will also be produced when accelerated protons interact with the intense radiation field of the burst, following Eq. 2.2. The neutrino flux can be calculated as a function of the relative ratio of protons and electrons in the fireball. If it is assumed, for example, that GRBs are the source of the observed flux of the highest energy cosmic rays [25, 26, 27], then energy must be approximately equally transferred to electrons and protons in the fireball [28]. Based on the observed gamma spectra, one might suppose GRB neutrinos to be generated following a broken power law energy spectrum [29]:

$$\frac{d\phi}{dE_\nu} = \begin{cases} \frac{A}{E_B E_\nu} & \text{for } E_\nu < E_B \\ \frac{A}{E_\nu^2} & \text{for } E_\nu > E_B \end{cases} \quad (2.4)$$

where E_B is the energy of the break in the typical GRB two-power-law spectrum. Its value depends on the boost factor of the fireball, among other things. The normalization constant A can be determined, for example, from the assumption that GRBs are the source of the highest energy cosmic rays. The opacity of the source depends strongly on the Lorentz factor of the outflow; cf. Equation 2.3. The boost factor has been only indirectly determined by GRB follow-up observations, and it may be expected to vary somewhat from burst to burst, but is probably in the range 10^2 – 10^3 [30, 31].

The expected neutrino event rate in AMANDA-B10 can be determined by a Monte Carlo simulation of the GRB signal and the detector. The number of events triggering AMANDA-B10

GRB spectrum at trigger level - $\Gamma=300$




Figure 2.2: GRB spectrum at AMANDA trigger level. The distribution is a convolution of the neutrino flux and the probability of conversion to a muon within the range of the detector.

under these assumptions is shown in Fig. 2.2; the total prediction for 78 bursts in 1997 is about 0.1 events, taking $\Gamma = 300$ for all bursts and E_B around 700 TeV [32]. However, an important consequence of fluctuations in Γ is that the signal is dominated by a few very bright bursts. Although we expect less than one neutrino event in AMANDA-B10 from a typical GRB, because the neutrino flux varies as $\phi_\nu \propto \Gamma^{-4}$ a single burst with favorable characteristics could produce multiple events in the detector [30]. Moreover, the spatial and temporal information provided by satellite detection greatly reduces the potential background and permits coincident searches with much larger effective area than for diffuse or point sources of neutrinos.

In addition, AMANDA-B10 may be able to constrain other theories of GRBs. Although the fireball model as presented above is the most widely accepted theory, there are other models that may also be able to explain the observed gamma ray fluxes. Some of these models predict very different neutrino fluxes, which may be observable or constrainable with AMANDA-B10 [33, 34, for example].

2.1.4 Diffuse High Energy Neutrinos

Whether or not individual sources of high energy neutrinos are sufficiently bright to be observed, a large neutrino detector will be sensitive to the diffuse flux of neutrinos from the background of unresolved sources [2]. Identifying this flux is of course more difficult than detecting point sources, because one cannot rely on directional discrimination. However, the diffuse astrophysical flux will have a significantly harder spectrum than the atmospheric foreground, and so the astrophysical diffuse flux can be measured as the high energy component of the isotropic flux.

2.2 Exotic Neutrino Physics

It is generally accepted that some 90% of the mass of the universe is composed of some unknown nonluminous material, creatively referred to as “dark matter” (or “dark energy”). One theory regarding the dark matter is that it is made up, at least in part, of weakly interacting massive particles (WIMPs), for which one candidate is the lightest supersymmetric particle, generally believed to be the neutralino. If all of these assumptions are correct, then WIMPs travelling through space should have some finite probability to undergo elastic neutral current scattering with nuclei and become trapped in gravitational wells. If this is so, then they will build up in massive objects until they reach equilibrium between collection and annihilation.

Neutralinos will annihilate into leptons, gluons, quarks, or gauge or Higgs bosons, depending on the neutralino mass. Most of these annihilation products will produce neutrinos by hadronization or decay. These WIMP-induced neutrinos would be distinguished by their origin in the center of the earth, the sun, or the galaxy; the center of the earth is attractive because the only background is the atmospheric neutrino flux, although the equilibrium population will be lower than the other sites due to the shallower gravitational well [35]. The greater proximity may be a mixed blessing: on the one hand the detector covers a greater solid angle, but on the other hand the neutralinos may have

sufficient thermal energy for the neutrinos to come from a relatively large solid angle themselves.

Neutrino telescopes are very complementary to other searches for neutralinos, in that the sensitivity of direct and collider searches to new particles decreases with the particle’s mass. Neutrino telescopes, on the other hand, are most sensitive to high mass neutralinos because the higher mass will give the annihilation products more energy, and as discussed in Section 3.4.3 the detectability of neutrinos rises with energy. Direct searches, whose sensitivity decreases with mass, are unlikely ever to cover the mass range above a few hundred GeV; nor will the Large Hadron Collider, which will operate below the threshold for producing neutralinos of high mass.

Neutrino telescopes can also search for ultra-high energy neutrino signatures of topological defects predicted by grand unified theories, and for magnetic monopoles. AMANDA has already established the current best limit on relativistic monopoles [36], one order of magnitude below the Parker bound. Further, neutrino telescopes with sensitivity to astronomical point sources will have a unique mode of observation of neutrino oscillations. Astrophysical sources are expected to produce only electron- and muon-flavored neutrinos. However, propagating over cosmological distances, the neutrinos should reach the equilibrium population of equal numbers of each flavor if oscillations are allowed, regardless of the mixing parameters. The observation of tau neutrinos from astrophysical sources would constitute very clear evidence for neutrino oscillations. Such ultra-long baseline ν_τ appearance studies could probe the neutrino mass differences down to the level of $\Delta m^2 \gtrsim 10^{-17} \text{ eV}^2$ [29].

2.3 Atmospheric Neutrinos

The earth’s atmosphere is constantly bombarded with cosmic rays, which interact with nuclei in the air to produce extensive air showers. At the energies of relevance to neutrino telescopes, the cosmic rays are composed mainly of protons and helium nuclei, shown in Fig. 2.3, with some contributions from heavier nuclei [37]. In this energy range the cosmic rays follow a spectrum of approximately $E^{-2.7}$, depending slightly on species. The air showers precipitated by cosmic rays contain two types of particles that can reach a deep neutrino telescope: muons and neutrinos. These particles are referred to as “atmospheric” because of their origin. The high energy neutrinos are primarily muon neutrinos; electron neutrinos are suppressed by almost two orders of magnitude [38].

The angular distribution of the cosmic rays is approximately isotropic, but the physics of

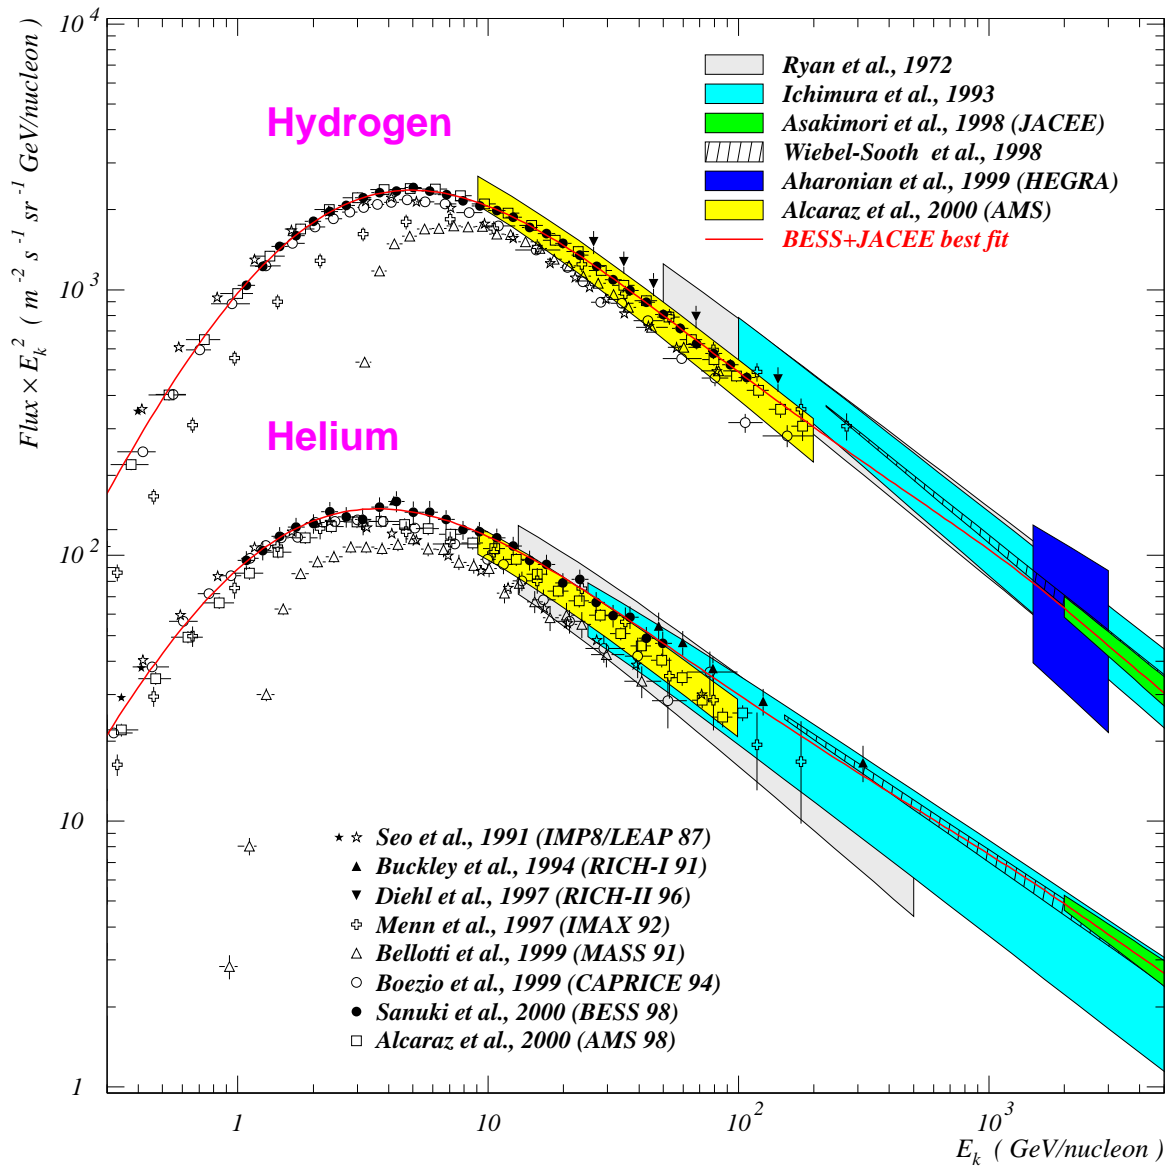


Figure 2.3: Differential spectra of the proton and helium components of the cosmic rays. The energy relevant to atmospheric muon and neutrino production in a large neutrino telescope is the TeV scale, at the far right of the plot. From [39].

meson decay modify the angular distribution. At high energies, the muon progenitors of atmospheric neutrinos are so strongly boosted that the vertical path length through the atmosphere is shorter than the muon decay length $x_0 = \gamma c \tau_\mu$. Once the muon hits the surface, energy loss becomes so rapid that the chance of decay to a high energy neutrino is very low. The high energy component of the atmospheric neutrino flux is thus generated primarily at high zenith angles relative to the surface under the air shower. This effect causes the horizontal flux at the detector to be several times as large as the vertical flux at energies relevant to a neutrino telescope [38], as seen in Fig. 3.8.

Atmospheric muons and neutrinos are formed primarily through the reactions described in Eq. 2.2. However, charmed particles are also formed in these high energy reactions, primarily D mesons. These charmed particles can decay semileptonically, producing neutrinos [38]. Because of the short lifetime of the charmed particles, these are referred to as “prompt” neutrinos. Prompt neutrinos constitute only a few percent of the neutrino flux at 1 TeV, but become more important at higher energies, as shown in Fig. 2.4. The charmed mesons responsible for the prompt flux do not reach the surface even at extremely high energies, so the prompt flux is isotropic, allowing it to be separated from the anisotropic conventional component.

Atmospheric neutrinos propagate some 10^4 km through the earth, depending on their zenith angle. Over this distance the neutrinos may oscillate between flavors; this is the presently accepted explanation for the atmospheric neutrino anomaly. The path length L traveled by a neutrino through the earth depends on the zenith angle θ , from 13000 km for a neutrino traveling the full diameter of the planet to only a few hundred km for a neutrino passing horizontally through a detector two kilometers below the surface. Furthermore, oscillations are sinusoidal in L/E_ν rather than L , meaning that high energy neutrinos effectively see a shorter path length in which to oscillate than do lower energy neutrinos. The oscillation probability is thus a function both of the neutrino’s energy and of its angle, as shown in Fig. 2.5, and becomes quite small for neutrinos with energies above 100 GeV (although a high value of Δm^2 would make oscillations more noticeable at higher energies). Measuring oscillation parameters with a neutrino telescope will thus require two things: a low energy threshold and good energy resolution. Smaller, denser detectors optimized for lower energies, such as SNO and Super-K, will be able to make more precise measurements on a per-event basis, but the much

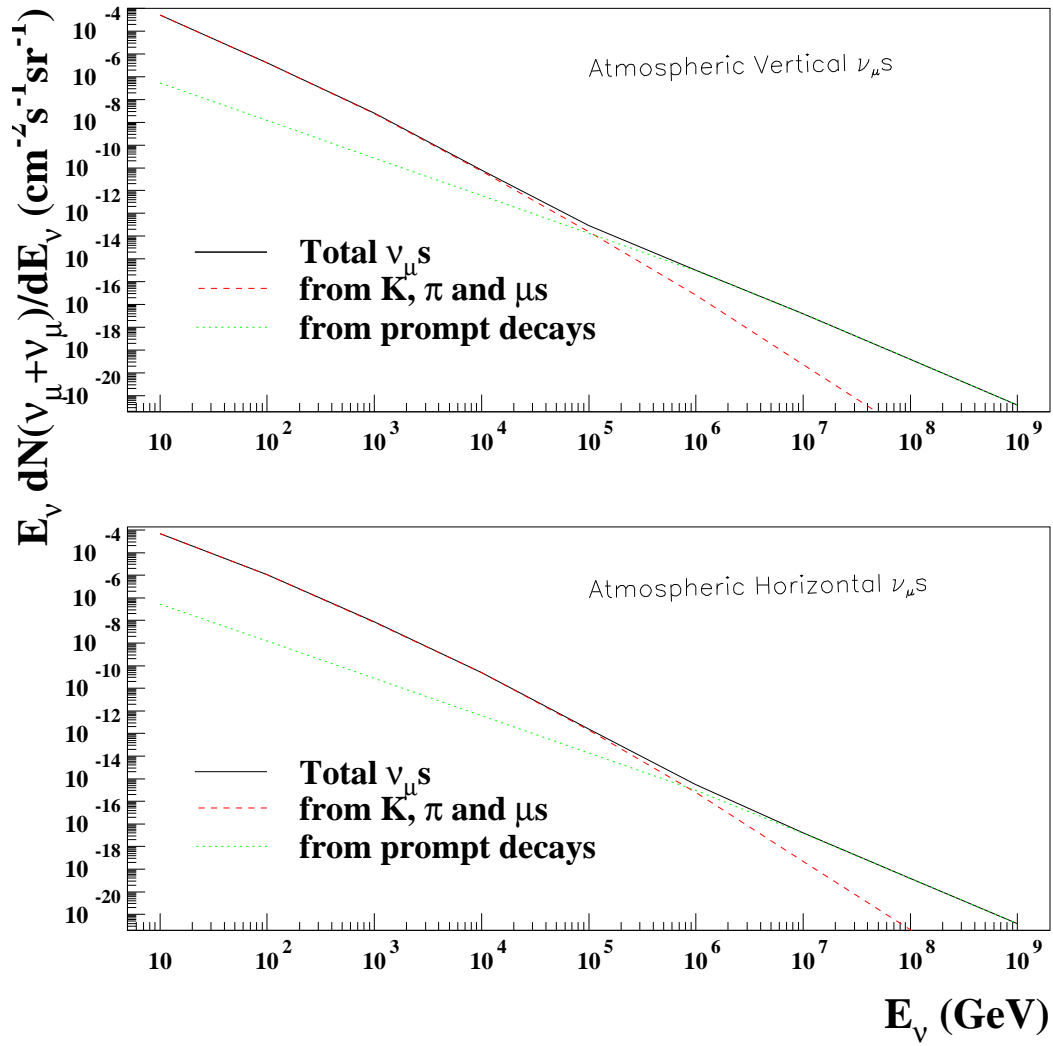


Figure 2.4: Vertical and horizontal muon neutrino fluxes. The prompt component becomes important at higher energies, especially in the vertical direction. From [40].




Figure 2.5: The probability that a neutrino is detected in the same flavor state that it was produced in, as a function of angle in a detector at 2 km depth. Maximal mixing between two states with $\Delta m^2 = 3.5 \cdot 10^{-3} \text{ eV}^2$ is assumed. From [40]; note that the angle is defined oppositely to the usage in this work.

larger effective volume of a detector like AMANDA may offset this advantage somewhat; nevertheless, measuring the effect will be difficult. For the high energy component of the atmospheric neutrino flux, for which neutrino telescopes are optimized, neutrino oscillations will have little or no effect.

The real importance of atmospheric neutrinos for a telescope like AMANDA lies in the fact that they are relatively well understood. At energies of tens of GeV to a few TeV their rate is known to within about 30% [41]. There are no benchmark astrophysical sources of high energy neutrinos, comparable for instance to the Crab for gamma ray telescopes, so the atmospheric neutrinos form an all-important calibration flux. As shown in Fig. 2.6, atmospheric neutrinos dominate all

expected astrophysical sources below TeV energies, and so to first order all neutrinos observed by a neutrino telescope with a threshold of order 100 GeV will be atmospheric in origin. For a mature neutrino telescope, atmospheric neutrinos will eventually constitute an irreducible low energy foreground to astrophysical neutrinos, and so telescopes like AMANDA (as opposed to detectors focused on lower energies like Super-K) are optimized for TeV neutrinos. In the meantime, the measurement of the expected number, angular distribution, and energy spectrum is thus a critical step in demonstrating the correct operation of a neutrino telescope; although the detector is optimized for somewhat higher energies, the large flux of atmospheric neutrinos means that an analysis specifically optimized to find them, as opposed to rejecting them as a foreground, will detect enough neutrinos to ensure that the detector's behavior is well understood.

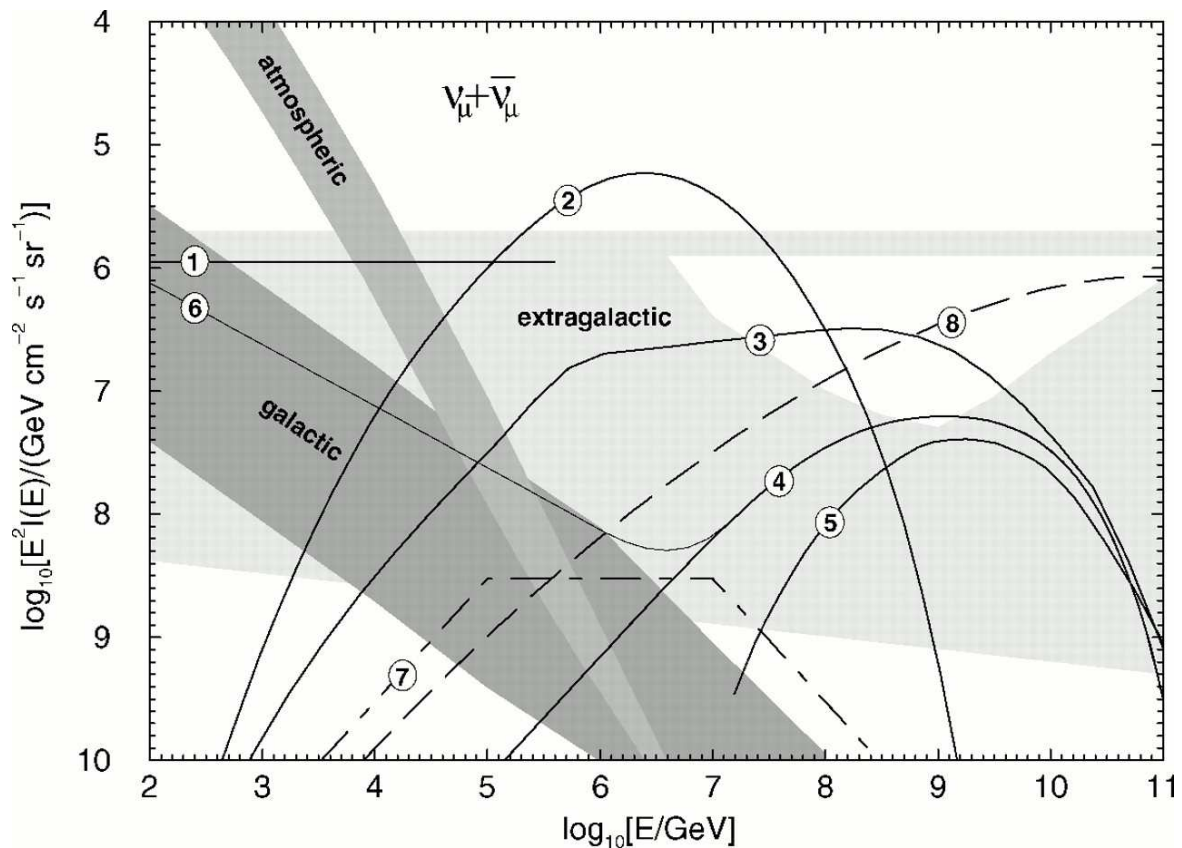


Figure 2.6: The magnitude of diffuse fluxes of muon neutrinos, from [2]. The fluxes are multiplied by E^2 for clarity. The width of the atmospheric band shows the angular dependence of the flux — the horizontal flux is the top boundary and the vertical flux is the bottom. The numbered lines indicate the combined fluxes from unresolved point sources under various models; see [2] and references therein for details.

Chapter 3

Neutrino Detection

3.1 Optical Cherenkov Detectors

All existing neutrino telescopes are based on the same technique, optical Cherenkov detection¹. The neutrinos themselves, of course, are unobservable. However, a neutrino or antineutrino passing through matter has some probability to interact via charged-current scattering



where l is one of the lepton flavors and N is a nucleon in the target. At high energies, the leptons will carry about half of the neutrino's energy, meaning that from the kinematics of the reaction, the neutrino and the lepton will be collinear to a mean deviation of

$$\sqrt{\langle \theta_{\mu\nu}^2 \rangle} \approx \sqrt{m_p/E_\nu},\tag{3.2}$$

about 1.75 degrees for a 1 TeV neutrino [12]. The other half of the energy will be released in the hadronic cascade X , producing a bright but relatively localized flash of light.

The scientific potential of neutrino astronomy arises from the great penetrating power of neutrinos, which allows them to emerge from dense inner regions of energetic sources. The unfortunate corollary is that the expected neutrino detection rate is small, meaning that extremely large detectors are required to observe reasonable event rates. For most astrophysical sources, the requisite scale is

¹Some pioneering early experiments, and experiments designed for other purposes but used secondarily to detect high energy neutrinos, were based on other techniques [2]

on the order of 1 km^3 [42]. For the observation of atmospheric neutrinos, effective volumes as much as two orders of magnitude smaller will still yield significant numbers of events provided that the energy threshold is 100 GeV or less, but nevertheless the scale of the detector must be at least tens of meters.

There are two approaches to building such a neutrino detector. The first, typified by Super-Kamiokande [43] and the Sudbury Neutrino Observatory [44], is to build a densely instrumented detector on the scale of tens of meters, with efficiency near unity. This approach offers the advantage that the detector is (barely) small enough to be custom-built. Such a detector will be effective at low energies, but the falling spectra of high energy neutrino fluxes mean that the geometric size of the detector will be insufficient to collect many events at high energies. Furthermore, the detector is not large enough to track muons for distances sufficient for energy measurement, or to contain tau events. The effectiveness of such a detector is thus quite limited at energies above a few hundred GeV.

The second approach is “to build a detector that barely works” [45]. Rather than building an artificial detector, one instruments a naturally occurring medium. Since the price is determined by the number of instruments used rather than the size of the detector, the density of instrumentation is a totally free parameter of the design, and one can achieve kilometer-scale detectors at relatively low cost if one accepts a very sparse detector, with a typical spacing between photosensors of tens of meters. The threshold of such a detector will be quite high, of course, and the efficiency for events near threshold will be low, but the amount of light generated by very high energy events is such that the detector will retain reasonable efficiency for the high energy neutrinos typical of astrophysical sources.

Two suitable natural media have been identified for the optical detection of neutrino-induced leptons: deep oceans or lakes [5] and the austral polar ice cap [46]. Both of these media are available in bulk, of course, and both are extremely clear, with typical attenuation lengths of 25–50 m or more in the blue wavelengths most important for neutrino telescopes [47]. To build a neutrino detector, one embeds a three-dimensional lattice of photosensors in the medium. Time resolutions of a few nanoseconds (corresponding to geometric resolution of about a meter) allow the track of a high energy

lepton to be reconstructed from the Cherenkov and stochastic light it emits, and as discussed above the lepton will be nearly collinear with the incident neutrino. Depending on the medium and the details of the hardware, a muon track can be reconstructed with an accuracy of better than a degree [9, 48], sufficiently precise for astronomical measurements.

Ocean water and polar ice are optically comparable media. In ocean water, the absorption length is much shorter than the scattering length, meaning that photons are very rarely scattered in the medium between emission and detection. In ice, on the other hand, the attenuation is dominated by scattering. This means that ice detectors are in principle capable of making better calorimetric energy measurements, but on the other hand care must be taken in an ice detector to distinguish “direct” or unscattered photons, which are useful for track reconstruction, from scattered photons which are useful only for energy measurement. It can be expected that a mature ice detector will have better energy resolution but slightly poorer angular resolution than a water detector of similar size.

3.2 Lepton Signatures

For a muon neutrino or antineutrino, the secondary muon will travel in a roughly straight line, losing energy at the rate of 0.2 GeV/m (rising with E_μ), as discussed in Section 3.4. For high energy neutrinos, this implies a muon path length of hundreds of meters, kilometers, or even tens of kilometers. The experimental signature of a muon track will thus be a long, linear deposition of light, as shown in Fig. 3.1(a). The misalignment angle in Eq. 3.2 decreases with the fraction $(1 - y)$ of the neutrino’s energy carried by the muon, so that longer-range muons will in general be better aligned with the original neutrino than shorter-range muons. Since the probability to detect a muon depends on its range, this means that the average misalignment of *detected* muons will actually be less than given in Eq. 3.2 [9]. Because the long muon range increases the effective volume of the detector, and because the muons preserve directional information about the neutrinos, muon detection is the most straightforward mode of operation of a neutrino telescope. This work will therefore focus on a search for muon neutrinos.

Nevertheless, neutrino telescopes are sensitive to the other leptons, and flavor tagging is possible based on the characteristic signatures of the different leptons. Electrons produced according

to Eq. 3.1 will not have the long ranges of muons. Rather, they will quickly cascade in the target medium via bremsstrahlung and pair production, depositing their energy exponentially over a typical length scale of approximately 36 cm [49]. This cascade will produce a great deal of light in a region of a few meters' length. In an ice detector, the light will be strongly scattered, producing an effectively spherical pattern, as opposed to the linear trace of a muon.

The most interesting signature comes from tau leptons. The short lifetime of the tau will be relativistically dilated by its high energy, but will still be too short for the tau to travel large distances except at very high (PeV scale) energies. For moderate energies, then, the tau neutrino will interact, producing a hadronic cascade at the interaction vertex. The tau will travel some distance comparable to or shorter than the size of the cascade, whereupon it will decay. The decay will produce a second cascade which will be extremely difficult to resolve from that produced in the original neutrino interaction, so that the event will be essentially indistinguishable from an electron event. At very high energies, however, the tau may travel hundreds of meters. Energy loss along the track is suppressed by the tau mass, meaning that the track will be dimmer than a muon track, but the second cascade will be far enough separated from the interaction vertex to be clearly resolved, as illustrated in Fig. 3.1(b). This “double-bang” topology forms a very distinctive signature of contained tau events [50].

3.3 Muon Neutrino Cross Section

The inclusive cross section for charged-current muon neutrino-nucleon interactions is [51]


$$\frac{d^2\sigma}{dx dy} = \frac{2G_F^2 M_N E_\nu}{\pi} \left(\frac{M_W^2}{Q^2 + M_W^2} \right) [xq(x, Q^2) + x\bar{q}(x, Q^2)(1 - y^2)], \quad (3.3)$$

where $-Q^2$ is the invariant momentum transfer from the neutrino to the outgoing muon, q and \bar{q} are the quark and antiquark structure functions of the nucleon, G_F is the Fermi constant, M_N and M_W are the masses of the nucleon and W boson, and x and y are the Bjorken scaling variables


$$x = \frac{Q^2}{2M_N(E_\nu - E_\mu)}$$

and

$$y = 1 - \frac{E_\mu}{E_\nu},$$



(a) A muon event. The muon passes through the array emitting light at a relatively constant rate.



(b) A tau event. The two spheres of light are caused by the initial neutrino-nucleon interaction and the subsequent decay of the tau lepton.

Figure 3.1: Simulated leptonic events in Ice Cube. The larger scale of the proposed detector makes the different topologies clearly visible.




Figure 3.2: Neutrino and antineutrino cross sections as a function of energy, from [51]. The solid lines are based on the CTEQ3 parton distributions; the dashed and dot-dash lines are neutrino cross sections based on older measurements.

the fraction of the nucleon's four-momentum carried by the interacting quark and the fraction of the neutrino's energy deposited in the interaction, respectively. At low energies the antineutrino cross section is roughly a quarter that of neutrinos, but at high energies the cross section is dominated by interactions with sea quarks in the nucleon rather than valence quarks, and the cross sections become equal, as shown in Fig. 3.2.

At low neutrino energies, $-Q^2 \ll M_W^2$, so the term in parentheses in Eq. 3.3 may be neglected. The cross section is thus seen to rise linearly with the neutrino energy. $-Q^2$ becomes comparable to M_W^2 at about 3.6 TeV, causing the growth of the cross section to slow. However, the average y begins to fall, as shown in Fig. 3.3, meaning that the muons carry off a higher fraction of the neutrino energy [51]. This faster-than-linear rise in muon energy, and thus muon range, with neutrino energy partially offsets the slower growth in neutrino cross section in terms of the detectability of the




Figure 3.3: Average inelasticity of neutrino-nucleon interactions, from [51]. The falling $\langle y \rangle$ in the region from 1 TeV to 1 GeV compensates for the falling cross section by increasing the muon range.

neutrinos².

3.4 Muon Energy Loss

Relativistic muons travelling through matter lose energy through a variety of channels. Because a muon is detected by the energy it gives off, via the light emitted in the various loss processes, and because a muon's range is determined by its rate of energy loss, it is essential to understand these mechanisms of energy loss.

3.4.1 Cherenkov Radiation

Ice is a relatively dense optical medium, with an index of refraction³ $n \simeq 1.33$ in the blue and green wavelengths at which the optical modules are most sensitive [53]. Highly relativistic muons, then, will have a velocity well over the Cherenkov threshold of c/n ; the energy threshold for muons

²For antineutrinos the fall in y is less drastic. However, the turnover in the cross section is lower, so the combined effect is similar to that for neutrinos.

³This is the index based on the *phase* velocity of light. Properly it is the group velocity that is important in most cases, but it is shown in [52] that the differences are relatively minor.

in ice is [54]

$$E_{\text{Ch}} = \frac{m_{\mu}}{\sqrt{1 - \frac{1}{n^2}}} = 160 \text{ MeV}.$$

The Cherenkov radiation is strongly peaked at an angle of

$$\cos \theta_{\text{Ch}} = \frac{1}{\beta n} = 41^{\circ}$$

with respect to the muon direction for energies relevant to AMANDA.

The total continuous energy loss of a relativistic muon in ice is about 0.2 GeV/m [55]. The energy emitted in Cherenkov radiation is only a small part of the continuous loss, about 2 MeV/m, the rest going into ionization of the medium. Nevertheless, a muon emits a few hundred Cherenkov photons per cm, enough for detection. The Cherenkov radiation follows a $1/\lambda^2$ spectrum, meaning that bluer wavelengths are the most important for detection of the light, up to the ultraviolet cutoff imposed by the glass components in the optical modules [54].

3.4.2 Stochastic Energy Deposition

In addition to the Cherenkov light radiated by all relativistic muons, very high energy muons will undergo stochastic (or ‘catastrophic,’ or ‘discrete’) energy loss. The main stochastic processes are bremsstrahlung, direct e^+e^- pair production, and (somewhat less important) hadronization of nuclei [56], shown in Fig. 3.4. As the name suggests, these processes are relatively rare but involve large amounts of energy, producing sharp bursts of light at discrete points along the track⁴. The average rate of stochastic energy loss is nearly proportional to the muon energy, so the total rate of energy loss per unit length travelled can be parametrized by

$$-\frac{dE}{dx} = a(E) + E b(E) \tag{3.4}$$

where both $a(E)$ and $b(E)$ are approximately constant at the energies of interest [55]. As discussed above, $a \approx 0.2$ GeV/m. The value of b is about $3.4 \cdot 10^{-4} \text{ m}^{-1}$ in ice, so that catastrophic events are the main component of energy loss for muons above about 600 GeV [54].

⁴At very high energies, these processes become so frequent that they cannot be resolved, resulting in quasi-continuous emission of light. Even at these energies, however, extremely bright single events — particularly bremsstrahlung — will be distinguishable from the mass of lower energy losses, producing the same stochastic pattern of light emission.




Figure 3.4: The linear coefficient of muon energy loss, showing the contributions from pair production, bremsstrahlung, and nuclear hadronization. β in the figure corresponds to b in this work. The values given are for rock; ice is less dense and so dE/dx is lower. From [56].

The center of mass frame for these stochastic processes is strongly boosted along the muon direction due to the extremely high momentum of the muon. The particles produced in these processes are thus kinematically constrained to the same direction as the muon, and the Cherenkov light they emit is also peaked around the Cherenkov angle of the muon. There is a smearing of a few degrees around the Cherenkov cone [57], which must be accounted for in simulation and reconstruction, but this is a higher-order correction to the general picture of conical light emission, with stochastic events varying the intensity, but not the direction, of emission.

3.4.3 Muon Range and Detectability

Equation 3.4 can be solved to find the approximate range of a muon of initial energy E_0

$$R_\mu(E_0) \approx \frac{1}{b} \ln \left(\frac{bE_0}{a} + 1 \right) \quad (3.5)$$

assuming of course that the muon does not decay in flight. For low energy muons, this provides a fairly good estimate, but for high energy muons the energy loss is dominated by catastrophic loss, and the random nature of the process becomes important, as shown in Fig. 3.5. A full Monte Carlo simulation is thus necessary to take into account fluctuations in the loss rate. On average, however, the muon range rises linearly with energy up to nearly a TeV, after which the growth of range is only logarithmic.

A muon can be detected at any point along its track, of course⁵. We can therefore use the product of the neutrino’s cross section and the muon’s range as a rough index of the relative “detectability” of a muon neutrino. The longer range of high energy muons, in conjunction with the rising cross sections discussed in Section 3.3 above, will offset the falling energy spectra of neutrinos. The combination of these effects causes the detectability of muon neutrinos to rise as E^2 (at least up to energies of a few TeV) meaning that for a hypothetical astrophysical E^{-2} neutrino source the number of neutrinos detected as a function of (neutrino) energy will be roughly constant in this energy region. Atmospheric neutrinos, on the other hand, have a softer spectrum, and so the number of muons detected from atmospheric neutrinos decreases rapidly with energy, as shown in Fig. 3.6.

⁵Ionization losses are nearly constant down to muon energies of about 1 GeV, so the picture described above describes all but the last few meters of a muon’s range. An energy threshold below several hundred GeV is nothing more than the requirement that the muon not range out before being tracked for a minimum distance within the detector.




Figure 3.5: Distributions of muon range in standard rock as a function of initial muon energy (in GeV). The vertical lines indicate the analytic solutions to Eq. 3.5, taking into account the running of b shown in Fig. 3.4. At higher energies the muon range is systematically overestimated by the analytic formula. From [56].

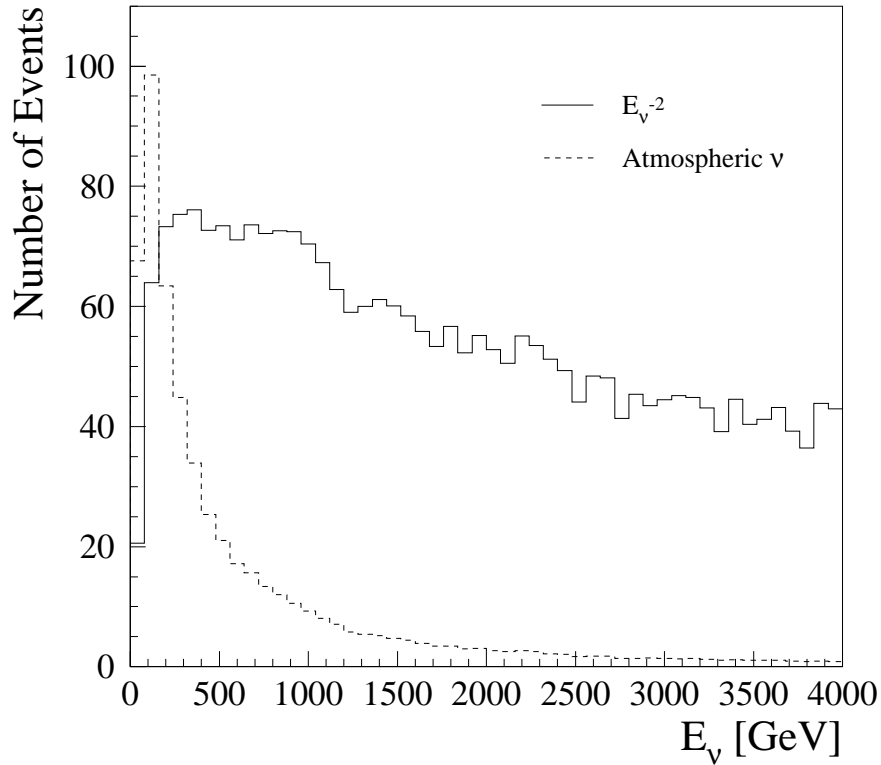


Figure 3.6: The energy of detected neutrinos from a hypothetical E_ν^{-2} source of arbitrary magnitude, as compared to those from atmospheric neutrinos. For the E_ν^{-2} source, the detection rate is approximately constant from the detector threshold to 1 TeV, whereas the rate of atmospheric neutrinos falls quickly.

High energy muons radiate energy rapidly, so even muons generated at very high energies traverse most of their range at TeV energies. Nonetheless, muons from sources with hard spectra will typically reach the detector at higher energies than those from atmospheric neutrinos, as shown in Fig. 3.7. Searches for atmospheric neutrinos must concentrate on low energies, whereas for sources with hard spectra one benefits from a high energy threshold that reduces the atmospheric neutrino background.

3.5 Background

The main background with which neutrino telescopes must contend is that of cosmic ray muons. Muons are produced when cosmic ray primaries (protons and heavier nuclei) impact on the earth's atmosphere, mainly through hadronization into pions and kaons which then decay to muons. This is

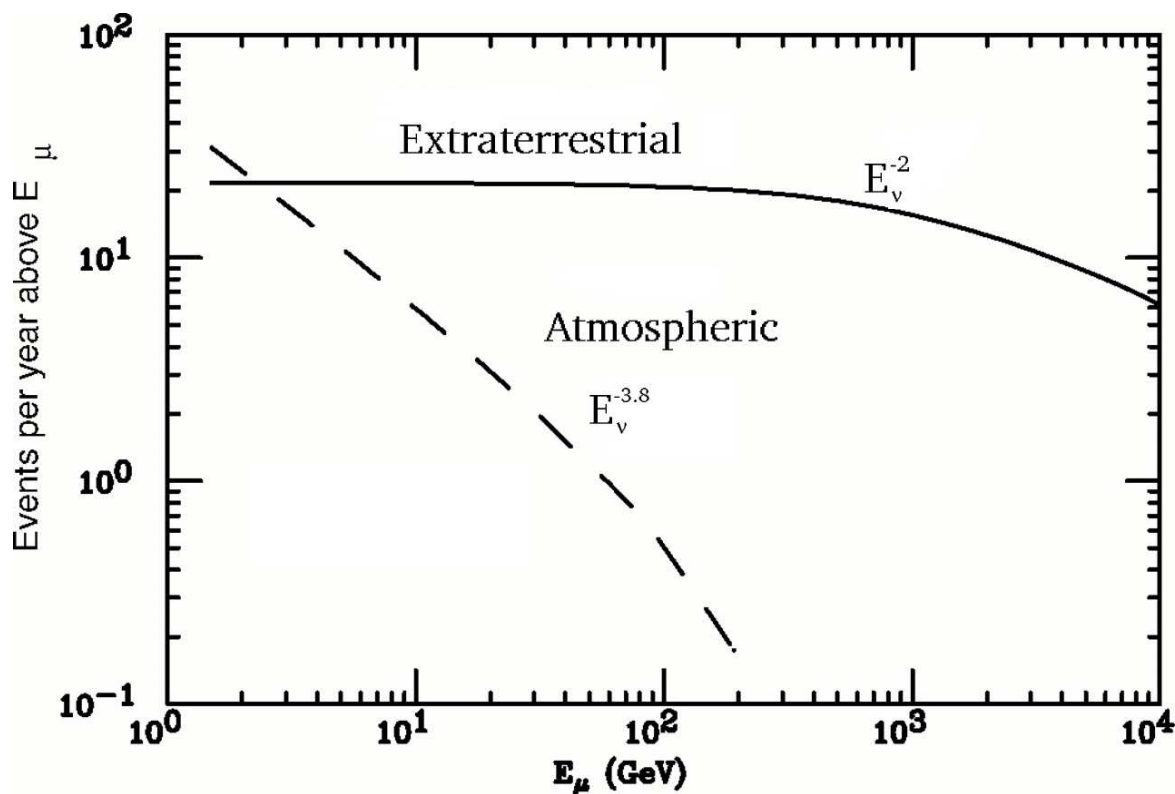


Figure 3.7: The (integral) muon energy spectrum for an E_ν^{-2} source of arbitrary magnitude, compared to muons from atmospheric neutrinos. Raising the threshold to 1 TeV has little effect on sensitivity to astrophysical sources, but greatly reduces the number of atmospheric neutrinos seen. For this reason, AMANDA is optimized for TeV energies. From [2].

the same process that produces the atmospheric neutrinos that are so useful as a calibration beam⁶. High energy muons, as noted above, have a range of several kilometers, and so muons produced in the atmosphere above a neutrino telescope have some probability to penetrate to the depth of the detector. For this reason neutrino telescopes are built as deep as possible, to attenuate the downgoing muon flux.

The vertical muon flux at the top of AMANDA-B10 (1500 m below the surface) is approximately $3 \cdot 10^{-3} \text{ m}^{-2} \text{ sr}^{-1} \text{ sec}^{-1}$, falling by a factor of 3 over the 500 m to the bottom of the detector [58]. These muons are responsible for the vast majority of AMANDA's triggers, which occurred at a true (deadtime-corrected) rate of 100 Hz with the 1997 set-up. This rate completely dominates any expected signals, making extensive background rejection necessary to obtain reasonable signal-to-noise. These muons are of course physically identical to comparable neutrino-induced muons, so in searching for relatively low energy atmospheric neutrino events we restrict our search to upgoing muons. The bulk of the earth then acts as a filter against directly produced muons, and the neutrino events become detectable, as shown in Fig. 3.8. At higher energies, above a few TeV, the falling spectrum of atmospheric muons has reduced the flux to levels low enough that astrophysical sources could be seen even in the Southern sky. Of course, the separation of signal and background is still only possible on a statistical basis for downgoing muons; one can never be absolutely certain that a given muon is not atmospheric, although the odds against such an origin may be vanishingly small. The downgoing hemisphere becomes even more important at very high energies, because the rising neutrino cross section will render the earth opaque even to neutrinos.

⁶These two methods of muon production lead to a somewhat confusing nomenclature. The term “atmospheric muon” refers only to a muon that is produced directly in the atmosphere and penetrates to the detector. Muons produced by neutrinos will always be referred to as neutrino-induced, even though the neutrinos themselves may be atmospheric in origin.




Figure 3.8: The angular distribution of the muon flux triggering AMANDA. The solid line is the atmospheric muon flux as predicted by the `corsika` air shower simulation [59]. The dashed line shows the rate of muons produced by atmospheric neutrinos, from the `nusim` simulation [60].

Chapter 4

The AMANDA Detector

4.1 Optical Modules

The AMANDA detector consists of a three-dimensional lattice of optical modules (OMs). Each OM consists of an 8" Hamamatsu photomultiplier tube (PMT) in a glass pressure housing¹. The OM is connected to the surface by an electrical cable, which provides high voltage to the module and also transmits the anode signal of the PMT back to data acquisition (DAQ) electronics on the surface.

Several types of OMs may be distinguished. Those on the inner four strings of the AMANDA detector, collectively referred to as AMANDA-B4, were deployed in the 1995–96 austral summer. These OMs were connected by coaxial cable to the surface, which provides protection against cross talk in the cables. Coaxial cable is highly dispersive, however, which results in a great deal of pulse distortion in the course of transmission (10 ns PMT pulses arrive at the surface with a width of more than 400 ns). Coax is also quite thick, which limits the number of cables that can be bundled together in a single string.

For these reasons, the six strings deployed in 1996–97 used twisted pair cables. Twisted pair cables produce less dispersion (pulses are typically stretched to 150–200 ns), and allow more cables per string. However, there is a great deal of cross talk observed in these strings. Although the site or sites of the cross talk are not clear at the time of this writing, it is believed based on the correspondance of cross talk patterns to the arrangement of twisted quads in the strings that at least

¹In AMANDA there is a one-to-one correspondance between OMs and PMTs, and the two terms will be used interchangeably in this work. In water detectors, bioluminescence and the decay of radioactive potassium produce noise rates orders of magnitude higher than in ice, and so OMs in such detectors generally comprise several PMTs operated in local coincidence to reduce the rate of false signals.

a portion of the cross talk occurs within the cable [61].

The strings deployed since 1997 have incorporated a number of other signal transmission technologies. All OMs on strings 11–19 have included an optical fiber for pulse transmission. The signal from the PMT is sent through an LED (or in some cases a laser diode) whose signal is transmitted over the fiber to an optical receiver at the surface. The fiber is essentially dispersion-free, meaning that pulses with separations of as little as 10–15 ns can be resolved. However, the optical fibers have proven somewhat more vulnerable to damage during refreezing, with a loss rate of nearly 10%. Those OMs whose fibers were broken are read out in the traditional manner, over twisted pair electrical cable.

Another technology which was investigated during the 1999–2000 season is based on digital optical modules (DOMs). These modules contain analog transient waveform digitizers (ATWDs), which record and digitize pulses *in situ* and transmit them to the surface asynchronously. This results in full retention of information, and obviates the need for expensive and vulnerable optical fibers, but means that the DAQ electronics are buried in the ice beyond possibility of repair or upgrade.

4.2 Geometry

AMANDA’s geometry is constrained by the need to deploy modules in strings, as opposed to some water detectors (notably NESTOR [10]), which plan to deploy three-dimensional structures. This is not a large disadvantage, though, because to a good approximation the effectiveness of a detector of given optical medium, number of OMs and intermodule spacing is independent of the precise arrangement of the modules. In fact, most of the water detectors have also chosen to deploy their modules on strings [9]. However, the string-based design means that the vertical spacing of modules is quite a bit closer than the horizontal spacing.

The first major deployment at the South Pole, in 1993–94, was of four strings at depths of 800–1000 m. The ice at these depths was found to contain a large residual population of air bubbles, which scatter light very strongly. Though excellent as a calorimeter [62], the ice at these depths cannot be used effectively for track reconstruction, and further deployments were made at depths of 1500–2000 m. The shallow detector, known as AMANDA-A, is operated in coincidence with the

deeper AMANDA-B, permitting the study of vertically downgoing atmospheric muons.

AMANDA-B is deployed in concentric circles of strings, one string in the approximate center of the array and three more on a circle of 35 m radius composing AMANDA-B4. An additional ring of three pairs of strings at 60 m radius makes up AMANDA-B10, the 1997 configuration on which this work is based. Eight more strings at 100 m radius and one just within this outer circle round out the full 19-string AMANDA-II. A schematic of the full detector is shown in Fig. 4.1.

The central four strings, based on coaxial cables, each contain twenty OMs, plus six test modules at the bottom of the center string which are not used in the data analysis. The vertical spacing between OMs on these strings is 20 m. On each of the six strings from 1997, 36 OMs were placed at 10 m intervals, bringing the total to 302 modules for AMANDA-B10. The sparsity of modules on the central strings means that the most precise information generally comes from the outer strings, although the central strings provide important proof that an event is a single continuous track and not caused by independent simultaneous muons.

The ice in which the AMANDA detector is embedded is very nearly uniform. However, climatological events in the planet's past, such as ice ages, have left their marks in the form of layers of ice with different amounts of inclusions of dust, soot, and so forth. These dust layers alter the optical properties of the ice, affecting photon propagation. The layers have been surveyed using calibration light sources [63] and downgoing muons [64]. The variation in effective scattering length λ_s with depth is shown in Fig. 4.2 for the depths of AMANDA-B10. These variations are included to first order in the current detector simulation; efforts to incorporate them more properly into the simulation and reconstruction programs are underway.

4.3 Data Acquisition

An event in AMANDA is triggered by the simultaneous observation of light by several OMs. For each event, the entire array is read out by electronics on the surface, and the information from the DAQ electronics is written to disk. Muon events have a typical duration of $\lesssim 5\mu\text{sec}$, although the propagation delays of PMT signals in the cables leading to the DAQ cause the signals from an upgoing track to arrive in a shorter interval.

In general, each OM in an event will produce a series of pulses corresponding to a series of




Figure 4.1: Schematic of the geometry of AMANDA-II. AMANDA-B10, on which this work is based, is shown in expanded view in the center, and an optical module is blown up on the right. The Eiffel Tower is shown to illustrate the scale.




Figure 4.2: Effective scattering length as a function of depth in AMANDA-B10, showing the horizontal dust layers running through the detector. On the right the depths of the OMs on the 10 strings are shown for reference. The straight lines in the left plot are the values used in the detector simulation. From [65].

photoelectrons produced at the photocathode. Because the pulses are heavily smeared, primarily by dispersion in the cables leading to the surface, individual pulses can be resolved only if they are widely separated — by more than about 500 ns for the central strings, or 200 ns for the outer six strings. Because so much information is lost in transmission, there was no attempt in AMANDA-B10 to record the complete waveforms from the OMs. Rather, for each channel the times at which the waveform crosses a discriminator threshold are recorded. The TDCs used have a buffer depth of eight pulses, each composed of a leading edge and a trailing edge. In normal operation most channels record no pulses, and even modules near the center of an event typically record no more than three, sometimes four, pulses². Each channel is also read out by a peak-sensing ADC, which records the maximum amplitude of the pulse train as a whole. Because the probability of two or more highly scattered photons arriving at a module simultaneously is negligible, the amplitude may safely³ be assigned to the first pulse in the series, and any pulses which are recorded after the Cherenkov time may be assumed to be due to single photoelectrons.

Triggers are formed based on multiplicity. Each channel, on detecting a pulse, sends a 2 μsec square pulse to the trigger unit. The inputs from all channels are added and compared to a preset threshold, which was set at 16 channels for 1997. When the sum crosses the threshold, a stop signal is sent to all TDCs and a veto of several μsec is sent to the trigger. All channels are then read out, and the system is reset. In fact, due to high-noise channels, the effective trigger multiplicity was somewhat less than the nominal level, and after hit cleaning the actual threshold is seen in Fig. 6.1 to be about 13 real hits.

The TDCs in AMANDA have a buffer length of 32 μsec . They are operated in common-stop mode, meaning that when they are read out they provide the history of the 32 μsec before the stop signal arrived. It takes approximately 9 μsec for the common stop to be sent after the trigger threshold is crossed, and events typically extend from 2 μsec before the trigger to 2500 ns afterward. The time history for an event thus consists of some 21 μsec of random noise, 2000 ns of the beginning of the event, a trigger hit at 23 μsec , another 2.5 μsec of event, and then some 7 μsec of afterpulsing

²There are channels which exhibit high noise rates, either occasionally or over the entire year, and these channels will frequently fill up the buffer. The information from these channels is not useful, however, and is ignored in the data analysis.

³Unless the first pulse is due to dark noise, which is very rare after hit cleaning.




Figure 4.3: Recorded times of leading edges in OM 140. Hits from muons fall in a window of a few μsec around the spike at 23 μsec . The spike is formed by events in which OM 140 contributed the triggering hit, and reflects the time delay in the triggering electronics. The low rates before the muon window are dark noise, and the excess above the dark noise level after the event window is due to afterpulsing.

and random noise, as shown in Fig. 4.3.

Chapter 5

Event Reconstruction

Reconstruction in AMANDA is in many respects an *ad hoc* process, which has evolved gradually from quite simple techniques designed for water neutrino telescopes as the need for additional complexity slowly became more apparent. This chapter will attempt to discuss the theory of reconstruction in a systematic fashion. It is hoped that the discussion here and in App. F will illuminate areas in which the reconstruction can be improved or extended in the future. The low efficiency of ongoing muon identification (see Chap. 7) means that in principle the effective volume of the detector could be increased by as much as an order of magnitude simply through improvements in the reconstruction algorithms. In terms of cost effectiveness this avenue seems by far the most promising way of increasing AMANDA's capability in the near future.

Mathematically, reconstruction is the process of choosing, for an observed event E , an hypothesis H to explain the event. Here $E = E(t_i, a_i, d_i, \dots)$ includes all of the information recorded about the event; in AMANDA-B, this information consists of the times t_i , amplitudes a_i , and durations (TOTs) d_i of all observed pulses (or "hits") i . H is in principle much more complex, containing all relevant information about numbers, paths, and energies of muons and other particles, positions of stochastic energy deposition, detector effects, photon propagation, etc.. In practice, of course, much of this information is not of interest, and in any case the AMANDA detector is incapable of making fine distinctions¹. Various approximations can thus be made which greatly simplify the range of hypotheses under consideration; reconstruction then consists of finding that point H_0 in the chosen

¹For instance, we can never know, and don't much care, whether the stochastic loss along some muon track was a 5 GeV bremsstrahlung photon or a 2 GeV brem with a 3 GeV pair production event 50 cm farther along the track.

hypothesis space which is preferred as an explanation for an observed event E_0 .

For the present analysis, we have used the simplest hypothesis space possible for muon tracks: the five-dimensional space consisting of three spatial coordinates (x, y, z) defining the position of the muon at some arbitrary time and two angles (θ, ϕ) defining the trajectory of the muon, which is assumed to be straight and of infinite length. Other dimensions to the hypothesis space can easily be envisioned, for example the muon's energy², stopping and/or starting points, the positions and energies of major stochastic energy deposition, the number of muons involved (in the context either of a downgoing muon bundle from a cosmic ray shower or of simultaneous independent muons), the flavors of particles, and so forth. However, incorporating these extra degrees of freedom will involve considerable effort, and at some point the amount of information delivered by the detector will be insufficient to constrain the event. Nonetheless, there is considerable room for the future development of more sophisticated reconstruction algorithms.

5.1 Maximum Likelihood Reconstruction

Event reconstruction in AMANDA has traditionally made use of a maximum likelihood method. The particulars of the method in use have been described in detail elsewhere [66], so only a brief review and some commentary will be given here.

The maximum likelihood method is a generalization of the χ^2 method, and in fact in the limit of Gaussian uncertainties the likelihood is simply related to χ^2 by $-2 \ln \mathcal{L} = \chi^2$. By using the likelihood method we are able to use detailed information about the optical properties of the ice; the measured hit times will not be normally distributed about the Cherenkov time³ but will instead have a broad tail at late times. This shift is due to scattering in the ice, and becomes stronger as the distance between the muon and the OM increases, and also as the angle between the PMT's optical axis and the track deviates from the Cherenkov angle.

A detailed photon propagation Monte Carlo [67] gives us a numerical description of photomultiplier response to muon tracks or electromagnetic showers of various energies at various distances

²The light deposition profile assumed for the present reconstruction is that of a bare muon, so in effect we have fixed the energy at around 100 GeV.

³The Cherenkov time is defined to be that time at which an unscattered photon emitted at the Cherenkov angle would be expected to hit the optical module.

and angles from an optical module in the ice. This information can be used to formulate the likelihood function $\mathcal{L}(R_i | H)$ that the hypothetical physical process would produce the response R_i in OM i , which contains all available information about light propagation both for Cherenkov light and for light emitted in catastrophic events. Given an event E_0 , which is the collection of responses R_i , and a hypothesis H_j , we calculate the likelihood⁴

$$\mathcal{L}(E_0 | H_j) = \prod_i \mathcal{L}_i(R_i | H_j) \quad (5.1)$$

that the hypothesis, if true, would have generated the observed pattern of hits. The hypothesis is then allowed to vary, and some optimization routine is used to find the location H_0 of the global extremum⁵ of \mathcal{L} .

In practice, a number of approximations are used to simplify and speed up the calculation and optimization. As described above, the space H is conventionally simplified by approximating all muons and muon bundles of whatever energy as single, minimum-ionizing muons without stochastic losses, reducing H to the five-dimensional function $H = H(\vec{x}, \theta, \phi)$. For ease of calculation, the numerical descriptions of photomultiplier response (the ‘photon tables’ [67]) are approximated by analytic functions [66]. Rather disturbingly, the timing likelihoods are all normalized to unity; this is equivalent to neglecting the probability that a tube be hit or not hit. This means that the reconstruction is not disturbed if a tube two meters from a track fails to fire, nor if twenty tubes are hit 200 m from the reconstructed track. Finally, the PMT pulses are normally assumed to be due to single photoelectrons irrespective of pulse amplitude and duration, which allows us to ignore the details of the hardware response and the statistics of multiple pulses. There has been some discussion of the validity and necessity of these approximations [49, 66, 68], but they are sufficient to allow track reconstruction.

There is a common but incorrect tendency to interpret H_0 as the hypothesis which was most likely responsible for the observed hits E_0 . $\mathcal{L}(E_0 | H)$ contains no information about the likelihood of H itself, only about the likelihood of E_0 *if H is taken as given*. H_0 represents the track hypothesis

⁴This formulation assumes that all tubes respond independently, an assumption which does not hold in the presence of cross talk.

⁵A note on terminology: we are interested in the maximum of the likelihood function, but the function actually calculated is $-\ln \mathcal{L}$, whose minimum corresponds to the maximum of \mathcal{L} . The terms ‘maximize’ and ‘minimize’ are therefore used, rather counterintuitively, as synonyms.

which, if true, would have had the highest likelihood to have produced the event E_0 . In other words, H_0 is the hypothesis for which the observed event is most consistent with the hypothesis. What we are really interested in, for the purpose of track reconstruction, is $P(H | E_0)$, the probability that the observed pattern of hits E_0 was in fact the traces of a muon $H(\vec{x}, \theta, \phi)$; we are interested in maximizing the probability of the reconstruction, not its consistency.

5.2 Bayesian Reconstruction

Bayes' theorem allows us to relate $P(H | E_0)$ to $\mathcal{L}(E_0 | H)$. Bayes' theorem states that

$$P(A|B) P(B) = P(B|A) P(A).$$

Identifying A with the physical hypothesis H and B with a hit pattern E and dividing, we have

$$P(H|E) = \frac{\mathcal{L}(E|H) P(H)}{P(E)}. \quad (5.2)$$

$P(E)$ is just the probability that a given pattern of hits is observed, which is constant for any particular event E_0 and thus is irrelevant to the question of event reconstruction. \mathcal{L} is the likelihood described in Eq. 5.1 that the given set of hits would be generated by the hypothesis of interest. The new factor in Eq. 5.2 is $P(H)$, the “prior” probability of the hypothesis $H(\vec{x}, \theta, \phi)$. $P(H)$ is called the prior because it does not depend in any way on the actual measurement, and so can be calculated prior to the measurement. $P(H | E)$ is the “posterior” probability, the probability of the hypothesis after the additional evidence E is taken into account.

The need to figure the prior into the likelihood stems from the fact that AMANDA is, by design, “a detector that barely works” [45]. In an ultra-high precision neutrino telescope, say one with a density of one optical module per cubic meter, a single track would be observed by so many modules that only one solution would be possible; the likelihood $\mathcal{L}(E_0 | H)$ would fall off extremely quickly away from the maximum at H_0 [69]. The vanishingly small \mathcal{L} will kill $P(H | E_0)$ except at H_0 , regardless of the shape of $P(H)$. But in the interest of maximizing effective volume, AMANDA was designed to collect information of the minimum precision necessary to allow track reconstruction. This means that events often have broad, shallow (or even multiple) maxima in \mathcal{L} . In such a detector, the shape of $P(H)$ can no longer be ignored.

$P(H)$ is a probability, and so must obey the relation

$$\int_H P(H') dH' = 1.$$

Since we have already decided to neglect the normalization constant $P(E)$, we need not be concerned about the absolute scale of $P(H)$, but it must be noted that $P(H|E)$ depends on the set of hypotheses $\{H\}$ allowed into consideration; this is in contrast to \mathcal{L} , which can be calculated from only the single hypothesis H . For example, if we were to search for neutrinos from Markarian 501, the assumed flux input to the prior would affect the significance of the detection — if we assumed a large flux we would tend to reconstruct more events as coming from the source. Clearly, extreme care in interpretation would be required in that case. However, for the present purpose, we are interested only in taking into account the relative frequency of downgoing muons, which are many orders of magnitude more common than upgoing muons. On this scale, even a simple approximation to the correct prior will lead to a huge increase in accuracy.

In the present analysis, we begin by assuming that all events are triggered by muons. This is in fact not the case, but the resultant prior should be effective in rejecting downgoing muon events. We further make the conventional approximation of all muons and muon bundles as single bare muons, and we neglect any dependence of the muon flux on depth over the height of the detector. Finally, we neglect the variation, approximately a factor of two, of the atmospheric neutrino flux with zenith angle over the upgoing hemisphere (cf. Fig. 3.8). These approximations, and the azimuthal symmetry of the muon flux, reduce the hypothesis space $\{H\}$ to a one-dimensional function $H(\theta)$.

To actually calculate $H(\theta)$ we rely on Monte Carlo. Properly, we should simulate cosmic rays and atmospheric neutrinos, propagate the resultant muons to the detector, and simply tally the number of muons which reach the array. In the present work, we have approximated this approach by using the trigger-level true angular distribution of the simulated events, shown in Fig. 3.8. This will distort the prior by emphasizing angles which contribute more high energy muons or in which the muon passes more OMs and therefore is more likely to trigger the array, but such distortions should be in any case smaller than the 10^6 up/down asymmetry observed at trigger level. The angular distribution is then parametrized with a high-order polynomial, as shown in Fig. 5.1, and used as a multiplicative factor to the likelihood \mathcal{L} in the reconstruction. As noted above, the prior




Figure 5.1: The model used for the prior function. The prior is flat over the upgoing hemisphere, and the simulated trigger-level angular distribution of downgoing muons is used for the downgoing hemisphere.

shown in Fig. 5.1 neglects the angular dependence of the atmospheric neutrino flux; this was done deliberately to allow direct comparison with the results of maximum likelihood analyses. In any case, this modulation of the neutrino flux is quite small on the scale of Fig. 5.1.

It seems that most misunderstandings regarding the Bayesian likelihood formulation derive from a confusion of the hypothesis H with the event E . It should be stressed that the hypothesis contains absolutely no information related to the actual experimental observation; it contains only knowledge of the prior conditions under which the experiment is performed. In the case of AMANDA, with the approximations described above, these prior conditions are simply the known fluxes of muons from cosmic rays and atmospheric neutrinos.

The Bayesian fit is attempting to find the direction that, given all the available information, most probably explains the observed event. One may think of the prior information as a sort of

“tie-breaker,” indicating a preference for certain reconstructions when the hit information is not sufficiently constraining to reconstruct a track precisely. In mathematical terms, the prior has an effect when either (a) there are two local minima of comparable depth in the likelihood function for which the prior has significantly different values, or (b) the value of the prior changes significantly over the breadth of a single minimum in the likelihood function.

The effect of the Bayesian fit is particularly strong near the horizon; cf. Fig. 7.18. For a horizontal track in the AMANDA-B10 detector, there is a considerable amount of ambiguity in the reconstruction; the short track length of a horizontal muon simply does not constrain the fit tightly; the hit information is not strong enough to rule out the possibility of an atmospheric muon from above $\cos(\theta) = 0.2$. Since the prior indicates that such tracks are more likely than shallower ones, the reconstruction properly chooses the steeper track.

Clearly, there are improvements which should be made to the Bayesian likelihood calculation. However, even the rather crude approximations which were used in this analysis are sufficient to produce a major improvement in the accuracy of the reconstruction, allowing for a much simpler analysis overall than is required with the conventional maximum likelihood method.

5.3 Minimization

Once the likelihood function (or for this analysis the posterior probability function) has been constructed, there is the question of actually finding the hypothesis H_0 which maximizes the function on the space. This is the task of the minimization routine, a general-purpose program which is interfaced with the AMANDA software. Minimization is a topic with many applications, and so there is a considerable literature on the question.

Minimization algorithms consist, in general, of two pieces: a global minimizer and a local minimizer. The global minimizer searches through the entire space on which the function is defined, attempting to find the rough position of the global minimum in a relatively efficient manner⁶. Once the approximate location of a minimum is found, it is passed to the local minimizer, which refines the solution. There are many competing algorithms on the market to accomplish each of these tasks.

In AMANDA, we use Powell’s direction set method, as implemented in Numerical Recipes [70],

⁶A straightforward grid search becomes prohibitively slow as the dimensionality of the problem increases.

for local minimization⁷. Powell's method consists, essentially, of finding a set of directions which are independent in the sense that minimization in one dimension does not spoil previous optimization in another direction. For a function of simple topology, the convergence is quite fast, nearly quadratic. The method may be confused by complex topologies, however, and there are reports of difficulties when minimizing in more than five dimensions [71], which may limit the future applicability of the method. Other methods may also be faster for typical AMANDA reconstructions [72]. For now, however, the method gives workable results.

Powell's method will find a minimum, but there is no guarantee that it is in fact the global minimum, not just a false local minimum. Indeed, the algorithm is often observed to take a down-going first guess from the line fit, and turn it around to produce an obviously incorrect upgoing reconstruction. AMANDA has thus taken to using a simple global minimization algorithm, a random multi-start scheme. After the first minimization (based on the line fit) a track extending in a random direction from the point of closest approach of the previous reconstruction to the event center of gravity is chosen as a new first guess and the event is given back to Powell's local minimizer. Typically the process is repeated a number of times; four to twenty successive local minimizations are typical. This method is not terribly efficient; the same minimum is generally found many times. However, with enough restarts the hypothesis space will be searched fairly thoroughly, and the chance of being deceived by a false minimum are quite small.

⁷There has also been experimentation with the MINUIT minimizer, and with a downhill simplex routine in the context of a simulated annealing global minimization (see Appendix F.3 for a further discussion), but Powell's method is used exclusively at present.

Chapter 6

Data

6.1 Triggering and Live Time

AMANDA was operated in 1997 with a nominal 16-fold majority-logic trigger, meaning that 16 channels were required to fire within a sliding window of $2 \mu\text{sec}$ to form a trigger. The *de facto* trigger level was somewhat lower than the nominal level because of instrumental effects. Electronic malfunctions caused some channels to “ring” at times, firing at up to 100’s of kHz instead of the normal 1 kHz. While a tube was ringing, it sent a near-constant signal to the trigger electronics, effectively lowering the trigger multiplicity by one channel. Furthermore, cross talk in the electronics or cables caused some pulses to be registered in more than one channel, which also lowered the effective trigger requirement. Finally, problems in the trigger electronics stretched the trigger pulses coming from channels on strings 9 and 10 to nearly $10 \mu\text{sec}$, rather than $2 \mu\text{sec}$ as was intended, which widened the trigger gate slightly [73]. With all of these effects, the actual trigger threshold observed in the 1997 data is approximately 13 hits. In fact the details of the trigger window are not particularly important for this analysis. The requirements imposed offline by the quality cuts are so stringent that events with less than 15 hits are effectively excluded, as shown in Fig. 6.1.

AMANDA-B10 recorded data over the austral winter of 1997. A bug in the data acquisition software rendered the first several weeks of data unusable, and equipment failure in midwinter introduced additional down time until repairs were made. In normal operation, detector deadtime due to event readout was approximately 25%. In total, $1.124 \cdot 10^7$ sec, 130.1 days, of exposure were collected from April to November of 1997 [74]. Improvements in data acquisition have reduced the deadtime

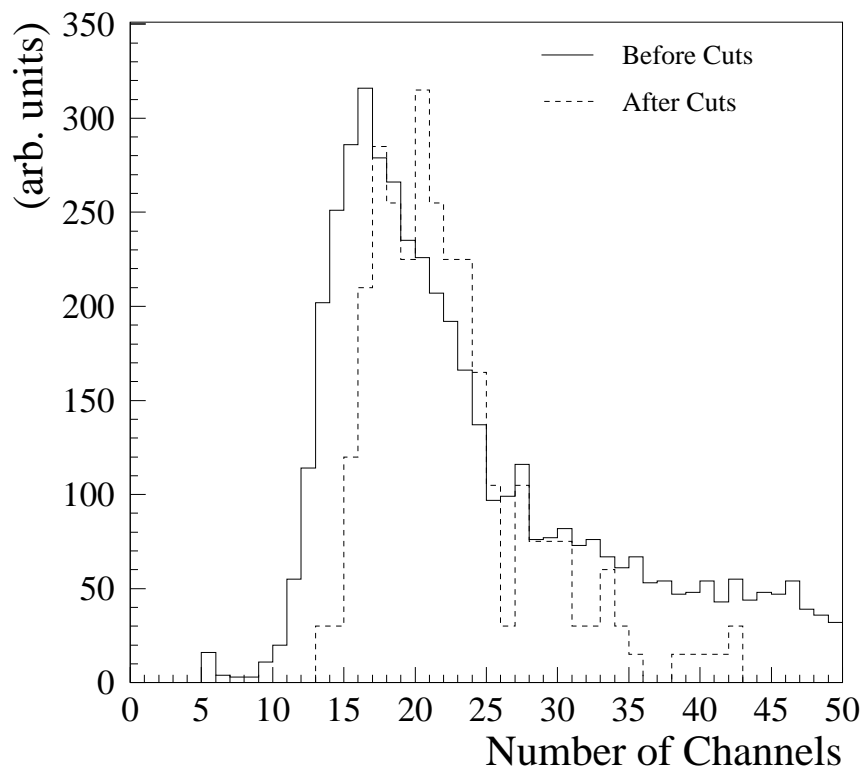


Figure 6.1: The number of channels participating in an event, after hit cleaning as described in Sec. 6.2. The solid line shows the distribution for events reconstructed as upgoing; note that the threshold is considerably below the nominal trigger level of 16 channels. However, very few low multiplicity events survive the cuts imposed during offline analysis, described in Chap. 7.

in the DAQ, and the total exposure in subsequent years has been higher.

6.2 Hit Cleaning

In practice the data recorded by the AMANDA detector are not perfect. In every event there are apparent hits which are actually due to noise of various types. Bioluminescence is obviously not a major problem for AMANDA, and the ice itself is extremely quiet. However, there is dark noise in the PMTs themselves as well as noise produced by radioactive potassium in the OM glass. These noise hits are at least approximately random, although non-Poissonian behavior¹ has been seen [71]. There is also cross talk between channels, induced either in the strings or in the surface electronics or both, which is highly nonrandom. PMTs are subject to pre- and afterpulsing, although prepulsing rates are very low and afterpulsing is easily identified by its characteristic time delay with respect to the primary pulse, which is longer than the duration of a typical event. Finally, there is electronic noise in the DAQ system.

The AMANDA reconstruction scheme begins with various hit “cleaning” algorithms, which attempt to remove these false hits from the event in order to prevent the reconstruction algorithms from becoming confused by spurious signals. It is trivial to include random noise in the maximum likelihood calculation, of course, but the first guess algorithms are more sensitive, and in any case cross talk is very difficult to incorporate without extremely detailed information about the hardware response, which is not yet available. Furthermore, the particular form of the likelihood function currently in use in AMANDA, based on timing alone, is peculiarly vulnerable to noise hits. One or two such hits can be enough to confuse the reconstruction, turning a downgoing cosmic ray muon into an apparently upgoing track. More sophisticated implementations of the likelihood function should be able to avoid such errors, but such algorithms are more time-consuming both in terms of computer cycles and of development effort.

The first stage of hit cleaning is to reduce the event duration to a window of 4500 ns around the trigger time. This eliminates random noise at the beginning of the event and afterpulsing at the end, while retaining essentially all interesting hits. Minimal amplitudes and times over threshold are then

¹This effect has not yet been satisfactorily explained. It may have to do with radioactive decays in the OM glass exciting atoms to metastable states which later decay in turn. However, the time scales of the anomalous noise are as long as milliseconds.

required of all hits, eliminating most electronic noise and cross talk. Any secondary pulses in a single channel are also discarded, as they will have been due to highly scattered photons which provide poor timing information at best, and because the reconstruction algorithms are not yet sufficiently sophisticated to make good use of the information. Finally, an isolation cut is applied, eliminating a hit if no other tube within 70 m fired within 500 ns. This cut is highly efficient at eliminating random noise.

Inspection of highly reduced data sets, however, reveals that most misreconstructed tracks are produced by the superposition of a few noise hits onto a real event, either a low energy downgoing muon track passing through the detector or a bright catastrophic event presumably produced by a higher energy track just outside. These noise hits will occasionally be random coincidences of dark noise — which would survive the isolation cleaning — but more often appear to be caused by cross talk. Cross talk by its nature is correlated with other hits, and so isolation criteria are ineffective at cleaning it from the event. Although much cross talk can be removed with TOT cuts, it has proven necessary to implement additional, more sophisticated algorithms to further reduce the number of such hits. Details on these algorithms can be found in Appendix B. The use of these algorithms was found to reduce the number of misreconstructions in the 1997 data set from nearly 60,000 events to only 3,000 events, and made possible the relatively simple analysis presented here. However, it should be stressed that these algorithms, in the absence of detailed cross talk measurements, were developed rather intuitively — they are imperfect and it is expected that significant improvements could be made.

6.3 Filtering

The 1997 AMANDA data set consists of approximately 0.5 TB of events, each event comprising about 0.5 kB of information. Maximum likelihood reconstructions as presently implemented consume some 0.2–0.5 CPU-sec per minimization, so a full reconstruction of each of the $1.2 \cdot 10^9$ events would require CPU-decades, well beyond the present computing resources of the AMANDA collaboration. It was thus necessary to apply some faster filtering algorithms to the data set to reduce it to a more manageable size.

The filtering on which this analysis is based was performed by the AMANDA group at Lawrence

Berkeley National Laboratory, using the resources of the NERSC computing center [75]. The filtering produced several data sets, one of which was designed to enrich the population of low energy upgoing muons. This filter proceeded in two steps; in the first stage, a line fit was calculated for each event. The line fit is an analytic form which produces a vector describing the flow of hits through the detector [76]; the times and positions of hits are assumed to follow a linear relation, $\vec{r}_i = \vec{r}_0 + (t_i - t_0)\vec{v}$. Minimizing χ^2 produces the solution

$$\vec{v} = \frac{\langle \vec{r}_i t_i \rangle - \langle \vec{r}_i \rangle \langle t_i \rangle}{\langle t_i^2 \rangle - \langle t_i \rangle^2}, \quad \vec{r}_0 = \langle \vec{r}_i \rangle - \vec{v} \langle t_i \rangle,$$

where brackets denote averages over the hits i . This solution gives an approximate direction for the muon track, as well as a vertex location and a velocity². The first stage of the filter removed events from the data set if the line fit direction was steeply downgoing, $\theta_{LF} < 50^\circ$.

This cut on the line fit reduced the data set by 78%. With this reduced number of events it was possible to perform a full reconstruction on the remaining set. The second stage of the filter was thus a maximum likelihood (*i.e.*, non-Bayesian) timing fit with a single minimization performed. Those events that were reconstructed as upgoing muons, $\theta_{like} > 80^\circ$, with three or more hits due to unscattered (in a window of $[-15 : 25]$ ns) photons passed the second stage of the filter. This second level of filtering rejected 91% of the remaining events.

Overall, the filter reduced the data set by about a factor of 50. The cost in signal was around 50%, based on Monte Carlo simulations of the atmospheric neutrino flux. This loss is unfortunate but unavoidable with the CPU resources presently available. Also, it should be noted that many of the neutrino events lost are probably unreconstructable in any case; for example, a low energy horizontal neutrino which interacts within the detector will produce a nearly spherical pattern of light from the interaction vertex with a radius comparable to the detector radius, with only a short track leading out. Such an event cannot be reliably reconstructed with the information available from AMANDA-B10. The total loss in number of triggered events is thus an overestimate of the number of “reconstructible” neutrinos lost.

²The muon itself, of course, travels at $\beta = 1$, but due to scattering in the ice the average velocity $\Delta\vec{x}/\Delta t$ between hits may vary. $|\vec{v}|$ for neutrino events ranges from about $0.25c$ to c .

6.4 Reconstruction

The LBL filter reduced the 1997 data set to some $2 \cdot 10^7$ events, few enough to permit relatively intensive reconstruction procedures. These data were subjected to a series of reconstructions, described in detail in Appendix C. The idea of the reconstruction chain was to find the best upgoing reconstruction for each event through a multistart minimization restricted to the upgoing hemisphere. After the best upgoing fit was found, a very thorough search for downgoing minima of higher quality was undertaken. These reconstructions used the Bayesian prior, which greatly reduced the number of false upgoing reconstructions by accounting for the large downgoing muon flux. Furthermore, the opportunity was taken to correct an error in the hit cleaning routines, which allowed some isolated noise hits to be retained. After this series of reconstructions, we reduced the data set to those events reconstructed as upgoing, $\theta_{best} > 90^\circ$. This requirement reduced the data set to only some $6 \cdot 10^4$ events, quite an impressive performance by the reconstruction.

Inspection of the remaining events revealed that the vast majority, almost 90%, belonged to a peculiar class of events, referred to as ‘COG’ events because the center of gravity of the hits in these events clustered very strongly in certain locations in the detector. An example of this type of event is shown in Figure 6.2. This geometric clustering suggested an instrumental origin, which was eventually identified as cross talk contamination of bright events presumably caused by catastrophic energy deposition by downgoing muons just outside the detector. To reduce this background, a new series of hit cleaning algorithms specifically designed to reduce cross talk was developed, as described in App. B. The remaining events were passed through another, slightly shorter, series of reconstructions with the new hit cleaning. Only $5.4 \cdot 10^3$ of the data were still fit as upgoing with the cross talk thus reduced.

In examining these five thousand events, it was found that there was another class of instrumental fake, characterized by simultaneous hits in many channels on the outer six strings with only a handful on the inner detector, as shown in Fig. 6.3. An event of this class was noticed in [36] but it was not realized that the malfunction was a recurring one. A plausible explanation is that these events are caused by oscillations in the high voltage levels supplied to the OMs, which would explain the simultaneity of the pulses and the fact that the OMs seem to fire in blocks of channels.

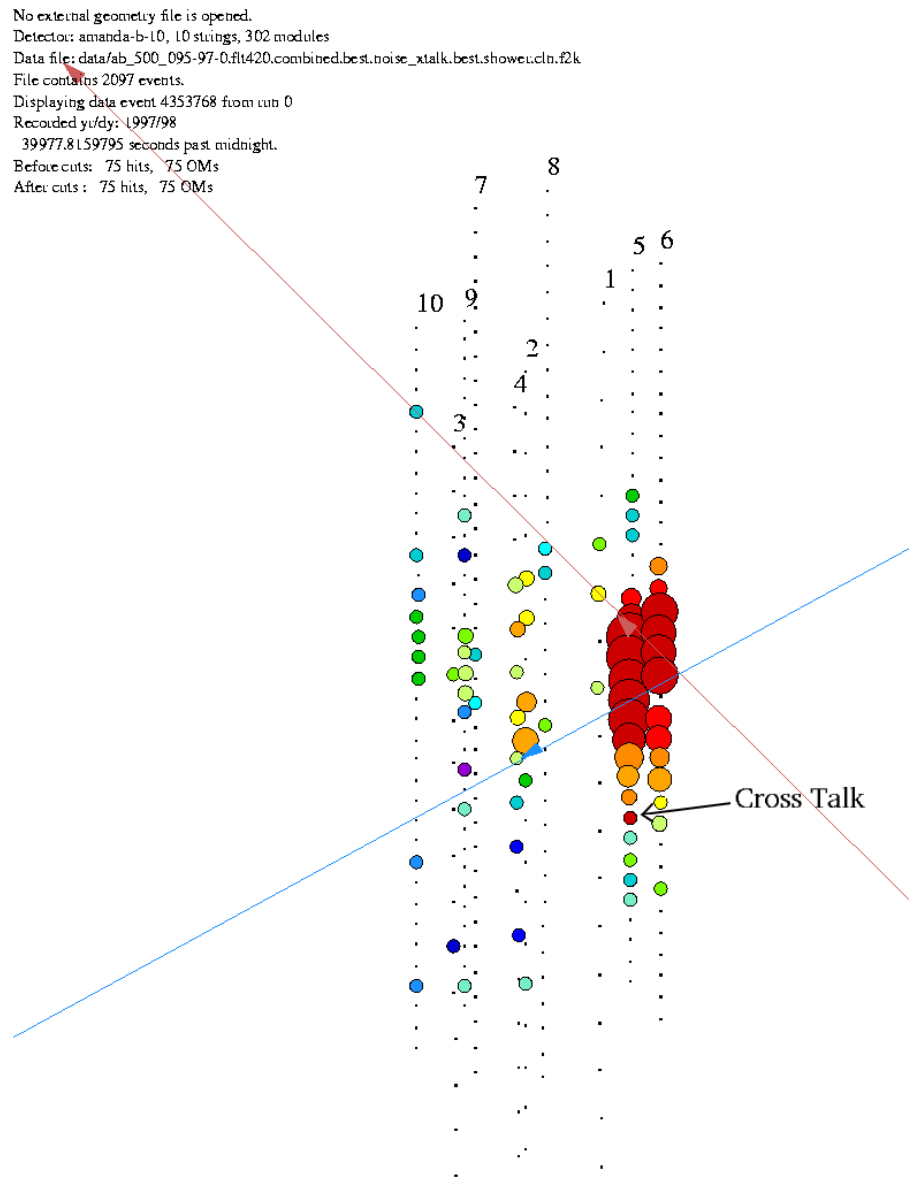


Figure 6.2: A ‘COG’ event. The event itself is real, probably caused by a bright stochastic event at the edge of the detector. The reconstruction has been confused by cross talk in channel 114, the fifth hit from the bottom on string 5, indicated by the arrow.

HV oscillation was observed in the 1997–98 austral summer [77], and was initially suspected as the culprit responsible for the ‘COG’ events, although that hypothesis was eventually rejected. These events populate a distinct region in the plane of the numbers of OMs hit in the six outer strings of the detector vs. the number hit in the four inner string, as seen in Fig. 6.4, and are easily eliminated with a cut on that plane.

With both the physical and the instrumental backgrounds greatly reduced, if not completely eliminated, only a few simple cuts were necessary to isolate a nearly pure subset of real upgoing events, as described in Chapter 7.

No external geometry file is opened.
 Detector: amanda-b-10, 10 strings, 302 modules
 Data file: data/ab_500_095-97-0.flt420.combined.best.noise_xtalk.best.showe.cln.f2k
 File contains 2097 events.
 Displaying data event 2426367 from run 0
 Recorded y/d/y: 1997/103
 35854.3124364 seconds past midnight.
 Before cuts: 34 hits, 34 OMs
 After cuts: 34 hits, 34 OMs




Figure 6.3: An ‘instrumental’ (non-physical) fake, probably caused by oscillating high voltage levels. The hits arrive at the DAQ at the same time, but are projected backward in time according to the cable delays on each channel, forming an apparently upgoing pattern that confuses the reconstruction.

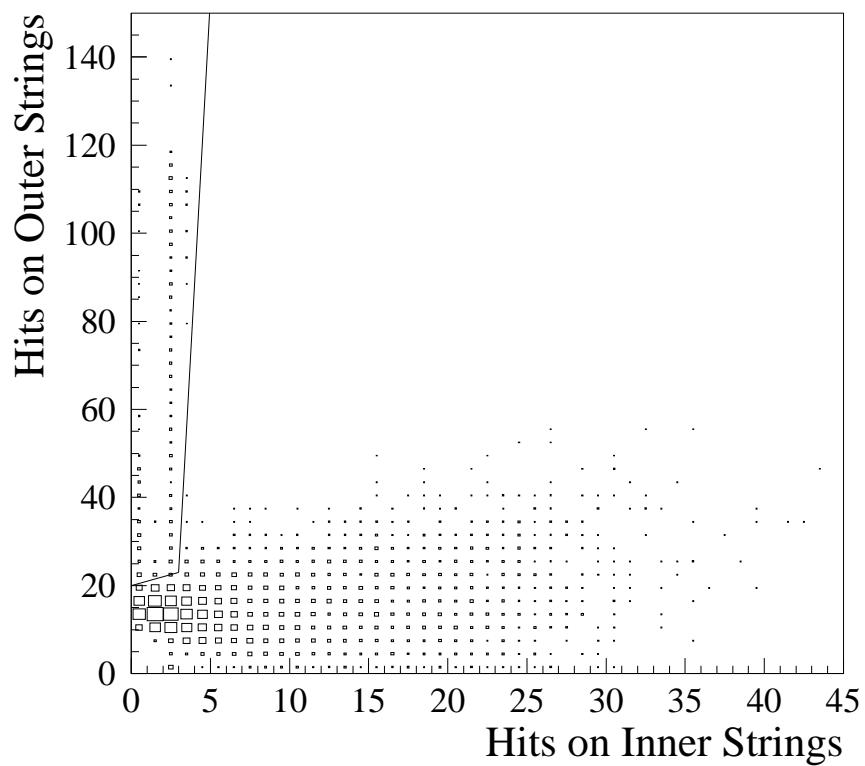


Figure 6.4: The number of channels hit in the inner part of the detector compared to the number hit on the outer strings. Real events form a broad band at the bottom of the plot. The band with large multiplicities on the outer strings and very few hits in the center are instrumental artifacts, probably caused by oscillating high voltage levels.

Chapter 7

Atmospheric Neutrinos

The main goal of this work is to demonstrate the observation of atmospheric muon neutrinos by AMANDA. These neutrinos constitute a convenient flux of (approximately) known number, angular distribution, and energy spectrum, which can be used to calibrate the response of the detector. Because there is no artificial source of high energy neutrinos that can be turned on and off, nor any known astrophysical source that can be tracked, proving that the detector actually works is difficult. The case ultimately rests on plausibility and consistency arguments: essentially, that the observed data agree with simulations of cosmic ray muons and atmospheric neutrinos.

The filtering and reconstruction algorithms described in Chapter 6 and App. C reduced the 1997 data set from more than a billion triggers to a manageable five thousand apparently upgoing events. However, these algorithms are not perfect; only some 570 atmospheric neutrino events are expected, and the remainder of the data must be misreconstructions, fakes, of some sort. The atmospheric muon Monte Carlo predicts that 2,000 of the fakes are caused by downgoing muons, (see Table D.2) and we have seen that instrumental effects such as cross talk and HV oscillations can also create apparently upgoing tracks, either by modifying real downgoing events in unexpected ways (in the former case) or by creating the events entirely (in the latter). Whatever the nature of the fakes, we must impose some additional requirements ('cuts') on the quality of the reconstructed events to get down to a relatively pure neutrino sample.

7.1 Cuts

No single measure of the quality of a reconstructed event has yet been identified for AMANDA data that could fill the role of, say, χ^2 in a simpler detector. There is the likelihood, of course, but as it is presently implemented the likelihood ignores much of the relevant information in the event, reducing it to at best a moderately useful parameter. A large number of other parameters have been identified and used in various analyses, with varying degrees of success [74, 78, 79, 80, 81]. One advantage of the sophisticated reconstruction algorithm used in this analysis is that it is not necessary to use large numbers of complicated cuts to isolate neutrino events. In fact, a simple set of cuts on six parameters, described below, will be seen to be sufficient.

Distributions of the cut parameters, comparing data to simulated atmospheric neutrinos, are shown in Figs. 7.1–7.8. In each of these figures, the cuts on the five parameters not plotted have been applied, and the level of the cut on the plotted variable is indicated with a dashed line. The plots are constructed for event quality (see Sec. 7.3) of $Q \geq 7$, the level at which a large, reasonably pure sample of atmospheric neutrinos is obtained, as will be shown in Sec. 7.5. The dots indicate the number of events observed in the 1997 data set. The size of the hatched boxes indicates the statistical precision of the atmospheric neutrino simulation; the lines show the 1σ Poissonian fluctuations expected. In each case the number of simulated events passing the $Q = 7$ cuts has been normalized to the number of data passing.

7.1.1 Reduced Likelihood

As discussed above, a Bayesian maximum likelihood fit is performed to fit muon tracks to the observed events. The value of the likelihood (actually, the Bayesian posterior probability) obtained for the best fit is the simplest indicator of the quality of the reconstruction. The actual functional form used is the negative logarithm of the likelihood, which in the case of Gaussian errors and a uniform prior is proportional to the χ^2 of the fit. There are $N_{\text{hits}} - 5$ degrees of freedom in the fit, a leading edge time for each hit less the five variables fit (a vertex position and two angles). We thus cut on the reduced likelihood

$$L = \frac{-\ln \mathcal{L}}{N_{\text{hits}} - 5},$$

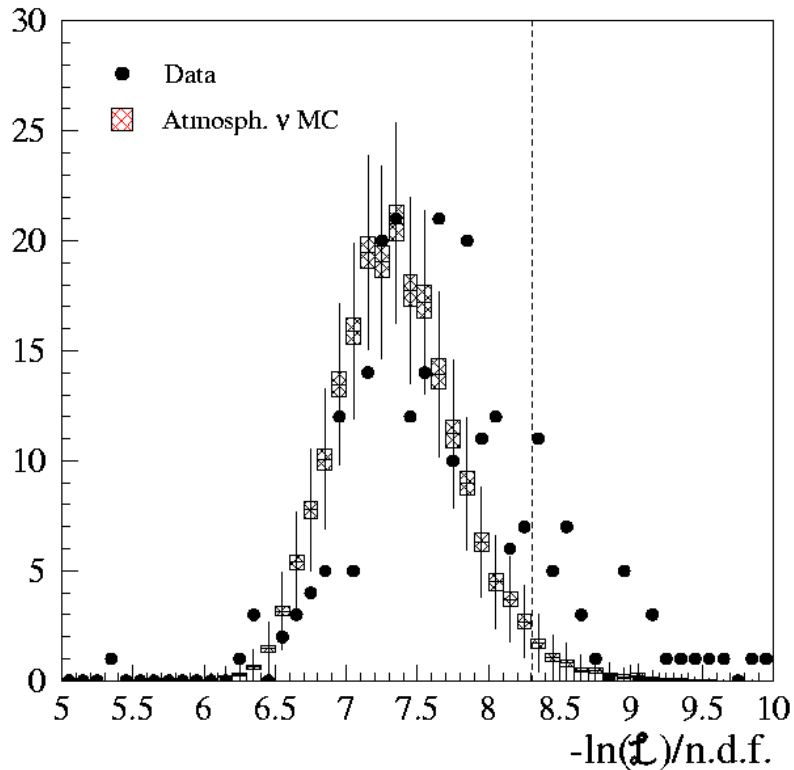


Figure 7.1: The reduced likelihood of the fitted track. The data are shifted to slightly higher (i.e., worse) values than the Monte Carlo, but there is a clear shoulder of misreconstructed tracks above about 8.4 that are removed with this cut. See page 60 for an explanation of the error bars.

as shown in Fig. 7.1.

For the upgoing hemisphere, the Bayesian prior has a constant value¹, producing no distortion of the traditional likelihood space, so the value of this parameter is directly comparable to those produced in traditional reconstructions. However, to keep the likelihood less than unity over the whole sphere (a technical requirement of the reconstruction software), an arbitrary normalization factor was applied to the prior. This factor, an additive offset of 13.71 in the logarithm, has therefore been subtracted from the value of $-\ln \mathcal{L}$ returned by the software, in order to make the values of L directly comparable to those obtained from traditional maximum likelihood fits.

¹This is a simplification of the actual atmospheric neutrino flux, which is a factor of two larger at the horizon than vertically upgoing.

7.1.2 Number of Direct Hits

Photons observed in the optical modules are considered “direct” if scattering in the ice between production and detection delays the photon by only a slight amount. The delay is measured relative to the time predicted for an unscattered Cherenkov photon emitted from the appropriate point on the fitted track. Different delay windows are possible; for counting direct hits, we have used a window of $[-15:75]$ ns. The negative extension of the window accounts for errors in geometry and calibration and for fluctuations in the rise time of the pulse due to PMT and TDC discriminator response.

This cut parameter entered the analysis twice. A requirement of three direct hits was imposed as part of the LBL filter [75], described in Section 6.3. In the filter the count was done including hits which were rejected for purposes of reconstruction by the noise cleaning; the cut was reapplied in this analysis after cleaning. The second cut is on the fraction of direct hits

$$\frac{N_{\text{dir}}^{[-15:75]}}{N_{\text{hits}}},$$

which we find more effective for background rejection than requiring an absolute number of direct hits. However, this formulation does reduce the sensitivity to high energy events, in which the mean visual radius is larger and more scattered light is collected; the effect of the cut on the energy spectrum can be seen in Figure 7.3. This is a conscious choice; high energy muons produce many late hits, but events with many hits reconstructed as late are usually fakes, not high energy tracks. We choose to follow the conservative approach of rejecting everything that does not look like an atmospheric (i.e., low energy) neutrino-induced muon, rather than tailoring the cuts to remove only the fakes predicted by the downgoing muon simulation. The latter approach is highly vulnerable to fakes that are not predicted by the Monte Carlo; since in searching for atmospheric neutrinos we have the luxury of sacrificing sensitivity to high energy events, we will do so to improve our background rejection.

7.1.3 Sphericity

High energy muon tracks are very long compared to the typical perpendicular distance at which their Cherenkov emission can be observed in the ice. The light from a muon track should thus be observed in a long, narrow cylindrical region of the detector. If a very bright stochastic event occurs along the track (emission of a hard bremsstrahlung photon, for example), the resulting cascade




Figure 7.2: The percentage of hits produced by unscattered photons. The cut parameter is restricted to rational values, which produces some spikiness. The data are systematically higher than the simulation, mostly due to discrepancies in N_{ch} — less scattered light is collected than is predicted.




Figure 7.3: Effect of the cut on the percentage of direct hits. Events in which large amounts of energy are deposited in the detector are lost, because so much scattered light is detected. The efficiency is shown for upward reconstructed tracks after all other $Q \geq 7$ cuts.

will produce a roughly spherical pattern of hits — the typical length of a cascade (a few meters) is less than the OM spacing, so the light appears to come from a point source. If the cascade is too bright, it will effectively obscure the underlying track, making reconstruction of the track difficult and producing an unreliable fit. We therefore require that the light observed in an event not be overly spherical.

The sphericity of the event is defined by thinking of each hit as a point mass, and calculating the tensor of inertia of the collection of hits. Diagonalizing the tensor, we obtain as eigenvalues I_i the moments of inertia about the principle axes of rotation. For a spherically symmetric collection of hits, these moments will be of equal magnitude; i.e., the eigenvalues will be degenerate. For a long cylindrical track, one moment will be much larger than the others. We can then reject spherical events by cutting on the normalized magnitude of the smallest moment $I_1/\sum I_i$. Spherical events have large smallest moments, so low values are required.

7.1.4 Track Length

We define the track length by projecting each of the direct hits onto the reconstructed track, and measuring the distance between the first and last hits. For this purpose we use a stricter definition of direct hits than when simply counting the direct hits: here the window is $[-15:25]$ ns. By requiring a reasonably long track length, we remove two classes of events. The more important class consists of misreconstructions: cases in which the event is very cascade-like and no track can be made out, or in which the fitter has converged to an incorrect local minimum of the likelihood function. In these cases the reconstruction typically finds a few direct hits — in the area where the incorrect fit crosses the true track — but they are highly localized. Requiring the direct hits be spread along the track is thus a consistency check on the reconstruction.

The second class of events rejected consists of approximately correct reconstructions of very short tracks (a few tens of meters), which because of the sparsity of the array cannot be reconstructed with precision. These events are produced either by low energy neutrinos or by muons which pass along the very edge of the instrumented volume. For the present analysis a low energy threshold is not necessary, and so the reduction in sensitivity to low energy events was considered acceptable.




Figure 7.4: The sphericity of events. This cut is largely parallel to other parameters but does reject some ‘COG’ fakes. High values indicate spherical events, low values more cylindrical topologies as expected for muon tracks.

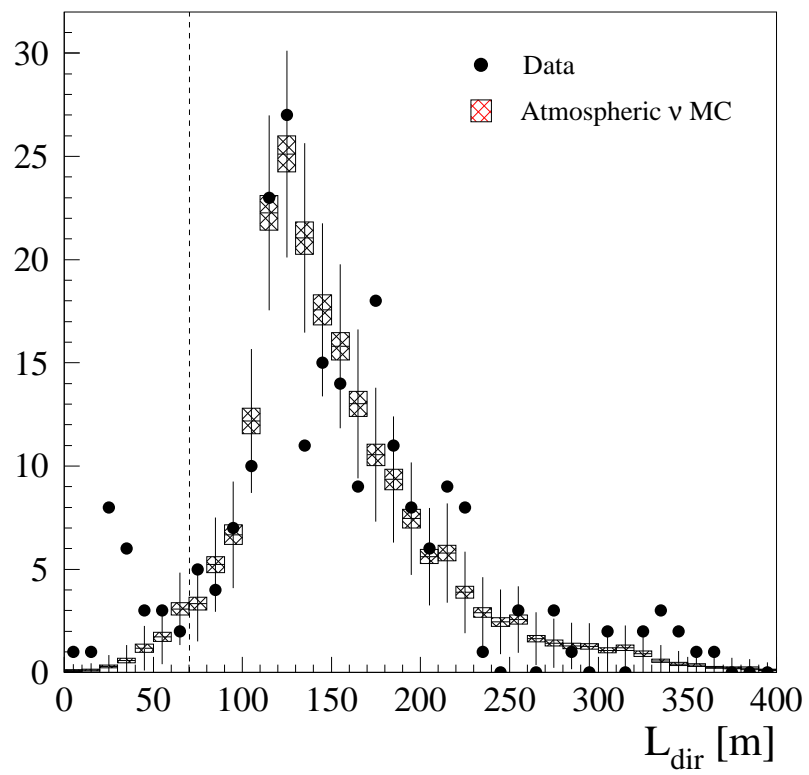


Figure 7.5: The distance over which muons were tracked through the detector. Many misreconstructed tracks have lengths of 50 m or less, although low energy neutrino events can also be quite short.

7.1.5 Line Fit Zenith Angle

The first fit performed on each event is the line fit, described in Sec. 6.3. This is an analytic calculation, as opposed to the minimization performed in the full likelihood reconstruction, and is used as the initial hypothesis for the likelihood minimization. Moreover, it provides a very simple description of the track, not dependent on a minimization and completely independent of all details of light propagation, optical module response, and so forth. The line fit thus serves as an important check of the full reconstruction, ensuring that the minimizer has not become confused by noise hits or cross talk.

In this analysis, events were rejected if the full fit found a solution considerably more upgoing than the initial guess of the line fit. One class of background in particular is reduced by this requirement: bright shower-like events in which cross talk causes one or more modules below the shower to fire tend to be reconstructed at steeply upgoing angles, although the line fit is rather horizontal. Note that there is no requirement that the line fit and full reconstruction actually agree, only that the full fit not be significantly more upgoing. This is in contrast to the usage in [74] and [82]. Either approach is reasonable, but in this analysis, we have used the cut to reject the specific class of background, and not more generally to select out events which are simple enough for the line fit to reconstruct.

7.1.6 Smoothness

The “smoothness” parameter is a check on the self-consistency of the fitted track. The event is reconstructed using timing information alone, whereas the smoothness is a purely topological measurement. The parameter measures the constancy of light output along the track; highly variable apparent emission of light usually indicates that the track has been completely misreconstructed, although it could also be that the track was obscured by bright stochastic light emission — which also tends to lead to poor reconstructions. We note that this cut, as with the fraction of direct hits, tends to reduce sensitivity to high energy muons, but in this analysis we are primarily interested in atmospheric neutrinos rather than hard sources. Smoothness was first defined in [80] and studied in some detail in [81]. The formulation of the smoothness parameter used in this analysis is that based on the predicted hit probabilities [81].

The smoothness parameter was inspired by the Kolmogorov-Smirnov test of the consistency




Figure 7.6: The difference in zenith angle between the line fit and the full Bayesian timing reconstruction; positive numbers indicate a shallow line fit that has been pulled to a more upgoing solution which is therefore suspect. The distribution of data is narrower than that of simulated neutrinos, but a clear shoulder due to fakes can be seen above 25° .




Figure 7.7: Definition of the “smoothness” parameter. The distribution of the observed hits is compared to that predicted for a muon emitting Cherenkov light. The prediction is shown as a flat line; in reality it depends on the distribution of modules around the reconstructed track.




Figure 7.8: Distribution of the absolute value of the smoothness parameter; smooth, high quality tracks have values close to zero. The peak of the distribution is slightly offset for data, but the long tail of misreconstructions is easily rejected.

of two distributions; we wish to measure how consistent the observed pattern of hits is with the hypothesis of constant light emission by a muon. The definition of the parameter is illustrated in Figure 7.7: the positions of the hits are projected onto the track, and the cumulative of their distribution along the track is compared to the predicted cumulative. The predicted cumulative (shown as a flat line in Fig. 7.7) is actually based on the distribution of modules around the track, and calculated using the probabilities that the modules fire given their distance from and orientation relative to the reconstructed track.

7.2 Systematic Uncertainties

There are a number of effects that lead to systematic uncertainties in the behavior of the AMANDA detector. The physical processes involved at the energies relevant to AMANDA are not

precisely known, and different models of cosmic ray interactions or of muon propagation give differing predictions of the number of events AMANDA will see. Furthermore, the detector is embedded in a natural medium whose optical properties are difficult to measure precisely, and the process of deploying instruments necessarily includes a melting/refreezing cycle that may drastically affect these properties; for a given physical event, different ice models will yield different predictions regarding the appearance of the event. And of course, as with any detector, the physical hardware is quite intricate and the description used in the simulations is only approximate. In the absence of a calibration source or test beam, these systematic effects are difficult to quantify, although considerable effort has gone into their study.

The physical processes that produce high energy neutrinos in cosmic ray showers are somewhat uncertain. Various extrapolations from measurements at lower energies produce different predictions for the atmospheric neutrino flux. The magnitude of this effect has been studied in [41], and is estimated to be approximately 30% for the energy range measured by this analysis. This is not a systematic uncertainty in the sense that we would like to measure this flux, but it puts a limit on the extent to which the observation of this flux can be used to measure the detector's response.

The rate of muon energy loss at very high energies is not precisely known. Based on comparison of different muon propagation simulations, the linear coefficient b in Eq. 3.4 is uncertain at the level of 5–10% [65, 83]. This uncertainty feeds approximately linearly into the muon range given in Eq. 3.5, for muons whose energy loss is dominated by stochastic effects (above about 500 GeV). The geometrical volume available for neutrino interactions producing muons that reach the detector is $V \sim R_\mu^3$, so the event rate will vary, to first order, as the cube of the energy loss coefficient. A 10% uncertainty on b thus would be expected to give a 30% uncertainty in the rate of high energy muons reaching the detector. For atmospheric neutrinos and atmospheric muons, one might expect the overall effect to be less because most events are of lower energy, but simulations with muon propagators based on [84] and [85] predict event rates differing by 27% after cuts.

The optical properties of the bulk ice have been studied in detail [63], using a number of optical emitters. However, none of these emitters can measure all of the properties of interest. The Cherenkov spectrum is peaked at short wavelengths, so that most detected light is in the deep blue. We have

embedded blue nitrogen lasers in the ice, but these lasers cannot produce sharp pulses, so only the total attenuation length, combining the effects of scattering and absorption, can be measured with this device. We can send sharp laser pulses generated at the surface down optical fibers to isotropizers in the ice, allowing us to separate these two effects. However, attenuation in the fibers limits us to longer wavelengths, so that some extrapolations must be made to the wavelengths of interest. More recent deployments have included blue LEDs inside the modules, so that further studies may produce better measurements. In the meantime, simulations with different models of the bulk ice (including the layered structure of Fig. 4.2 vs. using a homogeneous model) change the predicted rates by about 15% at the final cut levels [74].

More difficult to measure are the optical properties of the ice melted during the deployment process, which refreezes around the optical modules over the period of a week following deployment. This quick refreezing is very different from the long adiabatic formation of the bulk ice, and leads to the inclusion of air bubbles in the immediate vicinity of the modules. A television camera deployed with a string in 1997–98 indicated very strong scattering, although the issue was clouded by the possibility that the equipment simply failed. The nitrogen laser and isotropizers do not have good lines-of-sight through the bubbly ice to nearby modules, although the blue LEDs will. We therefore have no direct measurements of the optical properties of the hole ice. Recent studies of OM sensitivity using the muon events [86] have led to a hole ice model with much stronger scattering immediately below the OMs than was previously assumed, and this has had a strong effect on the predicted rate of neutrino events. Because the effects are very strong, and because the measurements of the hole ice are highly uncertain, the uncertainty in the overall neutrino passing rate due to the hole ice is a factor of two.

Another uncertainty which affects the predicted neutrino rates is the simulation of the thresholds of the surface electronics. Improvements in the description of the amplifiers and discriminators have produced strong variations in the number of triggers produced by neutrino events, because with the steeply falling energy spectrum small changes in the threshold behavior will have large effects, changing the trigger rates by as much as 50%. However, the events at trigger threshold are very difficult to reconstruct because of their low energy, and so almost all of the variation disappears

when the simulated events are filtered and reconstructed. This systematic is therefore important for calculating the trigger rate, but does not extend through the whole analysis.

Finally, a bug has been discovered in the program used to simulate photon propagation in the ice [68, 87]. This affects the muon light yield directly, of course, causing a slight (few percent) reduction in photons close to the track and significant overprediction for the number of scattered photons. For the minimum-ionizing muons characteristic of atmospheric neutrinos, the overall effect should be small, because most light is detected within a few tens of meters of the track. However, because this code was also used to interpret the measurements of ice optical properties, those measurements are also called into question. The overall effect of this error is not known at the time of this writing, although it is not expected to be larger than the other effects mentioned above.

7.3 Event Quality

The fundamental problem in demonstrating the observation of neutrinos is, of course, evaluating the reduced data set. There are basically three tools we can use to determine how pure the data are: simulations of cosmic ray muons, simulations of neutrinos, and the event viewer. One would like to use the Monte Carlo simulations of neutrinos and downgoing muons to make some quantitative measurement of the purity of the data sample. Unfortunately, the simulations are not of sufficient accuracy to be accepted *prima facie*. The atmospheric neutrino flux which is input to the signal Monte Carlo is itself uncertain at the level of 30% or greater [41] — cf. Fig. 2.3, where the uncertainties on the fluxes of cosmic ray primaries are quite large above about 100 GeV — not to mention the effects described in Sec. 7.2. One cannot therefore simply compare the numbers of data observed to neutrinos expected. The background Monte Carlo is even less precise; unlike the neutrino Monte Carlo, the cosmic ray simulation tries to describe the bizarre muons, the one muon in a million that appears to be something very different than it truly is. This is intrinsically a far more difficult problem, even leaving aside the systematics of Sec. 7.2. The background Monte Carlo also faces additional systematics: different air shower Monte Carlos yield very different rates of high energy muons, and the many technical details, such as cross talk, that are not included in the detector simulation are seen to have a much larger relative effect on the number of misreconstructions than on the signal efficiency. Furthermore, the surface layer of the ice cap (the firm layer of partially

compacted granular snow in an intermediate stage between snow and ice) has a much lower density than the ice itself; it has recently been found [65] that including the density profile of the firn layer can raise the background rates by as much as 40%, depending on the angle of the muon.

One wishes, then, to have some robust method based on the available simulations for estimating the purity of the data sample, which is however not dependent on the normalization of the neutrino Monte Carlo or on the ability of the background Monte Carlo to predict all of the bizarre ways in which a downgoing muon can be made to look like an upgoing one, and which is flexible enough to allow for the many systematic uncertainties in the simulations. To this end we have developed the concept of “event quality,” which combines, in a natural fashion, the information from the various parameters available into a single number.

Consider the six-dimensional space formed by the six cut parameters described in Section 7.1. Any given event can be described as a point in the space, with coordinates

$$\vec{q} = \left(N_{\text{dir}}^{[-15:75]} / N_{\text{ch}}, |S_{P_{\text{hit}}}|, L_{\text{dir}}^{[-15:25]}, I_1 / \sum I, \Delta\theta, L \right).$$

However, we would like to make these axes more symmetric; with the intuitive choice of units (meters, degrees, percentages, and so forth) the distribution of events in the various coordinates is very irregular, and worse, different dimensions increase in event quality in different directions — high numbers are indicative of high quality on some axes, but on others low values are better.

The cumulatives of the distributions provide a natural scale for the parameters. There are *a priori* three sets of cumulatives we could use: the data, the simulated signal, and the simulated background. We ought not use the data, because we would not know what fraction of the distribution was due to signal and what to background — that determination is the point of the entire exercise. The background simulation, as discussed above, is the least reliable of the three sets of events, as well as being the most limited in statistics. The signal Monte Carlo is not perfect, but we have more faith in it than in the downgoing simulation. We thus scale each axis of our space by the cumulative of the simulated atmospheric neutrino signal; the units are percentages of simulated neutrino events rejected. The origin, $\vec{q} = (0, 0, 0, 0, 0, 0)$, is thus no cuts, and the opposite corner of the unit hypercube, $\vec{q} = (1, 1, 1, 1, 1, 1)$, corresponds to cuts on each parameter strong enough to reject all simulated signal. The point $\vec{q} = (0.5, 0.75, 0, 0, 0, 0)$ would correspond to a cut on $N_{\text{dir}}^{[-15:75]} / N_{\text{ch}}$

which by itself would remove half the simulated atmospheric neutrinos, a cut on $|S_{P_{\text{hit}}}|$ which alone would remove three quarters of the signal, and no cuts on the other four parameters.

With the parameters scaled by the inverted cumulatives, the atmospheric neutrino Monte Carlo is distributed fairly uniformly in the space. The distribution would be perfectly uniform if the parameters were completely independent, but of course they are not; *e.g.*, events with large numbers of direct hits tend to have both high percentages of direct hits and long projected lengths L_{dir} . Perfect uniformity is not required, however, as long as the signal can be easily distinguished from the steeply falling background.

The space itself has too many dimensions to be examined easily. We therefore project the events onto a single dimension for clarity. The choice of the dimension, which is a line in the full space, is arbitrary, but we should pick a direction in which the background falls off quickly so that the neutrino events can be examined over a wide range of event qualities. This desideratum is equivalent to choosing a line through a good set of cuts, which reject the background while preserving most of the signal. In choosing a set of cuts, we have not attempted to optimize signal-to-noise in a quantitative way — our inability to do so is the reason for this exercise in the first place — but have rather tightened cuts gradually and intuitively in an attempt to reject events which were obvious fakes.

Having chosen a set of cuts, we draw a line from the origin, through the point corresponding to the cuts, to the edge of the hypercube. We then divide the cut space into nested rectangular shells of equal thickness with one vertex lying on the line, extending to the boundary of the space as shown in Fig. 7.9. Each shell represents a progressively higher level of cuts, and the events within the shell are said to be of a certain quality Q determined by the shell in which they are found. For clarity we have used not the natural scale of the unit hypercube in measuring the quality, but have rather defined the quality such that the set of events chosen by hand corresponds to a quality of 10; the highest quality, bounded by the border of the unit hypercube, has a value of 24 in these units. Details of these quality levels, including numerical values of the cut parameters and the number of events observed, are given in Appendix D.

With the events thus projected into a single dimension, we can examine the distribution of




Figure 7.9: An illustration in two dimensions of the division of cut space into shells of increasing quality. The actual cut space is six-dimensional, and the axes are rescaled according to the cumulative distributions of predicted atmospheric neutrino events.




Figure 7.10: The number of events passing various levels of cuts $N_{events}(\geq Q)$. The data are seen to turn over from rough agreement with the predicted background to relatively good agreement with the atmospheric neutrino prediction at higher Q .

the various sets of events — data, background Monte Carlo, and neutrino Monte Carlo — shown in Fig. 7.10. As expected, the simulated backgrounds are clustered at low quality, and fall steeply. The signal Monte Carlo are distributed more uniformly, although the reduced phase space at higher qualities leads to a slow fall in the neutrinos as well. The data begin by falling steeply with the backgrounds, but then level out and track the signal Monte Carlo to higher qualities.

By examining the ratio of the number of events observed to the number predicted, shown in Fig. 7.11, we can test the accuracy of the Monte Carlos. At low qualities, we know that the data set is dominated by backgrounds, and so a ratio far from unity indicates a poor description of the fakes. At high qualities we are similarly sensitive to the accuracy of the neutrino simulation. There




Figure 7.11: The ratio of the number of events observed to the number predicted by simulations of atmospheric neutrinos and downgoing muons.

is in both cases an uncertainty in the overall normalization, but if the simulations are qualitatively accurate then we should have a flat ratio, even if misnormalization causes the ratio to converge to a number slightly different from unity.

Note that we are comparing the number of events *at* each given quality (within finite bins, of course). This is fundamentally different, and more informative, than comparing the number of events passing progressively tighter sets of cuts. We are looking at the number of events within a narrow region ΔQ , rather than integrating the number of events above some fixed set of cuts. In the latter case events at many different qualities are convolved, and an excess at one quality can conceal a deficit at another. We are also fixing the relative contribution of each cut by following a straight line in the cut space. If cuts are applied sequentially rather than tightened simultaneously, it is possible

to hide disagreements between data and Monte Carlo by the ordering of the cuts; a cut that removes unsimulated instrumental effects which are dominant at high cut levels can conceal those effects if it is (retroactively) applied before the unsimulated events become the dominant component of the background; this procedure gives the appearance of better understanding of the data than is in fact the case.

We observe a number of features in Fig. 7.11. At very low qualities, below about $Q = 3$, we see a significant excess in the data, compared to the prediction from the Monte Carlo (primarily downgoing muons, with a small number of atmospheric neutrinos). This excess could be due to several factors. The excess is not huge, and could simply be due to uncertainties in the prediction of penetrating muons. However, the fact that the excess falls more quickly than the downgoing muons probably indicates that there is residual contamination from instrumental effects, or possibly could be related to the errors in the muon propagator via the angular distribution of muons that produce fakes. This hypothesis is supported by Fig. 7.12, which shows the distribution in depth of the events in the first bin of Fig. 7.11. Without the specialized cross talk cleaning algorithms, there were 45,000 unsimulated fakes, not 2700, so the relatively small excess observed at low Q in Fig. 7.11 would indicate better than 90% efficiency on the part of the cross talk cleaning algorithms. It would be surprising if the algorithms were more effective, since in the absence of a detailed cross talk map the cleaning algorithms were developed in a rather intuitive fashion.

Once the background has fallen away, we observe the ratio to settle to a value of approximately 0.6, over a broad range from about $Q = 7$ to $Q = 21$. At the very end, the ratio rises again; there are too many extremely good events. It is of course possible, in a logical sense, that there are downgoing muons or instrumental effects that produce events which look more like upgoing muons than do upgoing muons, in every quality parameter simultaneously, and also in the event viewer [88]. However, the more likely culprits are the systematics discussed in Section 7.2; it would be surprising if the qualitative description were perfect in the presence of such large uncertainties. It is hoped that reductions in these systematics will lead to a flatter Q curve; this requirement could be used to help evaluate future versions of the Monte Carlos. Two suggestions already investigated, changing the PMT quantum efficiency and angular response function, do not improve the shape.




Figure 7.12: The depth of the center of gravity of the events at $Q = 0$. The two peaks at the bottom and center of the detector are associated with channels known to cross talk extensively, an effect not included in the detector simulation. The peak at the top of the detector is caused by nearly horizontal muons; the discrepancy in magnitude could also be explained by systematics in the muon propagator or by the fact that the atmospheric muon prediction contains only proton air showers.




Figure 7.13: The number of channels hit per event. The Monte Carlo systematically overpredicts the number of modules reporting.




Figure 7.14: The observed zenith angle distribution of events passing the $Q = 7$ cuts; -1 is vertically upgoing and 0 is the horizon.

The fact that the ratio is flat above $Q = 7$ gives us confidence in the description of events passing the corresponding set of cuts. As shown in Figs. 7.1–7.8, the agreement of the data with simulated neutrinos in the cut parameters is reasonably good, though the systematics mentioned in Sec. 7.2 prevent a perfect agreement. Other characteristic observables plotted in Figs. 7.13–7.16, such as the number of channels firing, the zenith angle distribution, and the position of events in the detector, also show good agreement with the predictions for atmospheric neutrinos. Some questions have been raised as to whether structure is visible in Fig. 7.14; a Kolmogorov-Smirnov test of the consistency of the data and the Monte Carlo prediction gives a probability of 4.4% that the sets are drawn from the same parent distribution. As discussed in Section 7.5, contamination of the remaining data by misreconstructed downgoing muons is low.

7.4 Efficiencies and Effective Area

Effective area is not a well defined concept for a detector whose thickness is comparable to the range of the particles it detects. While high energy muons have ranges of many kilometers, the atmospheric neutrinos that are the goal of this analysis typically have much shorter ranges. In this case, muons have a reasonable probability of ranging out in the middle of the detector, or conversely

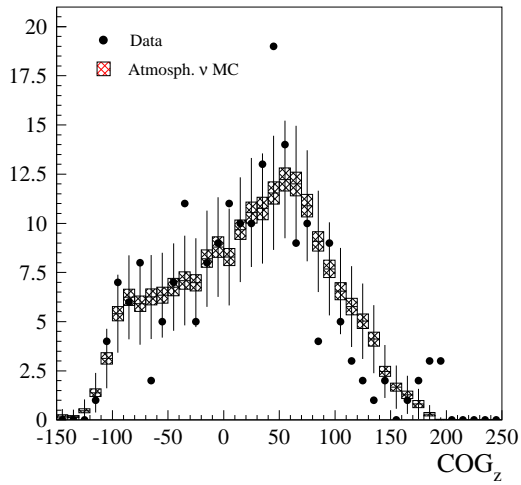


Figure 7.15: The depth of the center of gravity of the event. Some contamination from nearly horizontal muons or muon bundles at the top of the detector may be visible.



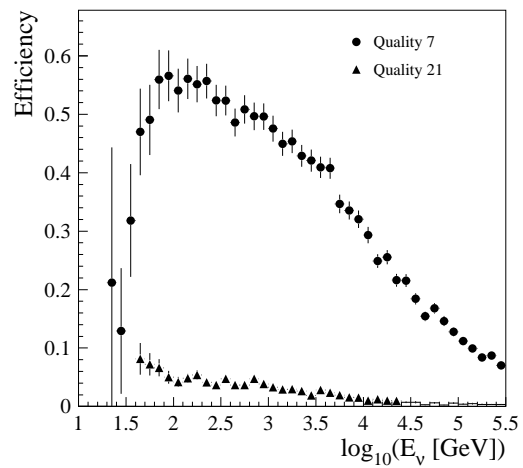


Figure 7.16: The radial coordinate of the center of gravity of the event.




(a) Efficiency for triggered events.




(b) Efficiency for upward-reconstructed tracks.

Figure 7.17: The fraction of neutrino events triggered which are identified with this analysis, for two different sets of cuts. The efficiency falls off sharply above a few TeV, primarily because of the cut on $N_{\text{dir}}^{[-15:75]}/N_{\text{ch}}$.



(a) Efficiency for triggered events.



(b) Efficiency for upward-reconstructed tracks.

Figure 7.18: The fraction of neutrino events triggered which are identified with this analysis, for two different sets of cuts. The low efficiency at the horizon is due to two factors: nearly-horizontal atmospheric muons form a large background, and the narrow detector geometry of the detector makes reconstruction difficult. Note that the efficiency of the cuts is nearly flat — the low efficiency is due to the Bayesian reconstruction being unable to guarantee effective rejection of nearly horizontal atmospheric muons. The full AMANDA-II detector will have much flatter angular response.

of being produced within the detector, in which case the light produced at the hadronic vertex will have a large impact on the detector response. The proper measure of the detector's size is thus effective *volume*, not area.

Nevertheless, most of the literature deals with effective areas, and it is useful to somehow convert the effective volume to an effective area for the purpose of comparison. One approach is to divide the effective volume by the range of the muon, which is correct for high energies but ignores the effects mentioned above. An alternative, given that we are interested in effective area mainly for comparison to theoretical predictions, is to directly compare the detector's actual performance to that of an ideal detector of given effective area; in this way the approximations are in some sense cancelled out of the final result [89]. With this definition, this analysis produces an effective area, integrated over the atmospheric neutrino spectrum, of approximately 7000 m² in the vertical direction (above about $\cos \theta = 0.8$), falling to only a few hundred m² near the horizon.

It is important to note that though this analysis was optimized for atmospheric neutrinos, the detector itself was not. Even for this analysis, the efficiency peaks at around 1 TeV (see Fig. 7.17(a)); the yield of 100 GeV neutrinos is significantly worse even though the cuts are most efficient for 100 GeV (Fig. 7.17(b)). The effective area integrated over a harder spectrum, like that of the astrophysical sources AMANDA is designed to detect, is much higher than that for atmospheric neutrinos. Furthermore, the completed AMANDA-II detector now taking data will have significantly flatter angular response — as shown in Fig. 7.18(a), the B10 detector has great difficulty in reconstructing muon tracks reliably unless they are rather steeply vertical. The sensitivity of the detector to astrophysical neutrinos can thus be expected to increase dramatically over the level of the analysis in Chap. 8 with an analysis of the 2000–2001 data optimized for higher energies.

7.5 Background Estimation

The central question of the analysis is to determine the purity of the data set for our chosen set of cuts. Based on Fig. 7.11, we believe that our simulations provide a reasonably accurate description of the data, but there are systematic effects which prevent us from treating the simulation as a quantitatively precise prediction. The number of events predicted from atmospheric neutrino Monte Carlo is nearly a factor of two above that actually observed, and because the background simulation

is intrinsically a far more difficult problem we ought not take that prediction as more precise than the signal Monte Carlo. Additionally, the background prediction of 11.7 ± 4.1 events is statistically limited; despite the simulation of a quarter trillion protons incident on the atmosphere, the prediction is based on only 68 events of various weights that pass the $Q = 7$ cuts.

Another tool, the event viewer, is obviously far less quantitative than the simulations, but it is also more robust: the human mind is far more capable of recognizing unexpected problems than are the AMANDA analysis programs. The difference between good upgoing tracks, like the ones shown in Figs. 7.19–7.21, and misreconstructions like those in Figs. 6.2 and 6.3 is clear to the human eye if not to the reconstruction software. Examination of the data set has led to the discovery of many strange classes of backgrounds produced by instrumental effects not included in the detector Monte Carlo, after which specific cuts or cleaning routines were designed to remove or at least reduce these classes of fakes. Examining the data also led to the development of more general cut parameters and, in the absence of a precise background Monte Carlo, guided the optimization of the cuts used in this work: obvious fakes were identified and cuts were selected to remove those fakes. This procedure clearly raises the danger of circular reasoning; all the recognizable fakes may have been removed, but there is no guarantee that all the fakes were recognized. Inspection of the events passing the cuts does indicate that the data remaining are of very high quality, but nine events are clear fakes, and another seven are unconvincing. This method is of course highly subjective, and because the cuts were designed precisely to remove subjectively unconvincing events one must expect that this method will tend to produce an undercount. Still, even a lower limit is informative, and even if we do not believe the precise count we may gain some estimate at least of the number of events involved.

A final approach is to use Figure 7.11 as a measurement of the misnormalization of the atmospheric neutrino Monte Carlo. The fact that the data and the simulated atmospheric neutrinos are distributed almost identically over a wide range of qualities is compelling evidence for the essential accuracy of the simulation. The Monte Carlo may not predict the absolute rate of events, but on the other hand we should not expect perfection with the large systematic uncertainties; overall, most of the events which are observed are described rather well. We can thus renormalize the atmospheric neutrino Monte Carlo to force the observed ratio to unity; i.e., we reduce the overall prediction of the

No external geometry file is opened.
Detector: amanda-b-10, 10 strings, 302 modules
Data file: /data/disk6/bayes/data/ab_801_199-97-0.ft420.combined.best.noise_xtalk.best.showet.cln.f2k
File contains 2456 events.
Displaying data event 296959 from run 0
Recorded y/d/y: 1997/293
5163.1811435 seconds past midnight.
Before cuts: 19 hits, 19 OMs
After cuts: 19 hits, 19 OMs




Figure 7.19: An upgoing muon in AMANDA-B10. The line indicates the reconstructed track.

No external geometry file is opened.
 Detector: amanda-b-10, 10 strings, 302 modules
 Data file: /data/disk6/bayes/data/ab_801_199-97-0.ft420.combined.best.noise_xtalk.best.shower.cln.f2k
 File contains 2456 events.
 Displaying data event 1197960 from run 0
 Recorded y/dy: 1997/285
 18 [32.009138] seconds past midnight.
 Before cuts: 41 hits, 41 OMs
 After cuts: 41 hits, 41 OMs




Figure 7.20: A higher energy muon in the detector, probably around a TeV. The amount of light emitted is clearly greater than in Fig. 7.19. Energy reconstruction based on this idea is presently under development.


```

No external geometry file is opened.
Detector: amanda-b-10, 10 strings, 302 modules
Data file: /data/disk6/bayes/data/ab_801_199-97-0.fit420.combined.best.noise_xtalk.best.shower.cln.f2k
File contains 2456 events.
Displaying data event 3114804 from run 0
Recorded y/dy: 1997/226
15661.2396395 seconds past midnight.
Before cuts: 19 hits, 19 OMs
After cuts: 19 hits, 19 OMs

```




Figure 7.21: A shallower upgoing muon. These events are more difficult to reconstruct, because the tracks are shorter and there is some background from nearly horizontal atmospheric muons. AMANDA-II and Ice Cube will have significantly higher efficiency at these angles due to the longer distance over which near-horizontal muons can be tracked.

number of neutrino events by $1.00 - 0.58 = 42\%$. With this renormalization, we have a prediction of 162 events from atmospheric neutrinos. We observe 204 events passing the cuts, so subtraction would imply that there are 42 fakes in the data set. This method is certainly inclusive, in the sense that anything not expected based on the renormalized signal Monte Carlo is considered to be a fake. However, we note that 22 of the excess events are found above $Q = 21$, on the far right of Figs. 7.10 and 7.11. If we accept that these events, predicted or not, are real neutrino events, then our count of the number of fakes becomes 20. It should be pointed out that this total is at best rather fuzzy; the 1σ uncertainty in the linear fit itself gives an error of 4.7% in the renormalization, which translates to ten events, even without discussing systematic effects.

All three of our methods of estimating the background suffer from large uncertainties and obvious limitations. Nevertheless, the fact that they all give approximately the same answer gives us some confidence in that answer. From the downgoing Monte Carlo there is a prediction of 11.7 fakes, with a caveat that there is probably a residual contribution from unsimulated events. From the hand count we have a tally of 16 questionable events, and from attempting to renormalize the signal prediction we have an estimate of 20 fakes. The estimate that background contamination is probably about 10% of the data set is thus rather robust, even if it is not overly precise.

With a better understanding of the quirks of the detector and of the ice, perhaps the Monte Carlos can be made more accurate, and better measurements can be made. Moreover, the larger detector used in later years will be much more effective at background rejection, and cleaner measurements will be possible. In the meantime, it is clear that the bulk of the data set does consist of neutrinos. While the overall effective volume of the detector may be uncertain to a factor of two, it is nevertheless also known to a factor of two, and this establishes AMANDA as the largest neutrino telescope in operation. With this caveat that the detector is still not precisely understood, it is yet possible to begin to use the detector to address the more exciting questions for which it was designed.

Chapter 8

Flux from Point Sources

The analysis presented in Chapter 7 is designed to detect the soft spectrum of atmospheric neutrinos. Most models of extragalactic neutrino emitters predict a considerably harder spectrum, typically E_ν^{-2} , and a separate analysis optimized to detect such sources has been undertaken [90]. Nevertheless, even for such a spectrum the most probable neutrino energy detected is only some 10–50 TeV. Furthermore, high energy muons from high energy neutrinos will deposit most of their energy very soon after production; most of the muon’s range is traversed at relatively low energies, for which the present analysis is suitable. This analysis thus retains sufficient sensitivity to place competitive limits on neutrino emission from candidate sources.

8.1 Candidate Sources

Rather than perform a full-sky search for point sources, with its attendant statistical complexity, we will look for emission from known sources of VHE gamma rays. With this concrete set of candidates the analysis is much simplified: we can avoid questions of binned vs. unbinned searches, the statistical penalties involved in searching large numbers of points, and so forth, and simply place limits on the neutrino flux coming from fixed directions.

The gamma ray telescopes now in operation have detected several sources of VHE ($\gtrsim 100$ GeV) emission. Six extragalactic sources are known, listed in Table 8.1. One of these sources is in the Southern sky, rendering it inaccessible to AMANDA at these energies, but the rest are well within AMANDA’s field of view, as shown in Fig. 8.1.




Figure 8.1: AMANDA events plotted in galactic coordinates along with the five candidate sources. The heavy line indicates the horizon.

Extragalactic VHE Sources				
Source	Rt. Asc. [h]	Decl. [°]	Gal. Long.	Gal. Lat.
Markarian 421	11.074	38.209	179.832	65.032
Markarian 501	16.898	39.760	63.600	38.859
1ES1959+650	20.000	65.149	98.003	17.670
1ES2344+514	23.785	51.705	112.891	-9.908
3C66A	2.378	43.036	140.143	-16.767
PKS2155-304	21.981	-30.226	17.730	-52.246

Table 8.1: Taken from Weekes [16].

8.2 Optimization and Background Estimation

In searching for emission from point sources, as opposed to the diffuse flux of atmospheric neutrinos, one is more willing to allow moderate levels of background due to misreconstructed down-going muons. Positional clustering is the signature used to identify sources from the backgrounds of misreconstructed events and of correctly reconstructed atmospheric neutrinos; since the backgrounds should be randomly distributed (in right ascension, if not in declination or even in azimuth) they are less troublesome than in the diffuse analysis. To determine the optimal cut level for the point source search, we plot in Fig. 8.2 the sensitivity of the search as a function of the minimal event quality, for an E^{-2} source at a declination of 40° . The sensitivity is the usual ratio of the number of signal events predicted per bin to the square root of the number of background events at that declination, under the approximation that all events in the data are background. The optimal sensitivity is achieved by requiring event quality $Q \geq 5$, as compared to the $Q \geq 7$ used for the atmospheric analysis.

In finding this optimum, two effects should be noted. First, as the cuts are tightened, the angular resolution improves and so the optimal bin size shrinks. Second, the decreasing background levels cause the optimal bin size to grow, as more and more signal can be retained for comparable amounts of noise. Both of these effects are small in the region of interest, and furthermore they tend to cancel each other out, so they have been neglected in finding the optimal set of cuts, although they are taken into account when determining the bin size used.

With our cuts set, we can calculate the angular resolution and optimal bin size. The RMS angular resolution in zenith and in azimuth¹ is shown for various declinations in Fig. 8.3. From these

¹Actually, $\phi \cos \delta$, not ϕ .




Figure 8.2: Sensitivity as a function of minimum event quality for a point source at a declination of 40° , with a spectrum of E_ν^{-2} . Plotted is the ratio of predicted signal to the root of the number of events observed, in arbitrary units.




Figure 8.3: The RMS angular resolution for an E^{-2} source. The lines are analytic parametrizations used to choose the bin size; they smooth out statistical fluctuations produced by the low population tails of the zenith angle distributions.

resolutions we find the number of background events expected in a $\pm 1\sigma$ elliptical search bin at the declinations of our source candidates. Alexandreas *et al.* [91] give the optimal bin size, as a function of this background expectation.

With the cuts and search bins set, we can look for clustering in the skymap around our candidate sources. Fig. 8.4 shows the positions of the 286 events passing the $Q \geq 5$ cuts, as well as the locations of the five sources and the search bins used in this analysis, in equatorial (J2000) coordinates on a Hammer-Aitoff projection.

We estimate the background expectation by examining off-source bins in the same declination band as the source. Because of AMANDA's unique location, the source appears at constant zenith angle, and time variations in sensitivity due to the position of the source are essentially nil². By using off-source bins in the same zenith band, any detector-related variations in sensitivity are properly accounted for. The declination band is defined by the limits of the source bin; no correction is made for the overlap of one source's off-source band with other source bins, because the statistical uncertainty of

²There is in fact a slight three-fold modulation in sensitivity with azimuth due to the geometry of the detector, but this will be averaged out unless emission comes in a flare of less than eight hours' duration.




Figure 8.4: The 286 events with $Q \geq 5$ plotted in equatorial coordinates along with the five candidate sources. The ellipses around the points are the search bins used in the analysis.

the background estimation is already a dominant factor in the calculation of significance, the expected signals are below the level of detectability, and any contamination of the off-source bands will give a conservative result by artificially elevating the background level. The expected background in the search bin is found by simply taking the ratio of the solid angle of the search bin to the solid angle in the off-source band, and counting the number of events in the off-source band. The declination bands are shown in Fig. 8.5.


8.3 Results

There are two calculations we can make for these candidate sources. First, we can calculate the statistical significance of the possible detections, using the likelihood-ratio test of Li & Ma [92]. This result is the probability that the observation indicates the presence of a signal, rather than pure background, and accounts properly for the statistical uncertainties in the background estimation. As shown in Table 8.2, there is no evidence for neutrino emission from those sources at the levels accessible to this analysis. The significances given in Table 8.2 do not include statistical penalties for the number of trials, which include not only the five sources tested here but also the fact that this data set has been examined repeatedly for evidence of point sources [80, 90, 93].


From the background estimates and the numbers of events in the source bins, we can also find 90% C.L. limits on the number of events from the candidate sources, according to the method of Feldman & Cousins [94]. The results are also given in Table 8.2. The Feldman-Cousins method does not account for statistical uncertainties in the background estimation (or for systematic uncertainties in the sensitivity — doing so in a strictly classical manner is not trivial). Following [95] and [96], we have used a Bayesian approach to account for this uncertainty, as well as the (much larger) systematic uncertainty in the detector’s sensitivity. We have constructed the confidence intervals following Feldman and Cousins’ prescription, except that the Poisson probability $P(n | \mu + b)$ of observing n events given known signal μ and background b has been replaced by the integral

$$\iint d\mu' db' P(n | \mu' + b') P(\mu') P(b').$$


The background level and detector sensitivity were assumed to be normally distributed about the central values. The width of the background distribution was taken from the number of counts in the



(a) The declination band containing 1ES1959.



(b) The declination band containing 1ES2344.



(c) The declination bands containing Mrk 421, Mrk 501, and 3C66A.

Figure 8.5: The declination bands used to estimate the backgrounds. The candidate sources are shown as diamonds inside ellipses representing the search bins. The different lengths of the bands indicate the reduced solid angle at high declinations.

Source	Bin Radius		N_{off}	N_{on}	N_{BG}	$S_{\text{Li-Ma}}$	90% U. L. [events]
	Rt. Asc.	Declin.					
Markarian 421	0.518	4.855	33	1	1.12	-0.11	3.24
Markarian 501	0.535	4.870	32	3	1.12	1.43	6.30
3C66A	0.574	4.867	39	1	1.47	-0.40	2.93
1ES2344	0.712	4.694	50	0	2.33	-2.14	1.21
1ES1959	1.092	4.123	25	2	1.79	0.15	4.12

Table 8.2: Results of the search for neutrino emission from candidate sources.

declination band of the source, and for the sensitivity the difference in the predictions of the standard Monte Carlo and that with the muon propagator of Lipari [85] and the OM angular sensitivity of [86] (the highest and lowest predictions, differing by approximately a factor of two) was taken as the 2σ level.

Having produced these limits on the number of events produced by our candidate sources, we must convert these event limits to limits on the neutrino flux. This process depends on the assumed spectrum of the source; as above when calculating the sensitivity we will assume an E_ν^{-2} spectrum. We find the normalization by simulating E_ν^{-2} sources with flux

$$E_\nu^2 \Phi_{\nu+\bar{\nu}} = 10^{-6} \text{ cm}^{-2} \text{ s}^{-1} \text{ GeV}$$

at the declinations of the various candidates, and dividing the event limits of Table 8.2 by the predictions; the same procedure is followed for an assumed E_ν^{-3} spectrum. The results are given in Table 8.3. The energy ranges for which the limits are valid, defined by the range of neutrino energies that contributes 90% of the predicted signal, are about 350 GeV – 175 TeV for the assumed E_ν^{-2} source, and 60 GeV – 6.3 TeV for E_ν^{-3} emission. The neutrino spectra from the Monte Carlo are shown in Fig. 8.6. The limits are calculated for steady-state emission over the period April 6 to November 14, 1997.

We have also included in Table 8.3 the experimental sensitivity, defined following [94] as the average upper limit that would be expected from this analysis in the absence of true sources; in some instances the limits outpace the sensitivity, and in these cases the sensitivity should be taken as the bound on the flux. As with the upper limits, the sensitivities have been calculated with uncertainties in background and sensitivity treated following [95] and [96].

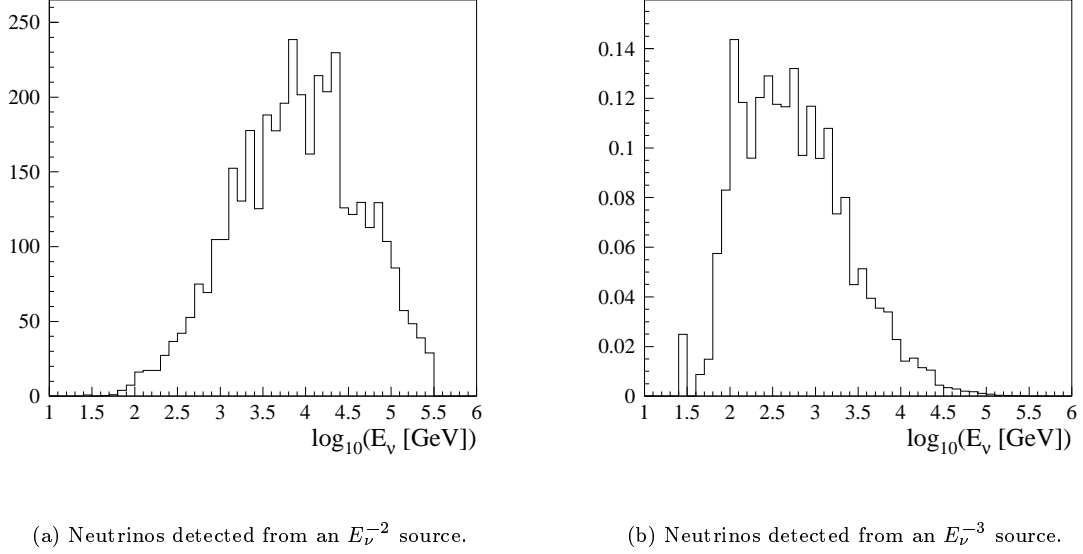


Figure 8.6: The energy of neutrinos identified by this analysis for assumed E_ν^{-2} and E_ν^{-3} spectra of arbitrary normalization. The cutoff at $10^{5.5}$ GeV is due to the limited range of the simulation; the tail is relatively small and makes the quoted limits slightly conservative.

Source	E_ν^{-2} spectrum		E_ν^{-3} spectrum	
	90% U. L.	Sensitivity	90% U. L.	Sensitivity
Markarian 421	$1.22 \cdot 10^{-6}$	$1.32 \cdot 10^{-6}$	$2.82 \cdot 10^{-3}$	$3.01 \cdot 10^{-3}$
Markarian 501	$2.14 \cdot 10^{-6}$	$1.12 \cdot 10^{-6}$	$4.96 \cdot 10^{-3}$	$2.63 \cdot 10^{-3}$
3C66A	$1.01 \cdot 10^{-6}$	$1.24 \cdot 10^{-6}$	$2.21 \cdot 10^{-3}$	$2.77 \cdot 10^{-3}$
1ES2344	$0.23 \cdot 10^{-6}$	$1.08 \cdot 10^{-6}$	$0.59 \cdot 10^{-3}$	$2.24 \cdot 10^{-3}$
1ES1959	$1.08 \cdot 10^{-6}$	$1.01 \cdot 10^{-6}$	$1.94 \cdot 10^{-3}$	$1.84 \cdot 10^{-3}$

Table 8.3: Sensitivities and flux limits on the candidate sources. The limits are on $(\Phi_\nu + \Phi_{\bar{\nu}})E^n$, where n is the spectral index of the assumed spectrum, in units of $\text{cm}^{-2} \text{sec}^{-1} \text{GeV}^{n-1}$.

Finally, we should note the appearance in Fig. 8.5(c) of a cluster of events near the position of 3C66A. If an analysis identical to that for 3C66A were performed with the search bin centered at right ascension 1.2 h, rather than 2.378 h, the search bin would contain 7 events, with an expected background from the declination band of 1.24 ± 0.22 events. This corresponds to a significance, according to Li & Ma, of approximately 3.4σ (a chance probability of $3 \cdot 10^{-4}$), and a most-probable flux of $E^2\Phi_{\nu+\bar{\nu}} = 1.91 \cdot 10^{-6} \text{cm}^{-2}\text{s}^{-1}\text{GeV}$. However, the number-of-trials factor is quite complicated in this case, because the search bin was placed by hand on the skyplot, and so the Li-Ma significance must be reduced somewhat. Naïvely, one could say that the search bin took up about 1/214 of the upgoing hemisphere, so that one must take a penalty for 214 trials; this brings the chance probability up to about 7%. But the bin was not the result of 214 random trials; the bin was placed by hand to maximize the signal, and so this treatment will overestimate the significance of the result. One must conclude, therefore, that there is a reasonably large probability that the apparent excess is simply a fluctuation of the background, and that the question cannot be settled with this data set. The additional data from 1998–2001 should certainly be enough to resolve the issue.

Chapter 9

Conclusions

The AMANDA detector has been in operation since 1996. Results from the first year of operation, with only four strings in place, were promising but failed to demonstrate conclusively that the detector was a working neutrino telescope. Improvements in reconstruction methods and in data analysis, applied to the data set recorded in 1997 with ten strings in operation, have resulted in the positive identification of upgoing neutrino events.

The AMANDA detector was optimized for the detection of TeV events; although some sensitivity is retained even below 100 GeV, the detector geometry was designed to minimize the observation of atmospheric neutrinos, which are an unwanted foreground to astrophysical sources. Moreover, the narrow geometry of the ten strings in place in 1997 greatly reduces sensitivity to horizontal muons. Nevertheless, the ten-string detector has been established as a working instrument of several thousand m² effective area for atmospheric neutrinos, observing neutrinos at a rate of more than 1.4 per day of exposure, with background rates below 10%. The response of the detector, while not yet perfectly modelled, is understood to within a factor of two, and the expanded detector now taking data will have a greatly increased effective volume and angular range.

With this demonstration of the proper operation of the detector, we can begin to place limits on astrophysical sources of neutrinos. Energetic extragalactic objects known from their gamma ray emission have been shown not to emit neutrinos significantly above the level of $10^{-6} E_{\nu}^{-2} \text{ cm}^{-2} \text{ s}^{-1} \text{ GeV}^{-1}$. With improvements in analysis techniques, optimization for high energy events, and the additional data already recorded with the expanded detector, the sensitivity of the detector will be greatly enhanced and the possibility of observing high energy extraterrestrial neutrinos, an elusive

goal for over 40 years, may finally become a reality.

Bibliography

- [1] R. Svensson and A. A. Zdziarski. *Astrophys. J.*, 349 (1990) 415–28.
- [2] J. G. Learned and K. Mannheim. *Ann. Rev. Nucl. Science*, 50 (2000) 679–749.
- [3] K. Greisen. *Ann. Rev. Nucl. Science*, 10 (1960) 63.
- [4] F. Reines. *Ann. Rev. Nucl. Science*, 10 (1960) 1.
- [5] M. A. Markov and I. M. Zheleznykh. *Nucl. Phys.*, 27 (1961) 385.
- [6] M. A. Markov, in *Proc. Intl. Conf. High Energy Phys. at Rochester*, E. C. G. Sudarshan, J. H. Tinlot, and A. C. Melissinos, eds. (1960) p. 578.
- [7] J. Babson *et al.* (the DUMAND Collaboration). *Phys. Rev. D*, 42 (1990) 3613–20.
- [8] V. A. Balkanov *et al.* (the Baikal Collaboration). *Nucl. Phys. B (Proc. Suppl.)*, 91 (2001) 438–44.
- [9] The ANTARES Collaboration. *A deep sea telescope for high energy neutrinos.* astro-ph/9907432 (July 1999).
- [10] L. K. Resvanis *et al.* (the NESTOR Collaboration). In *Proc. Honolulu Workshop on High Energy Neutrino Astrophys.*, V. J. Stenger, J. G. Learned, S. Pakvasa, and X. Tata, eds. (World Scientific, Singapore) (March 1992) pp. 325–53.
- [11] C. de Marzo for the NEMO Collaboration. *Nucl. Phys. B (Proc. Suppl.)*, 87 (2000) 33–5.
- [12] T. Gaisser. *Cosmic Rays and Particle Physics* (Cambridge University Press) (1990).
- [13] R. C. Hartman *et al.* *Astrophys. J. Suppl. Ser.*, 123 (1999) 79–202.
- [14] J. H. Buckley *et al.* *Astron. & Astrophys.*, 329 (1998) 639–58.
- [15] M. Catanese and T. C. Weekes. *Proc. Astron. Soc. Pacific*, 111 (1999) 1193–222.
- [16] T. Weekes. *Status of VHE astronomy c. 2000.* astro-ph/0010431 (October 2000).
- [17] A. A. Watson. *Phys. Rep.*, 333-34 (2000) 309–27.
- [18] X. Bertou, M. Boratav, and A. Letessier-Selvon. *Intl. J. Mod. Phys.*, A15 (2000) 2181–224.
- [19] P. Blasi, R. I. Epstein, and A. V. Olinto. *Astrophys. J. Letters*, 533 (2000) L123–6.
- [20] V. Berezhinsky, M. Kachelriess, and A. Vilenkin. *Nucl. Phys. B (Proc. Suppl.)*, 70 (1999) 500–2.
- [21] C. M. Hoffman, C. Sinnis, P. Fleury, and M. Punch. *Rev. Mod. Phys.*, 71 (1999) 897–936.

- [22] F. Halzen, in *Neutrinos in Physics and Astrophysics (Proc. Theor. Advanced Study Inst. at Boulder)*, P. Langacker, ed. (World Scientific, Singapore, 1998) pp. 524–69.
- [23] E. Waxman and J. Bahcall. *Phys. Rev. D*, 59 (1999) 023002.
- [24] T. Piran. *Phys. Rep.*, 314 (1999) 575–667.
- [25] E. Waxman. *Phys. Rev. Lett.*, 75 (1995) 386–9.
- [26] M. Milgrom and V. Usov. *Astrophys. J. Letters*, 449 (1995) L37–40.
- [27] M. Vietri. *Astrophys. J.*, 453 (1995) 883–9.
- [28] E. Waxman. *Nucl. Phys. B (Proc. Suppl.)*, 87 (2000) 345–54.
- [29] E. Waxman and J. Bahcall. *Phys. Rev. Lett.*, 78 (1997) 2292–5.
- [30] F. Halzen and D. W. Hooper. *Astrophys. J. Letters*, 27 (1999) L93–6.
- [31] J. Alvarez-Muñiz, F. Halzen, and D. W. Hooper. *Phys. Rev. D*, 62 (2000) 093015.
- [32] The AMANDA Collaboration. *Search for high energy neutrinos from gamma-ray bursts with the AMANDA South Pole neutrino telescope*. Forthcoming.
- [33] A. de Rújula. *The cannonball model of gamma ray bursts*. Presentation at the XXXVth Rencontres de Moriond “Very High Energy Phenomena in the Universe” (January 2001). <http://moriond.in2p3.fr/J01/transparent/DeRujula/>
- [34] P. Meszaros and E. Waxman. *TeV neutrinos from bursting and choked fireballs*. astro-ph/0103275 (March 2001).
- [35] L. Bergström, J. Edsjö, and P. Gondolo. *Phys. Rev. D*, 58 (1998) 103519.
- [36] P. Nießen. *Search for Relativistic Magnetic Monopoles with the AMANDA Detector*. Ph.D. thesis, Humboldt University, Berlin (2001).
- [37] B. Wiebel-Sooth, P. L. Biermann, and H. Meyer. *Astron. & Astrophys.*, 330 (1998) 389–98.
- [38] L. V. Volkova. *Sov. J. Nucl. Phys.*, 31 (1980) 784–790.
- [39] G. Fiorentini, V. A. Naumov, and F. L. Villante. *Atmospheric neutrino flux supported by recent muon experiments*. hep-ph/0103322 (March 2001).
- [40] I. F. M. Albuquerque and G. F. Smoot. *Measuring neutrino oscillations with neutrino telescopes*. hep-ph/0102078 (March 2001).
- [41] T. Gaisser. *Uncertainty in flux of atmospheric neutrinos: Implications for upward muons in AMANDA B10*. AMANDA Int. Rep. 20001201.
- [42] F. Halzen. *Nucl. Phys. B (Proc. Suppl.)*, 38 (1995) 472–83.
- [43] H. Sobel. *Nucl. Phys. B (Proc. Suppl.)*, 91 (2001) 127–33.
- [44] A. B. McDonald. *Nucl. Phys. B (Proc. Suppl.)*, 91 (2001) 21–8.
- [45] F. Halzen. Lectures to the Theoretical Advanced Studies Institute at Boulder (1998).
- [46] D. M. Lowder *et al.* (the AMANDA Collaboration). *Nature*, 353 (1991) 331–3.

- [47] P. B. Price. In *Proc. Zeuthen Workshop on Simulation and Analysis Methods for Large Neutrino Telescopes*, C. Spiering, ed. (July 1998) pp. 120–131.
- [48] The IceCube Collaboration. *IceCube: a kilometer-scale neutrino observatory*.
<http://pheno.physics.wisc.edu/icecube/>
- [49] M. Kowalski. *On the Reconstruction of Cascade-Like Events in the AMANDA Detector*. Diploma thesis, Humboldt University, Berlin (1999).
- [50] J. G. Learned and S. Pakvasa. *Astroparticle Phys.*, 3 (1995) 267.
- [51] R. Gandhi, C. Quigg, M. H. Reno, and I. Sarcevic. *Astroparticle Phys.*, 5 (1996) 81–110.
- [52] P. B. Price and K. Woschnagg. *Role of group and phase velocity in high-energy neutrino observatories* (2000). AMANDA Public Manuscript 20000603.
<http://area51.berkeley.edu/manuscripts/>
- [53] S. G. Warren. *Appl. Opt.*, 23 (1984) 1206–25.
- [54] A. Bouchta. *Muon Analysis with the AMANDA-B Four-String Detector*. Ph.D. thesis, Stockholm University (1998).
- [55] D. E. Groom *et al.* *European Phys. J.*, C15 (2000) 1. (The Particle Data Book.)
<http://pdg.lbl.gov>
- [56] G. Hill. *Experimental and Theoretical Aspects of High Energy Neutrino Astrophysics*. Ph.D. thesis, University of Adelaide (1996).
- [57] C. H. V. Wiebusch. *The Detection of Faint Light in Deep Underwater Neutrino Telescopes*. Ph.D. thesis, Rheinisch-Westfälische Technische Hochschule, Aachen (1995).
- [58] E. Andrés *et al.* (the AMANDA Collaboration). *Astroparticle Phys.*, 13 (2000) 1–20.
- [59] D. Heck. In *Proc. Zeuthen Workshop on Simulation and Analysis Methods for Large Neutrino Telescopes*, C. Spiering, ed. (July 1998) pp. 228–234.
- [60] G. C. Hill. *Astroparticle Phys.*, 6 (1997) 215–228.
- [61] J. Klug. *Enhancement of the Up/Down Signal Ratio and Crosstalk Studies in AMANDA*. Diploma thesis, Uppsala University (1997).
- [62] R. Porrata. *The Energy Spectrum of Pointlike Events in AMANDA-A*. Ph.D. thesis, University of California, Irvine (1997).
- [63] P. B. Price, K. Woschnagg, and D. Chirkin. *Age vs. depth of glacial ice at South Pole* (2000). AMANDA Public Manuscript 20000201.
<http://area51.berkeley.edu/manuscripts/>
- [64] S. Tilav. Presentation to the AMANDA collaboration meeting. Irvine, California (March 1998).
- [65] P. Desiati, personal communication.
- [66] C. Wiebusch. *Muon reconstruction with AMANDA*. DESY Internal Report, DESY-PROC-1-1999 (1999).
<http://www.ifh.de/~wiebusch/publications/amareco.ps.gz>
- [67] A. Karle. In *Proc. Zeuthen Workshop on Simulation and Analysis Methods for Large Neutrino Telescopes*, C. Spiering, ed. (July 1998) pp. 174–185.

- [68] A. Karle, personal communication.
- [69] G. C. Hill. *Application of Bayesian inference to event reconstruction and background rejection in AMANDA*. Forthcoming.
- [70] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press) (1995).
- [71] K. Rawlins, personal communication.
- [72] K. Hanson. *Fast, accurate reconstruction for AMANDA with a simplex minimizer*.
<http://uhe.hep.upenn.edu/~kaeld/amanda-int/simplex/simplex.html>
- [73] C. Perez de los Heros. private communication.
- [74] A. Biron, M. Gaug, C. H. Wiebusch, and R. P. Wischnewski. *Analysis of atmospheric neutrinos in 97 data*. AMANDA Int. Rep. 20010501.
- [75] J. Lamoureux and J. Jacobsen. *(Re)filtering/reconstructing 97/98/99 up-going AMANDA events*.
http://rust.lbl.gov/~jodi/amanda/97_refilter/filt97-upgoing.ps
- [76] V. Stenger. *Track fitting for the DUMAND octagon*. DUMAND Internal Report, HDC-1-90 (1990).
- [77] S. Barwick, personal communication.
- [78] T. DeYoung *et al.*. *Search for neutrinos with the AMANDA 10 string array — a first analysis of the 1997 data*. AMANDA Int. Rep. 19990101.
- [79] A. Biron, M. Leuthold, and C. H. Wiebusch. *Separation of muon neutrino event candidates in AMANDA-B*. AMANDA Int. Rep. 19990902.
- [80] G. Barouch *et al.*. *Search for atmospheric neutrino events in AMANDA-B10*. AMANDA Int. Rep. 20000101.
- [81] M. Gaug, P. Nießen, and C. Wiebusch. *Investigations on smoothness observables*. AMANDA Int. Rep. 20000201.
- [82] M. Gaug. *Detection of Atmospheric Muon Neutrinos with the AMANDA Neutrino Telescope*. Diploma thesis, Humboldt University, Berlin (2000).
- [83] W. Rhode, private communication.
- [84] W. Lohmann, R. Kopp, and R. Voss. *Energy loss of muons in the energy range 1-10000 GeV*. CERN Yellow Report, 85-03 (1985).
- [85] P. Lipari and T. Stanev. *Phys. Rev. D*, 44 (1991) 3543–54.
- [86] P. Olbrechts and C. Wiebusch. *On the angular sensitivity of optical modules in ice*. AMANDA Int. Rep. 20010102.
- [87] P. Mioćinović, personal communication.
- [88] R. Wischnewski, personal communication.
- [89] F. Halzen, personal communication.

- [90] The AMANDA Collaboration. *Search for point sources of high energy neutrinos with AMANDA*. Forthcoming.
- [91] D. E. Alexandreas *et al.*. *Nucl. Inst. & Meth. in Phys. Res.*, A328 (1993) 570–577.
- [92] T.-P. Li and Y.-Q. Ma. *Astrophys. J.*, 272 (1983) 317–324.
- [93] A. Biron and W. Rhode. Presentation to the AMANDA collaboration meeting. Newark, Delaware (March 2001).
- [94] G. J. Feldman and R. D. Cousins. *Phys. Rev. D*, 57 (1998) 3873–89.
- [95] R. Cousins and V. Highland. *Nucl. Inst. & Meth. in Phys. Res.*, A320 (1992) 331–335.
- [96] J. Conrad, O. Botner, A. Hallgren, and C. Perez de los Heros. *Integrating beliefs: The calculation of upper limits on number of events and flux including systematic uncertainties*. AMANDA Int. Rep. 20010101.
- [97] G. Hill. *Application of importance sampling to the atmospheric muon generator “basiev”*. AMANDA Int. Rep. 20000904.
- [98] E. Andrés *et al.* (the AMANDA Collaboration). *Nature*, 410 (2001) 441–443.
- [99] K. Hanson. *Crosstalk in 1997 AMANDA raw data*. AMANDA Int. Rep. 20000801.
- [100] *Simulation of Events with GEANT for Muon and Neutrino Detectors (SiEGMuND)*. <http://www.ifh.de/nuastro/software/siegmund/siegmund.html>
- [101] T. DeYoung, J. Edsjö, J. Kim, and C. Wiebusch. *Standard reconstruction script for AMANDA-B 1997 data*.
- [102] S. Gadomski *et al.*. *Nucl. Inst. & Meth. in Phys. Res.*, A320 (1992) 217–227.
- [103] D. Pandel. *Determination of water and detector parameters and reconstruction of muons up to 100 TeV with the Baikal neutrino telescope NT-72*. Diploma thesis, Humboldt University, Berlin (1996).
- [104] T. J. Loredo, in *Maximum Entropy and Bayesian Methods*, P. F. Fougère, ed. (Kluwer Academic Publishers, Dordrecht, the Netherlands, 1990) pp. 81–142.
- [105] N. Metropolis, *et al.*. *J. of Chemical Phys.*, 21 (1953) 1087–1092.
- [106] R. Battiti and G. Tecchiolli. *Neurocomputing*, 6 (1994) 181–206.
- [107] R. Battiti and G. Tecchiolli. *O. R. S. A. Journal on Computing*, 6 (1994) 126–140.
- [108] R. Battiti and G. Tecchiolli. *Annals of Operations Research*, 63 (1996) 153–188.

Appendix A

Combining Weighted Monte Carlos

This analysis relied on the 2000 `basiev` mass production for simulation of the downgoing muon background. This mass production used importance sampling techniques to improve the precision of the simulation while reducing simulation time, focusing the simulation on important regions of parameter space which would not be heavily sampled in a natural Monte Carlo. The Monte Carlo outputs the proper factors to use for unbiasing the output, which depend on the distributions from which the program sampled events. However, the optimization strategy changed from versions 0–2, which focused on near-horizontal muons, to version 3, which sampled uniformly in zenith angle up to a cutoff angle. This means that the effective sampling distribution, on which the weights of the individual events depends, must be recalculated.

The weight w of an event in a biased Monte Carlo is the ratio of the probability density function f which is to be simulated to the biased distribution p actually sampled by the Monte Carlo program. These functions are probability densities and so must be normalized, and the ratio is evaluated at the point corresponding to the event in question. For example, in AMANDA we wish to simulate an isotropic flux of protons producing downgoing muons, but for versions 0–2 we actually sampled from a $1/\cos^2\theta$ distribution. The normalized true and biased distributions were then

$$f(\cos\theta) = \frac{1}{\alpha_0} \quad \text{and} \quad p(\cos\theta) = \frac{1}{\beta_0 \cos^2\theta} \quad (\text{A.1})$$

respectively, where α_0 and β_0 are the normalization constants

$$\alpha_0 = \frac{1}{\cos\theta_{\text{high}} - \cos\theta_{\text{low}}} \quad \text{and} \quad \beta_0 = \int_{0.06}^1 \frac{1}{\cos^2\theta} d(\cos\theta) = \frac{47}{3} \quad (\text{A.2})$$

and $\cos\theta_{\text{high}}$ and $\cos\theta_{\text{low}}$ were the limits of the simulation, $[0.06, 1]$.

The weighting function was then

$$w(\cos \theta) = \frac{f(\cos \theta)}{p(\cos \theta)} = \frac{\beta_0 \cos^2 \theta}{\alpha_0}$$

and the weight of an event at an angle θ_0 was found by evaluating $w(\cos \theta)$ at $\theta = \theta_0$. For convenience, this calculation was performed by the Monte Carlo at run time, based on the sampling distribution then in use, and recorded in the F2000 data stream. For example, the above weight was output as the variable `zenithweight` (`znthwght` in `munt`).

However, the values calculated by the Monte Carlo depend on the sampling distributions actually being used at the time, which may not be correct for the full set of simulated events taken as a whole. To see this, consider that the sampling distribution for the zenith angle can be defined (in the limit of a large number of events) as the differential number of events n at a given angle, divided by the total number of events N :

$$p(\cos \theta) = \frac{n(\cos \theta)}{N}.$$

If simulated events from production runs with different sampling distributions are to be combined, we must calculate the *effective* sampling distribution

$$p_{\text{eff}}(\cos \theta) = \frac{n_1(\cos \theta) + n_2(\cos \theta)}{N_1 + N_2},$$

from which *de facto* the events were drawn.

As described in [97], the background simulation was biased in three distributions simultaneously. Each of these biased samplings — in zenith angle, threshold energy, and generation plane area — produces its own weighting factor w_i . The total weight w of an event is simply the product of the weights from each of the biased samplings:

$$w = \prod_i w_i.$$

Only one of the three sampling distributions — the zenith angle — was changed over the course of the simulation. The calculations of the other factors (`gnplnwgh` and `smpbvgh`) performed by the Monte Carlo at run time are therefore still valid, and we need only recalculate the zenith angle factor `znthwght`.

As described above, versions 0–2 of the mass production contain some $N_0 = 1.582 \cdot 10^{11}$ simulated protons chosen from a $1/\cos^2\theta$ sampling distribution. Beginning with version three, it was decided to flesh out the prediction for more vertical muons by removing the zenith angle weighting and sampling from a uniform distribution in $\cos\theta$. To avoid redundancy, the angular range of versions three and higher was restricted slightly to avoid the horizontal region already sampled heavily by versions 0–2. Versions 3–4 thus contain $N_3 = 9.955 \cdot 10^{10}$ protons simulated uniformly over $0.276 \leq \cos\theta \leq 1$.

Overall, then, we have simulated $N = N_0 + N_3 = 2.578 \cdot 10^{11}$ protons. The differential number of these protons simulated in an infinitesimal range about $\cos\theta$ is

$$n(\cos\theta) d(\cos\theta) = [n_0(\cos\theta) + n_3] d(\cos\theta)$$

where

$$n_i = N_i p_i.$$

For versions 0–2, p_0 is as calculated in Eqs. A.1–A.2, above. The sampling was uniform in $\cos\theta$ for versions 3–4, so p_3 consists only of the stepwise normalization factor

$$p_3 = \begin{cases} \frac{1}{\beta_3} = \frac{1}{\cos\theta_{\text{high}}^{(3)} - \cos\theta_{\text{low}}^{(3)}} = \frac{1}{0.724} & 0.276 \leq \cos\theta \leq 1 \\ 0 & \cos\theta < 0.276 \end{cases}$$

The effective sampling distribution used is thus

$$p(\cos\theta) = \frac{n}{N} = \frac{n_0(\cos\theta) + n_3}{N_0 + N_3} = \begin{cases} \frac{1}{N_0 + N_3} \left(\frac{N_0}{\beta_0 \cos^2\theta} + \frac{N_3}{\beta_3} \right) & 0.276 \leq \cos\theta \leq 1 \\ \frac{1}{N_0 + N_3} \left(\frac{N_0}{\beta_0 \cos^2\theta} \right) & 0.06 \leq \cos\theta < 0.276 \end{cases}$$

The weight of an event is, as usual, the ratio of the unbiased distribution to the distribution actually sampled

$$w = \frac{f(\cos\theta)}{p(\cos\theta)} = \frac{1}{\cos\theta_{\text{high}} - \cos\theta_{\text{low}}} \left(\frac{N_0 + N_3}{n_0(\cos\theta) + n_3} \right).$$

The other factor which is needed to make an absolute prediction is the simulated time, which sets the absolute scale of the weighting factor calculated above. The simulated time is given as usual by

$$T_{\text{gen}} = \frac{N}{\Phi A \Omega}$$

where N , as above, is the number of protons simulated, Φ is the integrated flux above the energy threshold of the simulation, A is the area of the generation plane, and Ω is the solid angle covered by the simulation.

We have used a flux $\Phi = 0.102 \text{ m}^{-2}\text{s}^{-1}\text{sr}^{-1}$, which is the flux of protons and helium above a nominal threshold of 1 TeV. The nominal generation plane was $A = 4.65 \cdot 10^5 \text{ m}^2$, and the solid angle is as given above¹. We therefore have simulated a total of

$$T_{\text{gen}} = \frac{N}{0.102 \cdot 4.65 \cdot 10^5 \cdot 0.94 \cdot 2\pi} = 9.203 \cdot 10^5$$

seconds of livetime. To make a prediction for the full 1997 data set, consisting of 130.1 days of livetime, we simply multiply each event weight by a scale factor $T_{\text{live}}/T_{\text{gen}} = 12.21$.

¹Note that the area and threshold given above are reference values; the actual simulation sampled above a different threshold and over a different area, depending on the angle of the incident proton. These terms are accounted for in the `basiev` output values of `smpbvwh` and `gnplnwh`.

Appendix B

Cross Talk

B.1 Channel Cuts

A large amount of cross talk can be seen in distributions of amplitude vs. pulse length (time over threshold) for various modules. The plots in Figures B.1 and B.2 show these distributions for a few modules in B-10 designated ‘good’ by the standard reconstruction script and used in the normal analysis. Each point in the plots represents one hit from the set of some $6 \cdot 10^4$ events passing the Bayesian reconstruction¹. The standard hit cleaning of Sec. 6.2 has been used in these plots, so that only the first pulse in each channel is plotted for any event, and the standard ADC, TOT, and isolation cuts have been applied.

Cross talk can be seen in many channels as points which lie above and to the left of the characteristic wing-shaped distribution. These short, high-amplitude pulses are much more frequent in this data set than in trigger-level data, which makes identification of cross talking channels easier. The first two channels in Fig. B.1 are normal channels, with low levels of cross talk. The other six channels exhibit high cross talk levels, and further cleaning is required. It seems that the cross talk in different modules has different characteristics; for example, in OM 133 the anomalous pulses are all of very short duration, although the amplitude is quite high. In channel 102 there is a wider variation in pulse duration, and in OM 155 there is a characteristic (low) amplitude. Whether these correlate with different physical processes is unknown.

Based on these distributions, and those of the rest of the modules in the array, we developed the

¹In fact, the event set is from an older version of the Bayesian filter, the set on which [98] is based. However, the set is very similar to that produced by the analysis described herein.




Figure B.1: Pulse amplitudes vs. duration for modules on the outer strings. Normal hits lie in the wing-shaped distribution extending to the right, as shown in channels 98 and 129. Channels with high levels of cross talk have large numbers of hits at high amplitude but with short duration, as in channels 95 and 102.




Figure B.2: Pulse amplitudes vs. duration for modules exhibiting high levels of cross talk. The pile-up seen in channels 95, 129, and 133 is caused by ADC saturation.

TOT Thresholds			
Channel	Min. TOT	Channel	Min. TOT
95	220	184	Removed
102	Removed	191	140
128	180	194	220
133	180	204	160
146	180	205	200
155	160	206	140
156	180	219	Removed
168	150	220	160
169	160	225	170
170	180	228	160
171	140	230	190
174	150	246	160
176	180	256	180
177	140	257	Removed
178	140	289	Removed
179	180	290	Removed
180	140	291	Removed
181	Removed	292	Removed
182	Removed	301	Removed

Table B.1: List of additional TOT cuts used to reduce cross talk.

list of channels given in Table B.1, to which additional TOT cuts, more stringent than the standard 125 ns minimum, were applied, or which were removed from the data stream entirely.

B.2 Hit Cuts

In an effort to further reduce the level of cross talk, we also implemented two new cross talk cleaning algorithms, based on the analyses by J. Klug [61] and K. Hanson [99]. These algorithms attempt to identify cross talk pulses by looking for low amplitude hits which come in (raw time) coincidence with hits of large amplitude. The first algorithm simply looks at all pairs of pulses on a string. However, the time coincidence window cannot be made too wide without losing large numbers of signal hits, and it has been found by both Klug and Hanson that some cross talk arrives with delays of up to a few hundred nanoseconds with respect to the inducing pulse. Therefore, the second algorithm uses expanded coincidence windows but, to preserve real hits, looks only at pairs of hits in channels identified by Klug or Hanson as being strongly coupled to each other.

Using these improved hit cleaning algorithms, we re-reconstructed the 58,000 events passing

the Bayesian filter (the “meat pie” reconstructions of Appendix C). With much of the cross talk removed, only 16% of the data were still reconstructed as upgoing. By contrast, some 85% of the simulated misreconstructed background survived, indicating that the bulk of the events removed were unsimulated fakes. Moreover, while peaks in COGz remained in the data, as seen in Fig. 7.12, they were much reduced in size relative to the ‘continuum’ of simulated background. We take the fact that this class of fakes is greatly reduced by the cross talk cleaning to be strong evidence that the cross talk hypothesis is correct.

Clearly, we have not completely eliminated the ‘COG’ fakes. Based on the number of simulated fakes, however, we believe we have reduced the class of fakes by at least 90%. While it is possible that the remaining unsimulated fakes are caused by something completely different, we believe that they are caused by cross talk that has escaped our cleaning routines.

Appendix C

Reconstruction Chain

The series of reconstructions applied to the data is described in this appendix. The analysis began with the data that passed the LBL filter [75], the first stage in Table C.1. The next two stages constitute a large multiple-restart search for the best upgoing and downgoing minima of the Bayesian posterior probability function, done with a slight correction to the standard hit cleaning.

By the time that the data set had been passed through these filters, the cross talk cleaning algorithms described in Appendix B had been developed. Rather than going back to the full set of data from the LBL filter, it was decided to simply refit the data set as it stood, some 50,000 events. It is only the fits from this final stage of the reconstruction that are used in this report, but the previous stages are described here because they will have some effect on the efficiency of the analysis.

In Table C.1, fits labelled Bayesian were performed with the Upandel likelihood multiplied by the prior described in 5.2. Fits referred to as ‘upgoing’ or ‘downgoing’ were seeded with random initial guesses chosen isotropically from the allowed hemisphere and passing through the center of gravity of the event. Unrestricted fits were based on the most recent line fit. It was necessary to redo the line fit after changing the hit cleaning, because even a single hit can cause a large change in the line fit solution. Changing the cleaning also made previous fits unusable, because when using timing likelihoods alone the minima found are not tied to any absolute scale; the values change sharply when a hit is added or removed.

Details of these reconstructions, including the software flags used by the reconstruction program `recoos` [100], are given in Tables C.2–C.5. Note that in all reconstructions, the time-over-threshold minimum of 125 ns was not applied to channels 291–302 after run 800, due to a malfunction

LBL Filter			
<i>This stage is described in [75]</i>			
Reco 1		Line fit	Single iteration
	Cut 1	$\theta_{LF} > 50$	
Reco 2		Upandel	
	Cut 2	$\theta_{Upandel} > 80$	Applied without hit cleaning
	Cut 3	$N_{dir}^{[-15:25]} > 3$	
Dr. Seuss			
<i>For these fits the isolation cleaning was done last</i>			
	Cut 4	$N_{dir}^{[-15:25]} > 3$	Reapplied with hit cleaning
Reco 3		Line fit	Single iteration
Reco 4		Unrestricted Upandel	
Reco 5		Upgoing Bayesian	
Reco 6		Downgoing Bayesian	
	Cut 5	Best of fits 4–6 is upgoing	
Meat Pie			
Reco 7		Upgoing Bayesian	4 iterations
Reco 8		Downgoing Bayesian	4 iterations
	Cut 6	Best of fits 4–8 is upgoing	20 iterations
Reco 9		Downgoing Bayesian	
Reco 10		Best upgoing	
Reco 11		Best downgoing	
	Cut 7	Best of fits 10, 11 is upgoing	
Final			
<i>For these fits the cross talk cleaning was applied</i>			
Reco 12		Line fit	Single, based on fit 12
Reco 13		Unrestricted Bayesian	
Reco 14		Upgoing Bayesian	
Reco 15		Downgoing Bayesian	
Reco 16		Best upgoing	
Reco 17		Best downgoing	
	Cut 8	Best of 16, 17 is upgoing	
Reco 18		Tensor of inertia fit	
Reco 19		Shower fit	

Table C.1: Outline of the reconstruction chain.

in the electronics. Also, prior to the final fit, a bug in the `recoos` program caused the dark noise rate (in hits/ns) used in the likelihood function to be effectively squared, causing the reconstruction to use very small probabilities for noise hits. These values are noted in the tables.

Table C.2: Parameters of the LBL Reconstructions

Parameter	Value	recoos flag
LBL Line Fit Parameters		
Reconstruction Type	Line fit	-r n -i m
Amplitude Weighting	Unweighted	-p w=0.
Time Shift	All residuals positive	-X g=f
Minimum Event Size	5 hits on 1 string	-p t=1:5
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
LBL Reconstruction Parameters		
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	LBL line fit	-i f -p f=1
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upandel	-z a_upandel
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	90 μ Hz over 4.5 μ s	-p n=300:4500 (with noise bug)
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1

Table C.3: Parameters of the Dr. Seuss Reconstructions

Parameter	Value	recoos flag
Dr. Seuss Line Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Line fit	-r n -i m
Amplitude Weighting	Unweighted	-p w=0.
Time Shift	All residuals positive	-X g=f
Minimum Event Size	5 hits on 3 strings	no flags (default)

Dr. Seuss Unrestricted Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	Dr. Seuss line fit	-i f
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upandel	-z a_upandel
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 mHz over 4.5 μ s	-p n=1000:4500 (with noise bug)
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2

Dr. Seuss Upgoing Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	Random upgoing	-i r -p Z=-1.:0.

continued on next page

Table C.3: *continued*

Parameter	Value	recoos flag
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 mHz over 4.5 μ s	-p n=1000:4500 (with noise bug)
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2

Dr. Seuss Downgoing Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	Random downgoing	-i r -p Z=0.:1.
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 mHz over 4.5 μ s	-p n=1000:4500 (with noise bug)
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2

Table C.4: Parameters of the Meat Pie Reconstructions

Parameter	Value	recoos flag
Meat Pie Upgoing Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.

continued on next page

Table C.4: *continued*

Parameter	Value	recoos flag
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Random multistart	-r g
Number of Iterations	4	-p M=3
Starting Hypothesis	Random upgoing	-i r -p Z=-1.:0.
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 mHz over 4.5 μ s	-p n=1000:4500 with noise bug
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2

Meat Pie Downgoing Fit Parameters		
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Reconstruction Type	Random multistart	-r g
Number of Iterations	4 or 20	-p M=3 or -p M=19
Starting Hypothesis	Random downgoing	-i r -p Z=0.:1.
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Upgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 mHz over 4.5 μ s	-p n=1000:4500 (with noise bug)
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2

Table C.5: Parameters of the Final Reconstructions

Parameter	Value	recoos flag
Final Line Fit Parameters		
Reconstruction Type	Line fit	-r n -i m
Amplitude Weighting	Unweighted	-p w=0.
Time Shift	All residuals positive	-X g=f
Minimum Event Size	5 hits on 1 string	-p t=1:5
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Final Unrestricted Fit Parameters		
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	Final line fit	-i f
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Bayesian Upandel	-z a_upandel+a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 kHz over 4.5 μ s	-p n=1000:4500
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Final Upgoing Fit Parameters		
Reconstruction Type	Random multistart	-r g
Number of Iterations	4	-p M=3
Starting Hypothesis	Random upgoing	-i r -p Z=-1.0:0.0
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p

continued on next page

Table C.5: *continued*

Parameter	Value	recoos flag
Function Minimized	Upgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 kHz over 4.5 μ s	-p n=1000:4500
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1

Final Downgoing Fit Parameters		
Reconstruction Type	Random multistart	-r g
Number of Iterations	4	-p M=3
Starting Hypothesis	Random downgoing	-i r -p Z=0.0:1.0
Time Shift	All residuals positive	-X g=f
Local Minimizer	Powell's	-m p
Function Minimized	Downgoing Bayesian Upandel	-z a_upandel+a_zenith_range +a_zwght2
Parametrization	Vertex, angles	no flags (default)
Free Parameters	Vertex, zenith, azimuth	-x x,y,z,zenith,azimuth
Noise Rate	1 kHz over 4.5 μ s	-p n=1000:4500
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Minimum Event Size	5 hits on 1 string	-p t=1:5
Hole Ice Optics	50 cm scattering	-X o=2
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1

Table C.6: Parameters of the Shower Reconstructions

Parameter	Value	recoos flag
Tensor of Inertia Fit Parameters		
Reconstruction Type	Tensor of inertia fit	-r n -i i
Particle Type	Electron	-p p=e-
Time Shift	All residuals positive	-X g=f
Amplitude Weighting	Linear	-p w=1.
Trafo	None	-X s=n
Minimum Event Size	4 hits on 1 string	-p t=1:4
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1
Shower Reconstruction Parameters		
Local Minimizer	Powell's	-m p
Reconstruction Type	Single minimization	-r m
Starting Hypothesis	Tensor of inertia fit	-i f
Particle Type	Electron	-p p=e-
Time Shift	None	-X g=n
Trafo	None	-X s=n
Function Minimized	Point source Upandel	-z a_pp_upandel
Likelihood Parametrization	Vertex, time	-x xyzt
Free Parameters	Vertex, time	-x x,y,z,time
Minimum Event Size	4 hits on 1 string	-p t=1:4
Hole Ice Optics	50 cm scattering	-X o=2
Noise Rate	1 kHz over 4.5 μ s	-p n=1000:4500
Pandel Jitter	15 ns	-p j=15
Pandel Absorption Length	96 m	-p a=96.
Rejected Channels	Standard	See [101]
Time Window	4.5 μ s	-y R=0.:4500.
Cross Talk Cleaning I	Hit Coincidences	-y x=50:8:0.5 -y X=8:0.5
Cross Talk Cleaning II	Bad Channels	See Table B.1
Amplitude	0.3–1000 p.e.	-y a=0.3:1000.
Time Over Threshold	125–2000 ns	-y b=125:2000.:1:302
Multiple Hits	First hit only	-y A=1
Hit Isolation	70 m, 500 ns	-y I=70.:500.:1

Appendix D

Event Quality Levels

This appendix contains the values of the cut parameters which were used to define the different quality levels Q referred to in Chapters 7 and 8, as well as the distribution of events for the data, simulated atmospheric neutrinos, and simulated downgoing muons.

Table D.1 shows the definitions of the different Q levels in terms of the six cut parameters defining the cut space. Each entry is the lower boundary of the Q level, so for example events which do not pass the cuts in the first row have quality $Q = 0$. The cut values for N_{dir} and L_{dir} are the minimum values required; all others are maxima.

Table D.2 gives the number of events found *at* each quality level (i.e., *not* inclusive of higher Q levels). The uncertainties quoted are taken from the observed variances of the sets,

$$\sigma = \sqrt{\sum w_i^2}$$

(where w_i is the weight of event i) and so are estimates of the statistical error only. Note that especially for the downgoing muon simulation, insufficient statistics are available for the higher quality levels, and so the quoted uncertainties are probably underestimates.

Event Quality Cuts						
Quality	$\frac{N_{\text{dir}}^{[-15:75]}}{N_{\text{hits}}}$	$ S_{P_{\text{hit}}} $	$L_{\text{dir}}^{[-15:25]}$	$\frac{-\ln \mathcal{L}}{N_{\text{hits}} - 5}$	$\frac{I_1}{\Sigma I}$	$\theta_{\text{like}} - \theta_{\text{LF}}$
1	0.2884	0.4565	30.3	8.890	0.2720	57.27
2	0.3296	0.4047	42.0	8.679	0.2579	47.26
3	0.3574	0.3679	50.2	8.572	0.2479	40.97
4	0.3793	0.3404	57.0	8.491	0.2403	36.35
5	0.3979	0.3202	62.6	8.416	0.2337	32.17
6	0.4144	0.3029	66.7	8.348	0.2283	28.96
7	0.4295	0.2887	70.4	8.304	0.2233	26.29
8	0.4436	0.2791	73.4	8.272	0.2185	23.83
9	0.4570	0.2689	76.6	8.236	0.2143	21.99
10	0.4700	0.2600	80.0	8.200	0.2100	20.00
11	0.4827	0.2516	82.6	8.169	0.2061	18.34
12	0.4952	0.2444	85.2	8.145	0.2022	16.94
13	0.5078	0.2375	87.8	8.123	0.1987	15.57
14	0.5206	0.2311	90.0	8.099	0.1951	14.45
15	0.5336	0.2262	92.3	8.080	0.1918	13.31
16	0.5472	0.2205	94.5	8.057	0.1882	12.27
17	0.5615	0.2156	96.6	8.039	0.1852	11.35
18	0.5769	0.2113	98.2	8.021	0.1818	10.53
19	0.5937	0.2072	99.9	8.003	0.1783	9.96
20	0.6127	0.2033	102.0	7.989	0.1749	8.85
21	0.6349	0.1987	103.2	7.972	0.1712	8.10
22	0.6627	0.1949	104.6	7.957	0.1680	7.38
23	0.7015	0.1910	105.8	7.942	0.1646	6.78
24	0.7750	0.1876	106.8	7.928	0.1612	6.12

Table D.1: Cuts defining the event quality levels. The cuts applied for the atmospheric neutrino analysis correspond to $Q = 7$ (underlined); for the point source search the $Q = 5$ cuts were used.

Event Quality Distributions			
Quality	Obs. Data	Atmospheric ν	Downgoing μ
0	4164	58.5 ± 1.3	1479.2 ± 65.2
1	597	43.8 ± 1.3	227.7 ± 24.4
2	225	43.5 ± 1.3	103.8 ± 15.9
3	78	36.9 ± 1.1	26.3 ± 6.8
4	58	35.2 ± 1.1	33.7 ± 11.7
5	38	35.7 ± 1.1	4.7 ± 2.1
6	44	31.7 ± 1.1	6.4 ± 3.5
7	12	26.3 ± 1.0	1.8 ± 1.2
8	20	27.3 ± 1.0	1.0 ± 0.6
9	13	24.8 ± 0.9	1.1 ± 0.8
10	20	24.7 ± 1.0	4.4 ± 3.1
11	9	16.8 ± 0.8	3.4 ± 2.2
12	11	25.5 ± 1.0	0.0 ± 0.0
13	11	16.7 ± 0.7	0.0 ± 0.0
14	15	19.2 ± 0.9	1.4 ± 0.9
15	7	14.7 ± 0.7	0.0 ± 0.0
16	10	14.4 ± 0.7	0.0 ± 0.0
17	11	11.4 ± 0.6	0 ± 0
18	8	13.1 ± 0.7	0 ± 0
19	13	11.8 ± 0.7	0 ± 0
20	6	9.1 ± 0.6	0.0 ± 0.0
21	7	8.6 ± 0.6	0.0 ± 0.0
22	15	7.8 ± 0.6	0.1 ± 0.1
23	9	5.2 ± 0.5	0 ± 0
24	7	1.6 ± 0.3	0 ± 0

Table D.2: Number of events found at each of the quality levels defined in Table D.1. The uncertainties on the simulated sets are statistical.

Appendix E

Events

Table E.1: List of Events

Event Num.	Day	Seconds	Zenith	Azimuth	Rt. Asc.	Declin.	Gal. Long.	Gal. Lat.
557701	97	76050	118.13	183.50	3.97	28.13	165.63	-18.86
2023536	99	69520	146.56	283.13	19.64	56.57	88.74	16.30
1528728	100	30458	156.75	354.36	4.08	66.75	140.39	10.67
253874	102	43758	135.43	93.74	1.29	45.42	127.67	-17.21
478874	105	8523	124.01	66.50	17.49	34.01	58.13	30.93
1260978	106	32729	103.83	7.15	4.25	13.83	179.81	-25.84
2682908	107	54029	163.36	165.08	23.72	73.36	118.03	11.14
1184319	108	51898	109.56	307.29	13.71	19.55	2.87	76.01
6693280	109	44184	163.78	121.46	0.02	73.78	119.37	11.25
3382545	110	30663	116.69	336.00	6.02	26.69	183.66	1.78
236184	111	60185	130.55	198.12	23.50	40.55	106.58	-19.71
2570660	112	4687	165.86	209.92	7.32	75.86	138.86	27.92
557244	114	54943	148.13	7.88	10.92	58.13	148.15	52.92
105696	115	1189	139.72	107.19	13.39	49.72	109.90	66.59
247763	116	11499	149.54	22.59	21.97	59.55	102.52	3.74
49481	116	15134	165.59	245.87	8.09	75.59	138.88	30.80
137711	118	27074	143.19	122.43	19.78	53.20	86.14	13.75
67863	118	53373	117.18	216.07	20.86	27.18	71.32	-10.87
978033	119	67866	142.41	84.90	9.71	52.41	163.87	46.89
11815	121	44011	142.81	193.65	19.95	52.81	86.52	12.23
609698	122	66693	143.89	27.22	13.43	53.89	112.03	62.51
1332456	122	76340	167.37	352.99	18.39	77.37	108.71	27.93
1391955	123	80281	102.58	5.94	18.69	12.59	43.10	7.85
297031	124	14136	162.63	256.68	7.62	72.63	142.51	29.20
48064	124	34983	125.04	114.03	22.94	35.04	97.72	-22.13
517683	124	61144	114.08	314.93	16.83	24.08	43.89	36.68
1129571	124	69313	154.82	304.71	19.79	64.82	97.08	18.75
333605	125	36977	115.91	130.35	22.47	25.91	86.83	-26.74
3634483	129	63934	146.60	131.32	6.18	56.60	157.46	17.09
397392	130	20631	130.99	246.74	10.49	40.99	178.25	57.73
4285271	130	73628	114.84	154.21	7.42	24.84	193.68	18.02
10604848	131	73508	147.65	244.34	1.44	57.65	127.61	-4.90

continued on next page

Table E.1: *continued*

Event Num.	Day	Seconds	Zenith	Azimuth	Rt. Asc.	Declin.	Gal. Long.	Gal. Lat.
7492181	133	17292	140.75	189.58	13.56	50.75	106.93	65.04
910225	133	52246	169.89	27.91	10.08	79.89	131.07	34.44
10708643	135	6065	145.47	65.64	18.83	55.47	85.16	22.28
282164	135	17544	138.77	331.15	4.33	48.77	154.25	-1.02
3596687	136	7172	151.20	226.29	8.49	61.20	155.40	35.30
2312069	138	67355	142.28	8.37	15.92	52.28	82.17	47.69
5362470	139	21345	148.53	304.23	7.44	58.53	158.46	27.38
6472354	139	35996	110.00	207.40	17.98	20.00	45.67	20.27
21328	141	50921	128.80	97.51	5.59	38.80	170.53	3.47
4196343	142	16981	121.91	84.39	21.08	31.91	76.87	-10.05
3391162	142	73198	111.17	1.62	18.26	21.17	48.41	17.12
3694138	142	76962	148.64	246.48	2.98	58.63	138.89	-0.23
9281799	144	68567	127.58	358.80	17.29	37.58	61.67	34.05
2313405	145	18597	98.00	200.49	13.99	8.00	346.12	65.01
3170212	145	29898	144.10	199.71	17.19	54.10	81.77	36.17
99192	148	13565	135.45	256.25	9.07	45.45	174.91	41.78
1774178	148	34687	97.49	259.39	14.74	7.49	1.90	56.80
2065117	148	38354	170.35	213.18	18.84	80.35	112.11	26.75
2143739	149	82346	113.16	356.02	21.64	23.16	75.45	-21.42
3277242	150	54435	150.95	125.48	5.30	60.95	150.17	13.14
350735	152	84818	145.87	34.11	19.99	55.87	89.44	13.42
1707787	153	63950	162.76	353.13	16.97	72.76	104.56	34.09
2998385	155	53249	122.33	328.13	15.79	32.34	51.65	51.73
3911386	155	65086	170.30	9.40	16.33	80.30	113.83	32.78
6963997	156	17518	102.82	34.81	1.46	12.82	136.38	-49.11
6967794	156	17566	146.87	166.68	16.68	56.87	85.95	40.13
232853	157	11403	101.17	196.20	13.06	11.17	313.63	73.79
70402	159	28937	120.52	268.12	13.28	30.52	65.88	83.50
4430785	159	84987	153.60	243.65	6.52	63.60	151.46	21.96
6942749	160	30877	149.82	8.25	7.21	59.82	156.72	25.82
540434	160	55317	116.44	200.84	1.18	26.44	128.28	-36.23
1122544	161	31336	153.70	153.40	21.73	63.70	103.84	8.02
3015752	162	79194	145.18	349.17	22.07	55.18	100.57	-0.27
3568652	163	71	132.33	58.23	19.49	42.33	74.89	11.41
578778	163	43554	133.35	14.12	10.55	43.35	173.64	57.72
7757984	164	50597	141.03	188.04	0.98	51.03	124.11	-11.83
10153642	164	80912	106.47	208.36	8.07	16.47	205.69	23.30
872673	166	59302	146.72	112.70	8.56	56.72	160.82	36.27
10428633	168	81141	160.05	1.62	22.18	70.05	109.95	11.34
486617	169	46474	162.22	2.33	12.54	72.22	124.97	44.84
4492679	170	12187	116.77	238.68	11.30	26.77	208.55	69.20
5178189	170	21153	160.08	245.59	13.34	70.08	119.36	46.84
2009955	171	36364	149.01	322.74	12.50	59.01	128.19	57.90
2636618	176	55019	115.90	24.57	13.90	25.90	31.43	75.98
2976350	176	59351	172.08	126.80	8.29	82.08	131.42	29.77
7530458	177	31087	144.95	79.30	3.65	54.95	145.48	-0.41
7599807	177	31966	143.18	95.47	2.82	53.18	140.12	-5.72

continued on next page

Table E.1: *continued*

Event Num.	Day	Seconds	Zenith	Azimuth	Rt. Asc.	Declin.	Gal. Long.	Gal. Lat.
1222689	177	79448	145.63	256.84	5.28	55.63	154.63	10.10
210637	179	24511	111.96	65.24	2.89	21.96	156.90	-32.76
438934	179	27293	122.77	61.80	3.89	32.77	161.57	-16.12
4760153	192	1280	162.65	21.64	0.18	72.65	119.85	10.02
2174805	194	44590	147.71	106.75	6.70	57.71	157.92	21.45
729394	195	40231	108.02	139.32	3.38	18.02	166.26	-31.84
1946174	196	30503	121.99	265.12	16.35	31.99	52.25	44.63
304623	196	77991	130.63	90.50	17.22	40.63	65.19	35.33
1720882	198	27065	144.83	327.44	11.37	54.83	147.26	57.87
1825306	200	26094	134.23	118.89	1.13	44.23	126.03	-18.54
2627233	200	73755	159.76	293.00	2.80	69.76	132.75	9.14
2995880	202	5311	133.50	332.97	5.20	43.50	164.15	2.45
6529842	202	51347	111.38	123.38	8.00	21.39	200.28	24.19
7085339	202	58586	150.71	222.09	3.43	60.71	140.75	3.30
8860815	205	48508	144.51	221.51	0.86	54.51	122.97	-8.36
14813245	206	39291	129.26	183.81	0.87	39.27	123.13	-23.61
22015370	207	45316	140.06	333.68	16.63	50.05	77.13	41.66
23946320	207	69877	158.59	146.46	11.95	68.59	130.28	47.75
356659	208	1217	123.43	1.96	2.52	33.43	145.81	-24.96
3309224	209	85244	143.16	122.91	17.93	53.16	80.94	29.55
4176635	210	9887	149.96	334.48	6.90	59.96	156.03	23.57
1960937	211	64009	138.28	218.08	5.80	48.27	163.53	10.30
10466084	213	55043	99.91	148.50	8.07	9.91	212.17	20.68
11245943	215	40087	106.00	199.30	0.65	16.00	118.64	-46.77
464325	218	70528	144.91	244.65	6.31	54.91	159.53	17.43
656833	220	51654	110.26	16.69	16.38	20.26	36.73	41.50
993066	220	56189	147.13	180.53	6.72	57.13	158.57	21.43
3800169	222	19869	100.08	20.36	7.41	10.09	207.60	11.91
1694897	223	80497	119.81	35.73	23.34	29.81	100.34	-29.04
3562282	224	18795	132.47	256.39	15.51	42.47	68.90	54.34
751212	225	14587	146.38	57.11	3.69	56.38	144.88	0.92
2949118	225	43476	119.67	104.77	8.56	29.67	193.94	33.95
3114804	226	15661	112.70	251.54	15.09	22.70	32.11	59.40
6569451	226	52678	115.23	352.15	18.69	25.23	54.74	13.23
2216336	227	55886	175.95	281.97	0.33	85.95	122.33	23.11
2641205	227	61524	159.12	162.17	9.89	69.12	142.19	40.71
819596	227	81668	140.73	91.78	20.19	50.73	85.88	9.22
2229429	228	13818	159.03	107.37	0.32	69.03	120.03	6.34
4366149	228	41888	166.30	161.48	4.53	76.30	134.77	18.81
217835	230	4125	139.47	261.38	11.48	49.47	152.74	62.61
530685	230	8098	152.52	68.73	1.43	62.52	126.91	-0.08
1762705	230	23657	132.81	136.95	1.22	42.81	127.17	-19.88
10705083	231	50261	143.28	336.88	19.37	53.28	84.61	17.14
11937623	231	66183	155.13	206.24	8.51	65.13	150.64	34.80
1080294	232	38572	103.01	202.76	1.12	13.01	128.81	-49.68
1665887	233	25394	176.59	45.25	8.01	86.59	126.62	28.10
2254851	233	32987	123.94	51.01	9.74	33.94	191.46	49.56

continued on next page

Table E.1: *continued*

Event Num.	Day	Seconds	Zenith	Azimuth	Rt. Asc.	Declin.	Gal. Long.	Gal. Lat.
1065047	233	57536	102.70	190.23	7.30	12.70	204.47	11.59
1208852	233	59393	154.71	286.15	1.42	64.71	126.55	2.08
3260742	233	85902	146.50	342.09	5.08	56.51	152.90	9.19
4274166	234	84162	112.24	86.57	21.69	22.24	75.31	-22.61
305628	236	20654	133.43	334.00	11.64	43.43	160.34	68.02
756991	238	33968	133.35	168.00	2.55	43.35	141.80	-15.78
1196304	239	61652	109.14	101.48	14.76	19.14	22.04	62.66
3570940	240	3760	130.72	299.68	9.49	40.72	181.27	46.54
11892631	241	19664	131.11	62.73	5.78	41.11	169.67	6.48
16302	241	40039	141.93	62.13	11.49	51.93	149.18	60.83
1936710	241	79745	178.79	106.99	19.56	88.79	121.60	26.90
835594	242	23596	117.34	81.72	5.67	27.34	180.75	-1.83
2830598	244	70706	128.52	105.74	17.32	38.52	62.88	33.77
269918	245	23518	163.65	298.95	15.37	73.65	110.06	39.49
466721	245	26068	113.77	290.60	16.63	23.77	42.52	39.18
2422641	246	25298	152.71	69.82	7.20	62.71	153.53	26.24
2172961	246	62099	158.56	6.02	21.71	68.56	107.02	11.75
4305038	250	19398	146.19	272.21	16.33	56.19	85.99	43.09
5569196	251	12248	102.70	28.37	6.66	12.70	200.28	3.18
14222751	252	21966	118.37	195.80	22.27	28.36	86.08	-23.19
15186523	252	34151	146.92	314.79	17.73	56.92	85.14	31.54
15655178	252	40071	99.64	9.27	15.75	9.64	18.22	45.58
3737276	256	15734	131.63	322.84	12.33	41.63	145.08	74.11
544695	256	44394	165.30	191.79	5.05	75.30	136.82	19.76
1724744	256	59407	123.38	214.15	7.74	33.38	186.59	24.88
2340945	256	67267	146.61	354.72	0.56	56.61	120.44	-6.17
9765274	258	4217	158.35	327.65	8.93	68.35	146.04	36.40
12039520	258	34181	155.54	71.78	10.34	65.54	143.95	44.92
5530269	260	81813	145.77	208.58	14.61	55.77	96.67	55.65
1021539	261	33065	112.53	238.63	23.10	22.53	93.18	-34.17
2376827	261	52028	163.07	203.85	6.70	73.07	141.68	25.18
753861	272	17008	128.98	356.93	11.46	38.98	173.44	68.91
569921	275	1037	122.19	80.04	1.67	32.19	134.78	-29.56
497811	276	23639	139.74	46.47	10.27	49.74	164.96	52.76
392929	277	71350	120.17	53.17	23.18	30.17	98.33	-27.84
82755	279	2063	135.62	149.52	21.59	45.62	90.92	-4.70
2347862	279	31401	165.93	290.69	20.35	75.93	108.98	21.10
112258	279	60044	144.01	160.18	13.03	54.01	119.66	63.05
177052	279	60890	162.07	110.23	16.59	72.07	104.55	35.94
221671	280	9418	131.78	55.25	5.99	41.78	170.24	8.84
265938	280	9983	167.11	141.96	0.36	77.11	121.23	14.33
1001647	280	46749	146.60	142.07	10.60	56.60	152.81	51.92
2312985	280	63564	131.77	350.35	1.39	41.77	129.33	-20.70
1197960	285	18132	156.97	261.56	18.99	66.97	97.57	24.07
1328993	287	30385	114.83	89.85	9.98	24.83	206.32	51.30
911354	288	4683	140.92	45.93	5.81	50.92	161.23	11.69
3464464	289	79430	160.37	69.82	1.11	70.37	124.20	7.53

continued on next page

Table E.1: *continued*

Event Num.	Day	Seconds	Zenith	Azimuth	Rt. Asc.	Declin.	Gal. Long.	Gal. Lat.
6089102	290	26256	143.41	194.69	2.04	53.42	133.47	-8.00
1229085	292	28720	143.66	341.83	17.04	53.66	81.31	37.46
2822463	292	48705	176.85	189.73	8.75	86.85	126.09	28.59
296959	293	5163	150.37	106.38	2.25	60.37	133.06	-0.86
2859301	294	25839	135.19	16.63	14.05	45.19	89.23	66.96
1603543	295	15291	134.98	172.03	0.82	44.98	122.53	-17.89
1292276	295	41331	118.00	279.75	0.89	28.00	123.51	-34.87
1552180	295	44611	142.26	257.01	3.32	52.26	144.60	-4.28
4200466	299	68214	144.00	125.83	18.91	54.00	83.83	21.24
1113492	300	2891	142.73	355.45	9.47	52.73	164.27	44.68
594071	300	54091	131.91	51.80	19.97	41.91	77.10	6.54
2071810	300	72466	147.53	73.20	23.66	57.53	113.41	-4.01
570834	302	23965	138.78	73.30	10.28	48.78	166.37	53.21
675775	302	25347	107.64	23.25	14.00	17.64	5.56	71.54
8301190	303	35604	164.53	32.69	16.29	74.53	107.97	35.95
4894124	304	27502	169.39	146.36	6.53	79.39	134.68	25.72
8452210	304	71447	119.23	56.53	0.75	29.23	121.32	-33.62
8248715	307	82811	156.24	165.34	20.86	66.24	101.90	13.74
21744	308	8809	109.97	43.79	8.42	19.97	204.14	29.24
633532	308	16445	150.31	267.31	19.65	60.31	92.32	17.81
535305	308	40316	156.61	51.47	16.68	66.61	98.00	37.51
7733050	310	7076	97.82	16.00	9.92	7.82	229.40	44.10
2377443	310	39963	156.49	128.74	11.57	66.49	134.54	48.85
149932	311	8939	115.02	334.40	13.28	25.02	13.97	83.94
1231561	311	79400	136.36	162.96	20.33	46.36	82.96	5.64
1236687	311	79464	106.76	266.74	13.43	16.75	342.37	76.90
1695439	311	85155	126.31	69.80	4.15	36.30	161.52	-11.37
5307050	312	43792	135.06	35.72	18.96	45.06	75.08	17.77
10775427	313	25373	97.44	313.99	19.35	7.45	42.92	-3.10
3368378	318	35157	110.40	324.35	21.71	20.40	74.07	-24.07
5602132	318	62817	138.32	110.37	19.68	48.32	81.25	12.35

Appendix F

On Reconstruction

The reconstruction methods currently in use within the collaboration, while sufficient to establish the observation of atmospheric neutrinos, are in many ways rather crude. There are at least three separate areas in which the reconstruction could be improved, and it is hoped that work on these topics would result in significant improvements in background rejection, in signal efficiency, and in the extraction of physically useful information from the detector. Considerable amounts of work will be required to produce such improvements, but it is hoped that with the establishment of the detector as a working instrument and with the increasing size of the collaboration some resources may be available for these projects. This Appendix is included in the hope that it may prove useful as a guide for future development work.

F.1 Calculation of the Likelihood

The heart and soul of any AMANDA event reconstruction, whether of muons or showers, whether using a maximum likelihood method or a Bayesian, is the likelihood function. This function calculates, given a test hypothesis (a track or a shower, say), the probability that the hypothesis would generate the observed detector response. Although AMANDA reconstruction is by necessity fairly sophisticated compared to many astrophysical detectors, our calculation of the likelihood function remains *ad hoc* in many ways. It seems that progress in this area will likely result in large improvements in physics analyses. I will therefore attempt to give a proper mathematical treatment of the likelihood function, and discuss a number of areas where improvements could be made.

F.1.1 The Math

AMANDA events are reconstructed by finding the set of parameters \vec{X} (for an infinite minimum ionizing muon, $\vec{X} = \{x, y, z, \theta, \phi\}$) which maximizes the likelihood \mathcal{L} of producing the observed detector response. If the responses of the phototubes are independent (that is, neglecting cross talk), \mathcal{L} is the product of the likelihoods of the responses of the individual phototubes

$$\mathcal{L} = \prod_{i=1}^{N_{\text{OM}}} \mathcal{L}_i, \quad (\text{F.1})$$

so the problem is to calculate the likelihood of the response observed from an individual tube.

Given the hypothesis \vec{X} and knowledge of the relevant physical processes (the optical properties of the ice, tube glass and gel, the quantum efficiency of the tube, and so forth), we can calculate $p(\vec{\tau} | \vec{X})$, the probability of having n photoelectrons produced at times $\vec{\tau} = \{\tau_1, \dots, \tau_n\}$. However, what is observed, at least with the present AMANDA electronics, is a single amplitude a and a series of leading edge times \vec{t} and pulse durations \vec{d} . These observations depend on the true series of photoelectrons via the hardware response function $p_i(a, \vec{t}, \vec{d} | \vec{\tau})$, which is unique to each channel and includes effects like noise. In general there are multiple series of photoelectrons which could generate any given observed tube response. The likelihood function for the tube response is thus an integral over the possibilities each weighted by its probability given the parameters of the hypothesis:

$$\mathcal{L}_i(a_i, \vec{t}_i, \vec{d}_i | \vec{X}) = \int d\vec{\tau} p_i(a_i, \vec{t}_i, \vec{d}_i | \vec{\tau}) p(\vec{\tau} | \vec{X}).$$

The integral over $\vec{\tau}$ is an integral over time series, but it can be made tractable by first summing over the number of photoelectrons composing the series

$$\mathcal{L}_i(a_i, \vec{t}_i, \vec{d}_i | \vec{X}) = \int d\vec{\tau}_n \sum_{n=0}^{\infty} p_i(a_i, \vec{t}_i, \vec{d}_i | \vec{\tau}_n) p(\vec{\tau}_n | n, \vec{X}) p(n | \vec{X}) \quad (\text{F.2})$$

(where $\vec{\tau}_n$ denotes a vector of n times) and truncating the series at some high value of n .

The standard AMANDA reconstruction technique takes a rather cavalier approach to Eq. F.2. We set the last term on the right hand side to unity for all n and set $n = 1$ in the second term, and for computational simplicity we use an analytic approximation to the form of $p(\vec{\tau}_1 | n = 1, \vec{X})$ predicted by Monte Carlo¹. We use a hardware response function which is flat in \vec{d} and a (i.e., we ignore the

¹No analytic solution for the arrival times of photons propagating through a medium intermediate between the scattering and non-scattering regimes is known, so the solution must be determined by numerical simulation.

amplitude and pulse length information completely) and for the resulting $p(\vec{t} | \vec{\tau}_1)$ we simply patch $p(\vec{\tau}_1)$ with a half-Gaussian at early times and introduce a minimum probability, a noise floor, to account for random noise hits. Somewhat surprisingly, the resulting likelihood function is able to reconstruct tracks sufficiently well for neutrino identification, but it seems clear that improvements can be made.

The problem of formulating the hardware response function $p(a, \vec{t}, \vec{d} | \vec{\tau})$ is of course very difficult. It may be that the integral over $\vec{\tau}$ called for in Eq. F.2 is computationally too intensive to be practical. The function is in principle not invertible (information is necessarily lost due to finite hardware response), but it may be possible, especially with the more advanced optical modules of AMANDA-II and IceCube, to at least partially invert the function — to allow the integral over $\vec{\tau}$ to be truncated by the exclusion of pulse trains which are clearly inconsistent with the observed response. Deconvolution methods (see, e.g., [102] for an impressive demonstration of the degree to which hardware response functions can be corrected for) may be useful in this effort, although of course they rely on precise knowledge of the hardware response and require a relatively high waveform sampling rate relative to the width of a pulse. At the least, these methods may be useful in obtaining data from the analog electrical strings of B10 of comparable quality to that from the modern strings, if flash ADCs are installed on the inner channels.

A proper calculation of the likelihood is of course far more involved than the methods presently in use, and the availability of computational resources may make a full treatment of every event impractical. Nevertheless, one can imagine a reconstruction scheme in which faster algorithms are used to progressively reduce the data set to the point that more sophisticated techniques become feasible.

F.1.2 Parametrization

At present we use an analytic parametrization to $p(\vec{\tau}_1 | n = 1, \vec{X})$. There have in fact been two such parametrizations suggested, the more commonly used one being a gamma distribution with an exponential absorption term, originally due to D. Pandel and named for him [103]. The other parametrization is an F-function with an exponential tail, and was developed by A. Bouchta [54].

Parametrizations offer one advantage over numerical descriptions of the timing distributions:

they can be calculated and differentiated on the fly, and therefore do not require large amounts of computer memory. However, parametrizing a distribution necessarily introduces inaccuracies, and with the increase in available memory and the use of on-the-fly interpolation techniques, these parametrizations no longer seem as worthwhile. Software which calculates the timing probabilities via direct references to the tabulated numerical distributions is now being developed, and promises to improve the accuracy, and possibly the speed, of the reconstruction.

F.1.3 Multiple Photoelectrons

At present we use only the first leading edge time t_1 recorded from each channel, and assume that it was produced by a single photoelectron sampled from the distribution² $p(t | n = 1, \vec{X})$, as described above³. One attempt to better model the case of multiple photons is to use instead of the single-photon time distribution $p(t | n = 1, \vec{X})$ the probability

$$p_1^{(\text{MPE})}(t | n, \vec{X}) = n p(t | \vec{X}) \left(\int_t^\infty dt p(t | \vec{X}) \right)^{n-1} \quad (\text{F.3})$$

that the first of n photons would arrive at time t ; this is referred to as the “multiphoton” or “MPE” distribution. The measured ADC response divided by that of a single-photoelectron pulse, rounded to the nearest integer, is assumed to be the number of photoelectrons n . A further refinement, more in line with the mathematically correct calculation of the likelihood, is to use the “Poisson saturated distribution”

$$p_1^{(PS)}(t | \vec{X}) = \sum_{n=1}^{\infty} p(n | \mu) p_1^{(\text{MPE})}(t | n, \vec{X})$$

a sum of the MPE distributions for various possible numbers of photoelectrons each weighted by the Poissonian probability

$$p(n | \mu) = \frac{\mu^n e^{-\mu}}{n!}$$

²Here $p(t | n = 1, \vec{X})$ will be used to denote $p(\vec{\tau}_1 | n = 1, \vec{X})$ modified by the addition of a half-Gaussian and a noise floor, as described in Section F.1.1.

³In some analyses subsequent leading edges are also used, with the overall likelihood being the product

$$\mathcal{L}_i = \prod_{j=1}^{N_{\text{pulse}}} p(t_j | n = 1, \vec{X}).$$

of observing n photoelectrons given an expectation of μ photoelectrons based on the track hypothesis \vec{X} .

Both of these methods begin to address the inaccuracy of using the single-photoelectron timing distribution without taking into account the number of photoelectrons, that is, the number of times the distribution was sampled. However, neither the MPE nor the PS techniques are now in common use for muon reconstruction. Furthermore, there is the problem that raising parametrizations of the single-PE distribution to high powers compounds the inaccuracies of the fit. It is possible to calculate Eq. F.3 numerically from the original numerical description of the single-PE time distribution, and this may improve the accuracy of the reconstruction.

F.1.4 Hit Probabilities

“Is there any point to which you would wish to draw my attention?”
 “To the curious incident of the dog in the night-time.”
 “The dog did nothing in the night-time.”
 “That was the curious incident,” remarked Sherlock Holmes.

Sir Arthur Conan Doyle
Silver Blaze
The Memoirs of Sherlock Holmes

Probably the most striking problem in our current scheme is that we ignore the probability $p(n | \vec{X})$ of a tube receiving any given number n of photoelectrons. This term is omitted from the product in Eq. F.2, which is equivalent to setting it to unity for all n . This means that a tube which fires far from the hypothetical track is not taken as evidence against the track, nor is a tube that fails to register a hit despite being very close to the track. In most analyses the failure to make use of this information is partially corrected in the analysis phase by making cuts on the smoothness parameter, which attempts to measure the self-consistency of the observed hit topology.

Some attempts have been made to incorporate this information into the reconstruction, but all attempts to date have used a fit to the expected number of photoelectrons, and have failed to account for the uncertain hardware response in any systematic fashion. The problem of inaccuracies in this probability function are particularly important because most modules fail to fire in any given event, and so any systematic errors in the probability that a module not fire introduced by incorrect modelling of the tails of the photon transport function will be raised to large powers.

F.2 Expanding the Likelihood Space

Probably the most important direction for further work is the expansion of the likelihood space itself. At present in the muon analysis, the likelihood space is five-dimensional: a three-dimensional vertex and two angles to describe an infinite, minimum-ionizing muon track. In some analyses a separate one-dimensional fit is done based on the results of the positional fit to attempt to reconstruct the energy. A fit to a “shower” (i.e., point-like) hypothesis is also frequently done, and the results of the shower and muon fits can be compared offline during the data analysis. However, the rather *ad hoc* manner in which these comparisons is done is unsatisfying; the reconstruction software ought to be able to determine whether an event is a muon track or a pointlike emission without user intervention.

The framework of the Bayesian reconstruction gives us the ability to expand the scope of the reconstruction in a straightforward fashion by comparing various models of the event. For example, the reconstruction program could determine whether the track was a throughgoing, starting, or stopping muon or a shower, including fitting the energy of the track and any large stochastic losses along the track. This type of fit is difficult in a frequentist framework, where the hypothesis must be specified *a priori*, and a hypothesis with more free parameters will almost always produce a solution with a better likelihood. There is no leverage to guide the reconstruction to simpler models, as the models are not directly compared.

In the Bayesian framework, however, it is quite natural to compare models. Ockham’s razor manifests itself quite naturally through the normalization of the prior functions of the different models — the posterior probability of the more detailed hypothesis will suffer a penalty for having additional parameters, and the simpler model will be preferred unless the more complicated one results in a significantly better fit to the data [104].

Consider the example of a muon event which is to be fit to either a throughgoing (i.e., infinite) or a stopping track. The stopping track has one more parameter than the infinite track hypothesis: the position of the muon’s decay. The frequentist maximum likelihood fit is able to fine-tune the additional parameter, so that if even a single tube near the end of the track failed to fire, the stopping muon hypothesis will be strongly favored. Intuitively, we know that it is highly unlikely that a muon

would decay just at the edge of the detector, and we would want much stronger evidence for this hypothesis than the failure of a single tube to fire. In the Bayesian approach, the normalization of the prior produces the desired preference for the simpler model. The prior for each model is normalized to unity, and so the prior probability for the muon to stop in any given differential element of a track of length l is reduced by a factor dl/l (assuming a uniform prior). The evidence from the data that the muon did indeed stop would need to be sufficiently compelling to outweigh this factor for the stopping muon hypothesis to be preferred.

Furthermore, the prior naturally encodes knowledge about the relationships between parameters of the fit. For example, a high energy track is less likely to decay and more likely to produce a stochastic deposition of energy than is a low energy track, and this information can be incorporated precisely into the reconstruction program via a multidimensional prior function. The Bayesian approach to event reconstruction thus gives a very straightforward prescription for greatly increasing both the sophistication and the simplicity of the analysis.

It can be argued that it is not necessary to approach these questions from a Bayesian perspective. One can certainly reconstruct contained muons, for example, by fitting to an infinite muon track and then inventing some more or less complicated series of quality parameters to attempt to identify those events which are likely to be contained muon tracks. But there are two disadvantages to this strategy: first, it is extremely time consuming, and second, there is no clear prescription for developing such parameters (or *statistics*, in the technical sense), and there thus will always be an element of arbitrariness in the analysis. Moreover, the full information about the event is contained in the likelihood function and the prior (properly formulated), whereas cut parameters and other estimators can never contain the full information.

What we have is indeed the fundamental question of Bayesian vs. frequentist theories of inference — whether it is better to directly approach the question of the superior model by calculating which hypothesis is more probable, given the data, or rather to sequentially assume the truth of various hypotheses and then to attempt to develop some collection of statistics of varying complexity to assess the self-consistency of the hypotheses. Note, however, that the classical objections to the Bayesian theory are completely irrelevant to this problem. We are not attempting to create a uniform

or unbiased prior in this case; the prior functions are highly nonuniform, and they are completely specified. We know the distributions of cosmic ray muons and atmospheric neutrinos to rather precise detail; their angular distributions and energy spectra are known from other experiments, as is the behavior of high energy muons travelling through matter. It is difficult to imagine a problem which is more perfectly suitable for the application of Bayesian techniques.

F.3 Minimization

The problem of reconstructing an event is a two-step process. The first step, as described above, is to specify the likelihood function and the prior — to describe what physical process may occur and how the detector responds to them. The second step is to determine, for the specific event, which of the many possible physical hypotheses was actually responsible for the recorded event. This step is accomplished by some algorithm designed to find the global optimum of the posterior probability function, for example as described in Sec. 5.3.

Optimization problems constitute a major focus of inquiry in some fields, although physicists seem to be rather slow in importing new techniques. Many optimization techniques from other disciplines are intended to address combinatorial problems (i.e., those defined over a discrete solution space, such as the travelling salesman), but a number of these can be adapted to continuous or mixed optimization problems.

Continuous minimization problems are often solved by a two-level process. The top level of the process is the *global minimizer*. The global minimizer is an algorithm which attempts to ensure that when a minimum is found it is in fact the global minimum, not merely a local minimum in the likelihood space. The global minimizer is typically a fast but imprecise algorithm, which only approximately locates minima. When the global minimizer finds an approximate solution, it feeds the solution to a *local minimizer*. The local minimizer takes this initial guess and attempts to refine it, improving on the precision of the solution. The local minimizer then returns the solution to the global minimizer, which may decide to continue searching for better minima. Local minimization is typically a time consuming operation, and so one desideratum for a global minimizer is that it call the local minimizer as infrequently as possible.

There are many global and local minimizers on the market. The local minimizer currently in

use in the collaboration is Powell's minimizer from [70], although there is some discussion of using the CERNlib MINUIT package for minimization. There are at least two global minimization algorithms in use. The first is simply to make an analytic first guess and give the event to the local minimizer. This has the advantage of being extremely fast, but is unreliable for two reasons. First, the initial guess may not be near the global minimum, and second, there is no guarantee even if the initial guess *is* near the mark that the local minimizer will converge correctly to that minimum. It is possible for a local minimizer to actually escape from the well in which it is started and converge to a completely different minimum, although this should be a rare case if the initial guess is reasonably accurate. The greater danger, especially if the likelihood function has multiple minima, is that the initial guess will be near a minimum that is shallower than the global minimum.

The second strategy is to restart the local minimizer a certain number of times, giving it a random initial guess at each restart. This is perhaps the simplest global minimization strategy that can be implemented which guards against the possibility of multiple minima. Nevertheless, restarting, say, five times is no guarantee that different solutions will be found even if there are multiple minima; it is possible for the five minimizations all to converge to the same solution, even if it is not the global minimum. As the number of restarts increases, of course, the chance of this happening is reduced. However, particularly in a high-dimensional likelihood space, it may take very many restarts and thus a very large amount of CPU time to guarantee with any reasonable confidence that the global minimum has been found. The random multistart algorithm is thus an effective but slow and rather crude global minimizer, which has been sufficient for simple five-dimensional AMANDA fits but which may become prohibitively expensive if more parameters (energy, track length, etc.) are to be fit.

Another approach which was investigated in AMANDA [54] but which has fallen by the wayside was the use of the Metropolis algorithm [105], also known as simulated annealing. Simulated annealing is a global minimization algorithm developed for combinatorial problems, but which was adapted to the continuous problem of track fitting by using a downhill simplex algorithm, as described in [70]. This algorithm is rather different from those described above in that it does not rely on a local minimizer but rather settles in on a (hopefully global) minimum itself. The basic idea is that the algorithm moves through parameter space, at each point being offered a possible next move. The algorithm decides

whether to take the proffered move based on the current ‘temperature’, an externally set parameter which is gradually decreased. Moves downhill in the likelihood (i.e., those which apparently are toward a minimum) are always accepted; uphill moves are accepted with probability $e^{-\Delta L/T}$. The idea is that as the temperature is reduced, the algorithm will settle into the global minimum, just as a crystalline solid will settle into its lowest energy state as its temperature is reduced after being heated. The primary drawback to simulated annealing is that the speed and dependability of the algorithm depends on the initial temperature and the rate of cooling, which must be optimized by trial-and-error to best fit the characteristics of the typical likelihood functions being minimized. Nevertheless, this approach may prove to be useful for AMANDA in the future.

One other global minimizer in common use in other disciplines takes its inspiration from nature, this time from biology rather than physics. Genetic algorithms are also designed for combinatorial problems, but they seem easily adaptable to the continuous case by means of Grey encoding⁴ [108]. In this algorithm, possible solutions are represented by strings of bits in analogy to chromosomes, and a population of solutions is allowed to evolve and caused to mutate, leading eventually to the maximally adaptive (i.e., optimal) solution. Genetic algorithms may be worth investigating, though they seem to be falling into disfavor in the disciplines most interested in optimization.

Finally, a global minimization algorithm which seems very well suited to AMANDA’s requirements is the Continuous Reactive Tabu Search (C-RTS) [106, 107, 108]. Tabu is the name of a general approach to combinatorial optimization, where the algorithm is prevented from falling into local minima by declaring these solutions temporarily off-limits (tabu). Battiti and Tecchioli have adapted this strategy to continuous problems by use of Grey encoding, using the Tabu algorithm to search for minima and then calling a local minimizer to refine the solutions. As the algorithm explores the search space it adaptively tunes its search strategy, optimizing the strategy to the shape of the likelihood function for the particular event. Battiti and Tecchioli have demonstrated impressive results for many continuous functions, and I believe that AMANDA could drastically improve its reconstruction speed and reliability by using this method. The amount of time spent searching for

⁴Grey encoding is a transformation of the normal binary representation of numbers such that a one-bit change of the number produces an adjacent number (as opposed to the binary case, where changing the first digit produces a very large change, the last digit a small change).

the minimum could also be directly set, allowing the technique to be used both for relatively quick filtering and for very thorough searches of the parameter space. The algorithm may also be useful because in the course of searching it adaptively samples the entire likelihood function, the results of which sampling could be used to roughly integrate the likelihood over various parameters.