Searching for Quantum Gravity with High-energy Atmospheric Neutrinos and AMANDA-II

by

John Lawrence Kelley

A dissertation submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

(Physics)

at the

University of Wisconsin – Madison

2008

Searching for Quantum Gravity with High-energy Atmospheric Neutrinos and AMANDA-II

John Lawrence Kelley

Under the supervision of Professor Albrecht Karle

At the University of Wisconsin – Madison

The AMANDA-II detector, operating since 2000 in the deep ice at the geographic South Pole, has accumulated a large sample of atmospheric muon neutrinos in the 100 GeV to 10 TeV energy range. The zenith angle and energy distribution of these events can be used to search for various phenomenological signatures of quantum gravity in the neutrino sector, such as violation of Lorentz invariance (VLI) or quantum decoherence (QD). Analyzing a set of 5511 candidate neutrino events collected during 1387 days of livetime from 2000 to 2006, we find no evidence for such effects and set upper limits on VLI and QD parameters using a maximum likelihood method. Given the absence of new flavor-changing physics, we use the same methodology to determine the conventional atmospheric muon neutrino flux above 100 GeV.

Albrecht Karle (Adviser)

The collaborative nature of science is often one of its more enjoyable aspects, and I owe a debt of gratitude to many people who helped make this work possible.

First, I offer my thanks to my adviser, Albrecht Karle, for his support and for the freedom to pursue topics I found interesting (not to mention multiple opportunities to travel to the South Pole). I would like to thank Gary Hill for numerous helpful discussions, especially about statistics, and for always being willing to talk through any problem I might have. Francis Halzen sold me on IceCube during my first visit to Madison, and I have learned from him never to forget the big picture.

Many thanks to Teresa Montaruli for her careful analysis and insightful questions, and to Paolo Desiati for listening patiently and offering helpful advice. I thank Dan Hooper for suggesting that I look into the phenomenon of quantum decoherence and thus starting me off on this analysis in the first place. And I thank Kael Hanson for providing fantastic opportunities for involvement in IceCube hardware and software development, and Bob Morse for a chance to explore radio detection techniques.

I offer my thanks as well to my officemate and friend, Jim Braun, for a collaboration that has made the past years much more enjoyable. The $O(1000)$ pots of coffee we have shared were also rather instrumental in fueling this work.

I would like to thank my parents, James and Lorraine, for their unflagging encouragement through all of my career-related twists and turns. Finally, I offer my deepest gratitude to my partner, Autumn, for her support, flexibility, and understanding — and for leaving our burrito- and sushi-filled life in San Francisco to allow me to pursue this degree. May the adventures continue.

Contents

Chapter 1

Introduction

Our Universe is a violent place. Stars burn through their elemental fuel and explode. Matter spirals to its doom around supermassive black holes at the center of galaxies. Space still glows in every direction from the primordial explosion of the Big Bang.

Born from these inhospitable conditions are the neutrinos. Anywhere nuclear reactions or particle collisions take place, neutrinos are likely to be produced — in the Big Bang, in stars, and even in our own nuclear reactors. The neutrino, having no electric charge, interacts only via the weak force, and thus normal matter is nearly transparent to it. Trillions of neutrinos pass through our bodies every second, and we never notice.

Wolfgang Pauli postulated the existence of the neutrino in 1933 to solve a problem with missing energy in radioactive beta decay [1]. Twenty years later, Reines and Cowan first detected neutrinos by placing liquid scintillator targets next to the Hanford and Savannah River nuclear reactors [2]. Today, we have also detected neutrinos from our Sun ([3]; see also fig. 1.1), from nuclear decay deep in the Earth [4], and even from a nearby supernova [5, 6]. Figure 1.2 shows the fluxes and energy ranges spanned by known and hypothetical neutrino sources.

Another product of the high-energy processes in the universe are cosmic rays: high-energy protons and atomic nuclei that are accelerated to energies far beyond that of any particle accelerator on Earth. These cosmic rays bombard the Earth continuously, producing yet more neutrinos that rain down upon us from high in our atmosphere.

Given a large enough target, we can detect these high-energy *atmospheric neutrinos*. The AMANDA-II neutrino detector employs the huge ice sheet at the South Pole as such a target, and

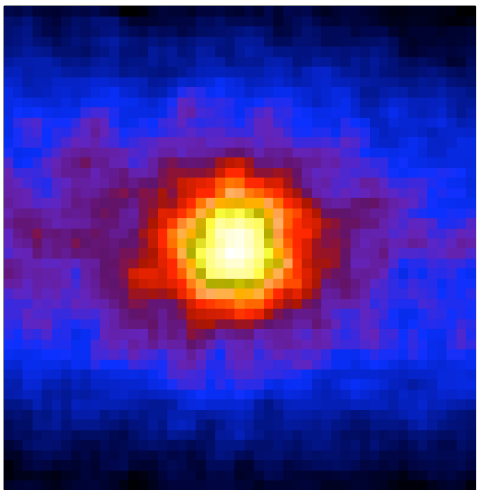

Figure 1.1: "Picture" of the Sun in neutrinos, as seen by the Super-Kamiokande neutrino detector. Image credit: R. Svoboda and K. Gordan.

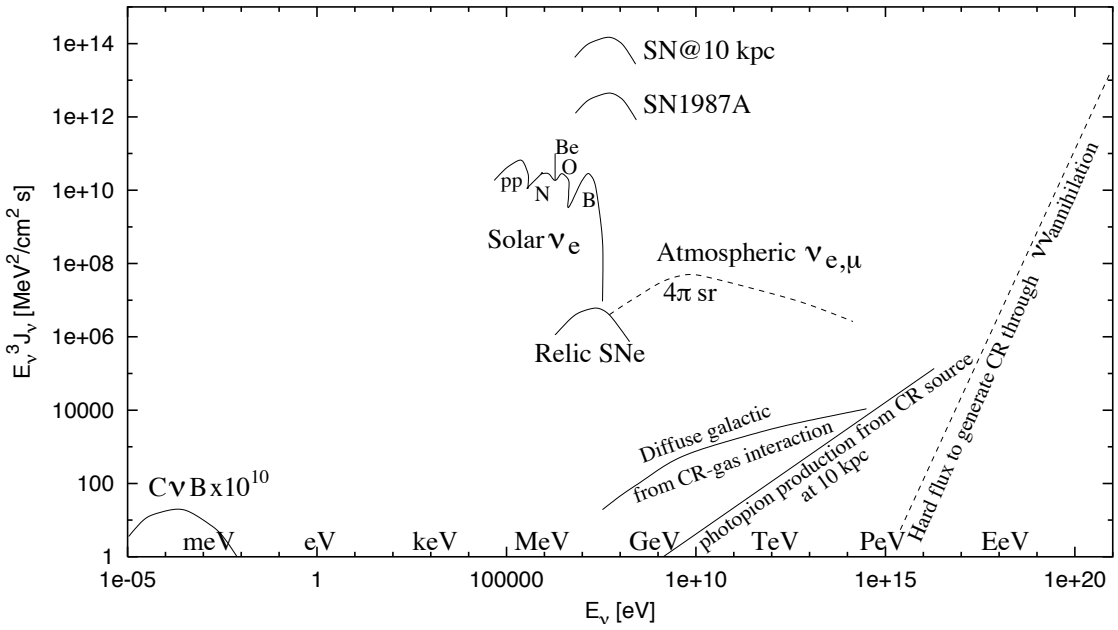

Figure 1.2: Neutrino fluxes as function of energy (multiplied by E^3 to enhance spectral features) from various sources, including the cosmic neutrino background from the Big Bang (CνB), supernovae neutrinos (SN ν), solar neutrinos, and atmospheric neutrinos (from [7]).

uses sensitive light sensors deep in the ice to detect the light emitted by secondary particles produced when a neutrino occasionally hits the ice or the bedrock. AMANDA-II accumulates such neutrinos at the rate of about 16 per day, about four of which are sufficiently high quality to use for an analysis.

Why study these neutrinos? Nature provides a laboratory with energies far above what we can currently produce on Earth, and studying these high-energy neutrinos can possibly reveal hints of surprising new physical effects. We know that our theories of gravity and quantum mechanics are mutually incompatible, but we have no theory of *quantum gravity* to unify the two. At high enough energies, we should be able to probe effects of quantum gravity, and neutrinos may prove crucial to this effort.

In this work, we examine atmospheric neutrinos detected by AMANDA-II from the years 2000 to 2006 for evidence of quantum gravitational effects, by determining their direction and approximate energy. We have found no evidence for such effects, leading us to set limits on the size of any violations of our existing theories. Finally, we determine the atmospheric neutrino flux as a function of energy, extending measurements by other neutrino experiments.

Chapter 2

Cosmic Rays and Atmospheric Neutrinos

2.1 Cosmic Rays

Cosmic rays are protons and heavier ionized nuclei with energies up to 10^{20}eV that constantly bombard the Earth's atmosphere. Exactly where they come from and how they are accelerated to such incredible energies are both open questions. Nearly 100 years after Victor Hess's balloon experiments in 1912 showed that cosmic rays come from outer space [8], we still do not know their source. One of the main difficulties is that the magnetic field of the Galaxy scrambles any directional information that might point back to a source. Still, all but the highest-energy cosmic rays come from within our Galaxy, and the expanding shocks around supernovae remnants are a likely candidate acceleration site [9]. Figure 2.1 shows a composite image of the expanding shock wave around the Tycho supernova remnant.

The cosmic ray energy spectrum is a power law with differential flux approximately proportional to $E^{-2.7}$ [11]. Figure 2.2 shows measurements of the flux from both direct measurements (space- and balloon-based instruments) and indirect measurements (air shower arrays). Above about 10^6 GeV, the spectrum steepens to approximately $dN/dE \propto E^{-3}$, a feature known as the knee. The exact mechanism for this transition is unknown, but one possibility is a rigidity-dependent cutoff of the spectrum as cosmic rays diffuse out of the Galaxy [12].

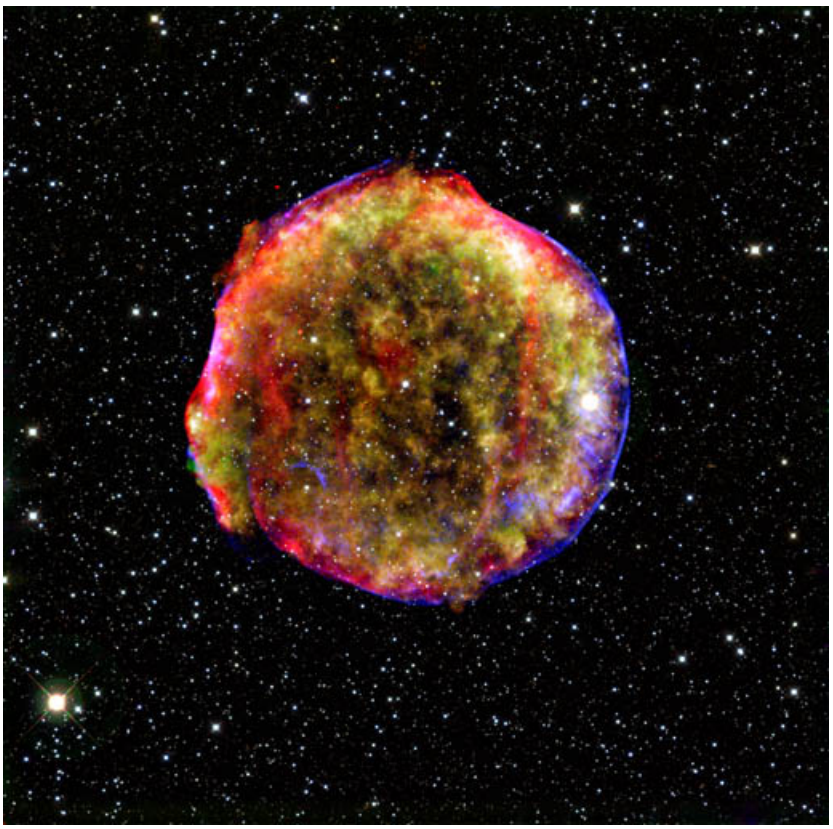

Figure 2.1: X-ray and infrared multi-band image of Tycho's supernova remnant (SN 1572), a possible source of galactic cosmic ray acceleration [10]. Image credit: MPIA/NASA/Calar Alto Observatory.

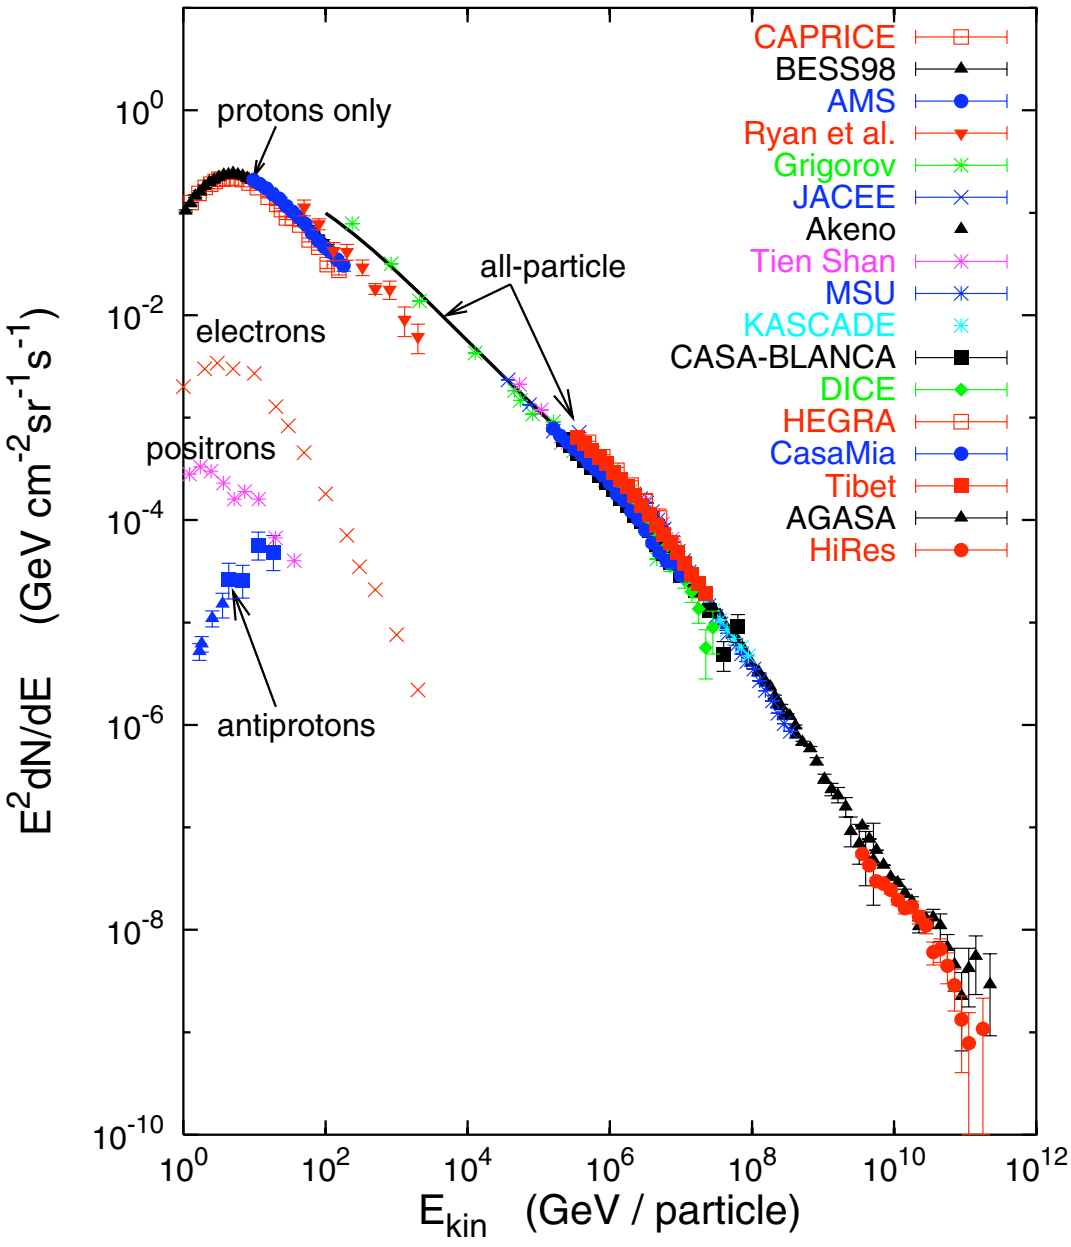

Figure 2.2: The cosmic ray energy spectrum as measured by various direct and indirect experiments, from [13]. The flux has been multiplied by E^2 to enhance features in the steep power-law spectrum.

2.2 Atmospheric Neutrinos and Muons

2.2.1 Production

As cosmic rays interact with air molecules in the atmosphere, a chain reaction of particle production (and decay) creates an extensive air shower of electrons, positrons, pions, kaons, muons, and neutrinos. Atmospheric neutrinos are produced through hadronic interactions generating charged pions and kaons, which then decay into muons and muon neutrinos:

$$\pi^+(K^+) \rightarrow \mu^+ + \nu_\mu \tag{2.1}$$

$$\pi^-(K^-) \rightarrow \mu^- + \bar{\nu}_\mu \ . \tag{2.2}$$

Some of the muons produced in the decay will also eventually decay via exchange of a $W^\pm$ boson, producing both electron and muon neutrinos:

$$\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e \tag{2.3}$$

$$\mu^- \rightarrow \nu_\mu + e^- + \bar{\nu}_e \ . \tag{2.4}$$

However, many of these atmospheric muons will survive to ground level and, depending on their energy, can penetrate kilometers into the Earth before decaying. The process of atmospheric muon and neutrino production is shown schematically in fig. 2.3.

2.2.2 Energy Spectrum and Angular Distribution

Atmospheric muon neutrinos dominate over all other neutrino sources in the GeV to TeV energy range. The flux of atmospheric electron neutrinos is over one order of magnitude smaller than the flux of muon neutrinos at these energies [15]. While the flux of parent cosmic rays is isotropic, the kinematics of the meson interaction and decay in the atmosphere alters the angular distribution of atmospheric ν_μ to a more complicated function of the zenith angle.

While elaborate three-dimensional calculations exist for the expected flux of atmospheric neutrinos [16, 17], an approximate analytic formulation is given by Gaisser [11]:

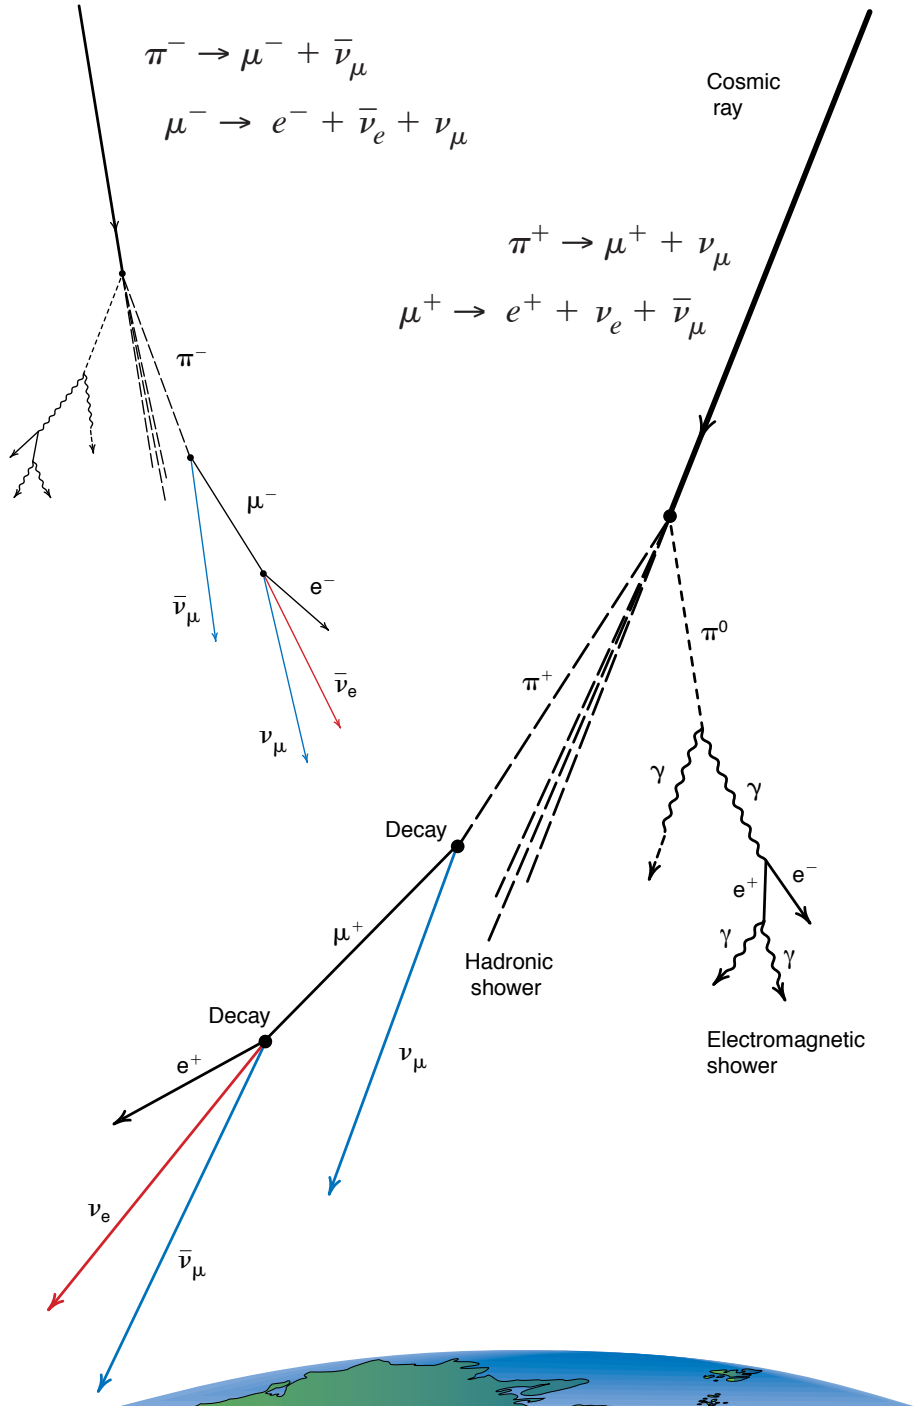

Figure 2.3: Atmospheric muon and neutrino production (from [14]).

$$\frac{dN}{dE_\nu} = C\left(\frac{\mathcal{A}_{\pi\nu}}{1+\mathcal{B}_{\pi\nu}\ \cos\theta^*\ E_\nu/\epsilon_\pi} + 0.635\ \frac{\mathcal{A}_{K\nu}}{1+\mathcal{B}_{K\nu}\ \cos\theta^*\ E_\nu/\epsilon_K}\right)\ , \tag{2.5}$$

where

$$\mathcal{A}_{\pi\nu} = Z_{N\pi}\frac{(1-r_\pi)^\gamma}{\gamma+1} \tag{2.6}$$

and

$$\mathcal{B}_{\pi\nu} = \frac{\gamma+2}{\gamma+1}\ \frac{1}{1-r_\pi}\ \frac{\Lambda_\pi - \Lambda_N}{\Lambda_\pi \ln(\Lambda_\pi/\Lambda_N)}\ . \tag{2.7}$$

Equivalent expressions hold for $\mathcal{A}_{K\nu}$ and $\mathcal{B}_{K\nu}$. In the above, γ is the integral spectral index (so $\gamma \approx 1.7$); Z_{ij} is the spectral-weighted moment of the integral cross section for the production of particle j from particle i; Λ_i is the atmospheric attenuation length of particle i; ϵ_i is the critical energy of particle i, at which the decay length is equal to the (vertical) interaction length; r_i is the mass ratio m_μ/m_i; and $\cos\theta^*$ is not the zenith angle θ at the detector, but rather the angle at the production height in the atmosphere.

The cosine of the atmospheric angle is roughly equal to that of the zenith angle for $\cos\theta \geq 0.5$. For steeper angles, we have a polynomial parametrization of the relation that averages over muon production height [18],

$$\cos\theta^* = \sqrt{\frac{\cos^2\theta + p_1^2 + p_2(\cos\theta)^{p_3} + p_4(\cos\theta)^{p_5}}{1 + p_1^2 + p_2 + p_4}} \tag{2.8}$$

where the fit constants for our specific detector depth are given in table 2.1.

Table 2.1: Fit parameters for the $\cos\theta^*$ correction (eq. 2.8), from [18].

p_1	0.102573
p_2	-0.068287
p_3	0.958633
p_4	0.0407253
p_5	0.817285

While significant uncertainties exist in some of the hadronic physics (especially production of

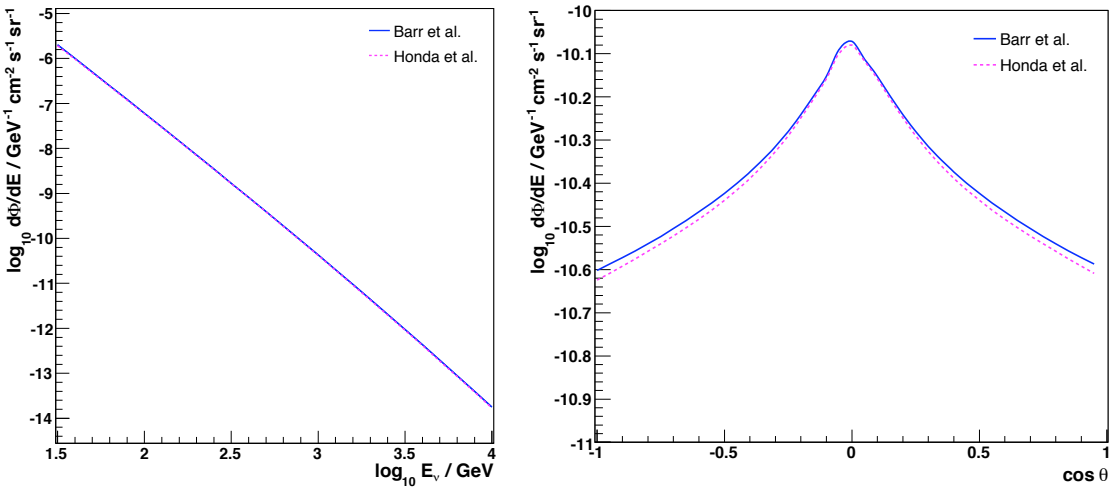

Figure 2.4: Predicted atmospheric neutrino flux as a function of energy and zenith angle, extended to high energies with the Gaisser parametrization. The flux vs. energy (left) is averaged over all angles, while the flux vs. zenith angle (right) is at 1 TeV.

$K^{\pm}$ and heavier mesons), eq. 2.5 is a useful parametrization of the flux in the energy range where neutrinos from muon decay can be neglected (E_ν of at least a few GeV, and higher for horizontal events).

We can also use eq. 2.5 to extend the detailed calculations of Barr *et al.* [16] and Honda *et al.* [17] to energies of 1 TeV and above, by fitting the parameters in an overlapping energy region (below 700 GeV). We show the extended fluxes for each of these models in fig. 2.4 as a function of energy and zenith angle.

2.3 Neutrino Oscillations

If neutrinos are massive, their mass eigenstates do not necessarily correspond to their flavor eigenstates. As we will show, this implies that neutrinos can change flavor as they propagate, and a ν_μ produced in the atmosphere may appear as some other flavor by the time it reaches our detector.

In general, there exists a unitary transformation U from the mass basis to the flavor basis. For oscillation between just two flavors, say ν_μ and ν_τ, the transformation can be represented as a rotation matrix with one free parameter, the mixing angle $\theta_{\rm atm}$:

$$\begin{pmatrix} \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} \cos\theta_{\rm atm} & \sin\theta_{\rm atm} \\ -\sin\theta_{\rm atm} & \cos\theta_{\rm atm} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} . \tag{2.9}$$

For free particles propagating in a vacuum, the neutrino mass (energy) eigenstates evolve according the equation

$$\begin{pmatrix} \nu_1(t) \\ \nu_2(t) \end{pmatrix} = \begin{pmatrix} e^{-iE_1 t} & 0 \\ 0 & e^{-iE_2 t} \end{pmatrix} \begin{pmatrix} \nu_1(t=0) \\ \nu_2(t=0) \end{pmatrix} . \tag{2.10}$$

Combining equations 2.9 and 2.10, and using the approximation that the mass of the neutrino is small compared to its momentum (so that $E_i \approx p + \frac{{m_i}^2}{2p}$), we find

$$\begin{pmatrix} \nu_\mu(t) \\ \nu_\tau(t) \end{pmatrix} = U_f(t) \begin{pmatrix} \nu_\mu(t=0) \\ \nu_\tau(t=0) \end{pmatrix} , \tag{2.11}$$

where the time-evolution matrix $U_f(t)$ in the flavor basis is given by

$$U_f(t) = \begin{pmatrix} \cos^2\theta_{\rm atm}\ e^{-i\frac{{m_1}^2 t}{2p}} + \sin^2\theta_{\rm atm}\ e^{-i\frac{{m_2}^2 t}{2p}} & \cos\theta_{\rm atm}\sin\theta_{\rm atm}\ (e^{-i\frac{{m_1}^2 t}{2p}} - e^{-i\frac{{m_2}^2 t}{2p}}) \\ -\cos\theta_{\rm atm}\sin\theta_{\rm atm}\ (e^{-i\frac{{m_1}^2 t}{2p}} - e^{-i\frac{{m_2}^2 t}{2p}}) & \cos^2\theta_{\rm atm}\ e^{-i\frac{{m_1}^2 t}{2p}} + \sin^2\theta_{\rm atm}\ e^{-i\frac{{m_2}^2 t}{2p}} \end{pmatrix} . \tag{2.12}$$

By squaring the appropriate matrix element above, this evolution equation can easily be used to obtain the probability that a muon neutrino will oscillate into a tau neutrino. Conventionally, the propagation time is replaced by a propagation length L, and the momentum can be approximated by the neutrino energy E, resulting in the following expression for the survival (non-oscillation) probability:

$$P_{\nu_\mu \to \nu_\mu} = 1 \ - \ \sin^2 2\theta_{\rm atm}\ \sin^2\left(\frac{\Delta m^2_{\rm atm} L}{4E}\right) , \tag{2.13}$$

where $\Delta m^2_{\rm atm}$ is the squared mass difference and L is in inverse energy units[1] (we continue this

[1] $L\ (\mathrm{GeV}^{-1}) = L\ (\mathrm{m})/(c\hbar) = L\ (\mathrm{m}) \cdot 5.07 \times 10^{15}\ \mathrm{m}^{-1}\ \mathrm{GeV}^{-1}$.

convention unless noted otherwise).

In practice, for atmospheric neutrinos the zenith angle of the neutrino relative to a detector serves as a proxy for the baseline L. Specifically, the baseline for a given zenith angle θ is given by

$$L = \sqrt{R_{\rm Earth}^2 \cos^2\theta + h_{\rm atm}(2R_{\rm Earth} + h_{\rm atm})} - R_{\rm Earth}\cos\theta \tag{2.14}$$

if the neutrino is produced at a height $h_{\rm atm}$ in the atmosphere, and where $\theta = 0$ corresponds to a vertically down-going neutrino. We assume that the Earth is spherical and set the radius $R_{\rm Earth} = 6370$ km, noting that the difference between the polar radius and equatorial radius is only about 0.3%. We use an average neutrino production height in the atmospheric $\langle h_{\rm atm}\rangle = 20$ km [19]. We note that any correction for detector depth is smaller than the error from either of these approximations.

A description of oscillation between all three flavors can be obtained as above, except that the transformation matrix U has a 3×3 minimum representation and has four free parameters: three mixing angles θ_{12}, θ_{13}, and θ_{23}, and a phase δ_{13} [20]. Fortunately, because of the smallness of the θ_{13} mixing angle and the "solar" mass splitting Δm_{12}, it suffices to consider a two-neutrino system in the atmospheric case.

Observations of atmospheric neutrinos by Super-Kamiokande [21], Soudan 2 [22], MACRO [23], and other experiments have provided strong evidence for mass-induced atmospheric neutrino oscillations. Observations of solar neutrinos by the Sudbury Neutrino Observatory (SNO) have also shown that the neutrinos truly change flavor, rather than decay or disappear in some other way [24]. A global fit to oscillation data from Super-Kamiokande and K2K [25] results in best-fit atmospheric parameters of $\Delta m^2_{\rm atm} = 2.2 \times 10^{-3}$ eV^2 and $\sin^2 2\theta_{\rm atm} = 1$ [26]. Thus from eq. 2.13, for energies above about 50 GeV, atmospheric neutrino oscillations cease for Earth-diameter baselines. However, a number of phenomenological models of physics beyond the Standard Model predict flavor-changing effects at higher energies that can alter the zenith angle and energy spectrum of atmospheric muon neutrinos. We consider two of these in the next chapter, violation of Lorentz Invariance and quantum decoherence.

Chapter 3

Quantum Gravity Phenomenology

Experimental searches for possible low-energy signatures of quantum gravity (QG) can provide a valuable connection to a Planck-scale theory. Hints from loop quantum gravity [27], noncommutative geometry [28], and string theory [29] that Lorentz invariance may be violated or spontaneously broken have encouraged phenomenological developments and experimental searches for such effects [30, 31]. Space-time may also exhibit a "foamy" nature at the smallest length scales, inducing decoherence of pure quantum states to mixed states during propagation through this chaotic background [32].

As we will discuss, the neutrino sector is a promising place to search for such phenomena. Water-based or ice-based Čerenkov neutrino detectors such as BAIKAL [33], AMANDA-II [34], ANTARES [35], and IceCube [36] have the potential to accumulate large samples of high-energy atmospheric muon neutrinos. Analysis of these data could reveal unexpected signatures that arise from QG phenomena such as violation of Lorentz invariance or quantum decoherence.

3.1 Violation of Lorentz Invariance

Many models of quantum gravity suggest that Lorentz symmetry may not be exact [31]. Even if a QG theory is Lorentz symmetric, the symmetry may still be spontaneously broken in our Universe. Atmospheric neutrinos, with energies above 100 GeV and mass less than 1 eV, have Lorentz gamma factors exceeding 10^{11} and provide a sensitive test of Lorentz symmetry.

Neutrino oscillations in particular are a sensitive testbed for such effects. Oscillations act as a "quantum interferometer" by magnifying small differences in energy into large flavor changes as the neutrinos propagate. In conventional oscillations, this energy shift results from the small differences

in mass between the eigenstates, but specific types of violation of Lorentz invariance (VLI) can also result in energy shifts that can generate neutrino oscillations with different energy dependencies.

The Standard Model Extension (SME) provides an effective field-theoretic approach to VLI [37]. The "minimal" SME adds all coordinate-independent renormalizable Lorentz- and CPT-violating terms to the Standard Model Lagrangian. Even when restricted to first-order effects in the neutrino sector, the SME results in numerous potentially observable effects [38, 39, 40]. To specify one particular model which leads to alternative oscillations at high energy, we consider only the Lorentz-violating Lagrangian term

$$\frac{1}{2}i(c_L)_{\mu\nu ab}\overline{L}_a\gamma^\mu\overleftrightarrow{D}^\nu L_b \tag{3.1}$$

with the VLI parametrized by the dimensionless coefficient c_L [39]. L_a and L_b are left-handed neutrino doublets with indices running over the generations e, μ, and τ, and D^ν is the covariant derivative with $A\overleftrightarrow{D}^\nu B \equiv AD^\nu B - (D^\nu A)B$.

We restrict ourselves to rotationally invariant scenarios with only nonzero time components in c_L, and we consider only a two-flavor system. The eigenstates of the resulting 2×2 matrix c_L^{TT} correspond to differing maximal attainable velocity (MAV) eigenstates. That is, eigenstates may have limiting velocities other than the speed of light and may be distinct from either the flavor or mass eigenstates. Any difference Δc in the eigenvalues will result in neutrino oscillations. The above construction is equivalent to a modified dispersion relationship of the form

$$E^2 = p^2c_a^2 + m^2c_a^4 \tag{3.2}$$

where c_a is the MAV for a particular eigenstate, and in general $c_a \neq c$ [41, 42]. Given that the mass is negligible, the energy difference between two MAV eigenstates is equal to the VLI parameter $\Delta c/c = (c_{a1} - c_{a2})/c$, where c is the canonical speed of light.

The effective Hamiltonian $H_\pm$ representing the energy shifts from both mass-induced and VLI oscillations can be written [43]

$$H_{\pm} = \frac{\Delta m^2}{4E}\mathbf{U}_{\theta}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}\mathbf{U}_{\theta}^{\dagger} + \frac{\Delta c}{c}\frac{E}{2}\mathbf{U}_{\xi}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}\mathbf{U}_{\xi}^{\dagger} \tag{3.3}$$

with two mixing angles θ (the standard atmospheric mixing angle) and ξ (a new VLI mixing angle). The associated 2×2 mixing matrices are given by

$$U_{\theta} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{3.4}$$

and

$$U_{\xi} = \begin{pmatrix} \cos\xi & \sin\xi e^{\pm i\eta} \\ -\sin\xi e^{\mp i\eta} & \cos\xi \end{pmatrix} \tag{3.5}$$

with η representing their relative phase. Solving the Louiville equation for time evolution of the state density matrix ρ

$$\dot{\rho} = -i[H_{\pm}, \rho] \tag{3.6}$$

results in the ν_{μ} survival probability. This probability $P_{\nu_{\mu}\to\nu_{\mu}}$ is given by

$$P_{\nu_{\mu}\to\nu_{\mu}} = 1 \; - \; \sin^2 2\Theta \; \sin^2\left(\frac{\Delta m^2 L}{4E}\,\mathcal{R}\right) \; , \tag{3.7}$$

where the combined effective mixing angle Θ can be written

$$\sin^2 2\Theta = \frac{1}{\mathcal{R}^2}(\sin^2 2\theta + R^2 \sin^2 2\xi + 2R \sin 2\theta \sin 2\xi \cos\eta) \; , \tag{3.8}$$

the correction to the oscillation wavelength $\mathcal{R}$ is given by

$$\mathcal{R} = \sqrt{1 + R^2 + 2R(\cos 2\theta \cos 2\xi + \sin 2\theta \sin 2\xi \cos\eta)} \; , \tag{3.9}$$

and the ratio R between the VLI oscillation wavelength and mass-induced wavelength is

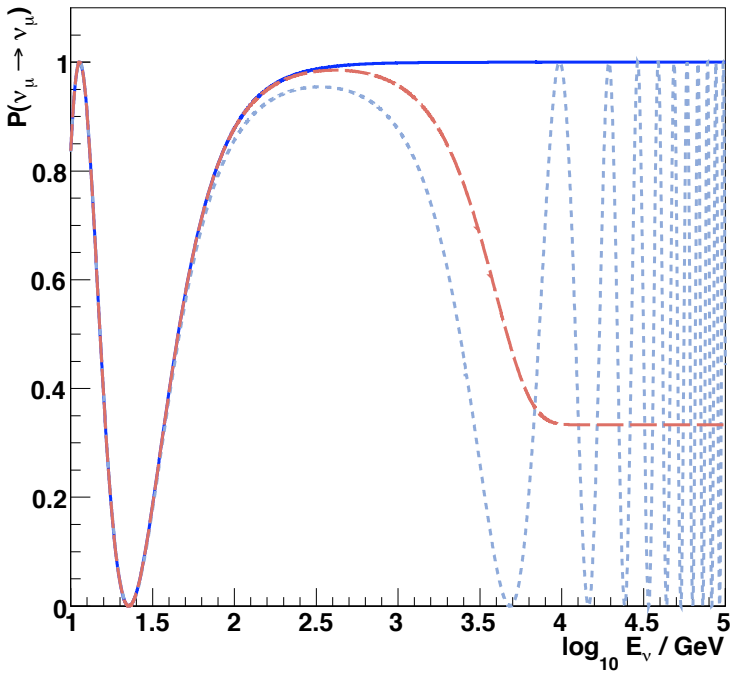

Figure 3.1: ν_μ survival probability as a function of neutrino energy for maximal baselines ($L \approx 2R_{\rm Earth}$) given conventional oscillations (solid line), VLI (dotted line, with $n = 1$, $\sin 2\xi = 1$, and $\Delta\delta = 10^{-26}$), and QD effects (dashed line, with $n = 2$ and $D^* = 10^{-30}$ GeV^{-1}).

$$R = \frac{\Delta c}{c}\frac{E}{2}\frac{4E}{\Delta m^2} \tag{3.10}$$

for a muon neutrino of energy E and traveling over baseline L. For simplicity, the phase η is often set to 0 or $\pi/2$. For illustration, if we take both conventional and VLI mixing to be maximal ($\xi = \theta = \pi/4$), this reduces to

$$P_{\nu_\mu \to \nu_\mu}(\text{maximal}) = 1 \; - \; \sin^2\left(\frac{\Delta m^2 L}{4E} + \frac{\Delta c}{c}\frac{LE}{2}\right) \; . \tag{3.11}$$

Note the different energy dependence of the two effects. The survival probability for maximal baselines as a function of neutrino energy is shown in fig. 3.1.

Several neutrino experiments have set upper limits on this manifestation of VLI, including MACRO [44], Super-Kamiokande [45], and a combined analysis of K2K and Super-Kamiokande data

[43] ($\Delta c/c < 2.0 \times 10^{-27}$ at the 90% CL for maximal mixing). In previous work, AMANDA-II has set a preliminary upper limit using four years of data of 5.3×10^{-27} [46]. Other neutrino telescopes, such as ANTARES, are also expected to be sensitive to such effects (see e.g. [47]).

Given the specificity of this particular model of VLI, we wish to generalize the oscillation probability in eq. 3.7. We follow the approach in [47], which is to modify the VLI oscillation length $L \propto E^{-1}$ to other integral powers of the neutrino energy E. That is,

$$\frac{\Delta c}{c}\frac{LE}{2} \rightarrow \Delta\delta\frac{LE^n}{2} \ , \tag{3.12}$$

where $n \in [1, 3]$, and the generalized VLI term $\Delta\delta$ is in units of GeV^{-n+1}. An $L \propto E^{-2}$ energy dependence ($n = 2$) has been proposed in the context of loop quantum gravity [48] and in the case of non-renormalizable VLI effects caused by the space-time foam [49]. Both the $L \propto E^{-1}$ ($n = 1$) and the $L \propto E^{-3}$ ($n = 3$) cases have been examined in the context of violations of the equivalence principle (VEP) [50, 51, 52]. We note that in general, Lorentz violation implies violation of the equivalence principle, so searches for either effect are related [31].

We also note that there is no reason other than simplicity to formulate the VLI oscillations as two-flavor, as any full description must incorporate all three flavors, and we know nothing of the size of the various eigenstate splittings and mixing angles. However, a two-flavor system is probably not a bad approximation, because in the most general case, one splitting will likely appear first as we increase the energy. Also, since we will search only for a deficit of muon neutrinos, we do not care to which flavor the ν_μ are oscillating (so it need not be ν_τ).

3.2 Quantum Decoherence

Another possible low-energy signature of QG is the evolution of pure states to mixed states via interaction with the environment of space-time itself, or quantum decoherence. One heuristic picture of this phenomenon is the production of virtual black hole pairs in a "foamy" spacetime, created from the vacuum at scales near the Planck length [53]. Interactions with the virtual black holes may not preserve certain quantum numbers like neutrino flavor, causing decoherence into a superposition of flavors.

Quantum decoherence can be treated phenomenologically as a quantum open system which evolves thermodynamically. The time-evolution of the density matrix ρ is modified with a dissipative term $\not{\partial} H\rho$:

$$\dot{\rho} = -i[H, \rho] + \not{\partial} H\rho \ . \tag{3.13}$$

The dissipative term representing the losses in the open system is modeled via the technique of Lindblad quantum dynamical semigroups [54]. Here we outline the approach in ref. [55], to which we refer the reader for more detail. In this case, we have a set of self-adjoint environmental operators $\mathcal{A}_j$, and eq. 3.13 becomes

$$\dot{\rho} = -i[H, \rho] + \frac{1}{2}\sum_j ([A_j, \rho A_j] + [A_j\rho, A_j]) \ . \tag{3.14}$$

The hermiticity of the A_j ensures the monotonic increase of entropy, and in general, pure states will now evolve to mixed states. The irreversibility of this process implies CPT violation [56].

To obtain specific predictions for the neutrino sector, there are again several approaches for both two-flavor systems [57, 58] and three-flavor systems [55, 59]. Again, we follow the approach in [55] for a three-flavor neutrino system including both decoherence and mass-induced oscillations. The dissipative term in eq. 3.14 is expanded in the Gell-Mann basis $F_\mu, \mu \in [0, 8]$, such that

$$\frac{1}{2}\sum_j ([A_j, \rho A_j] + [A_j\rho, A_j]) = \sum_{\mu,\nu} L_{\mu\nu}\rho_\mu F_\nu \ . \tag{3.15}$$

At this stage we must choose a form for the decoherence matrix $L_{\mu\nu}$, and we select the weak-coupling limit in which L is diagonal, with $L_{00} = 0$ and $L_{ii} = -D_i$. These D_i represent the characteristic length scale over which decoherence effects occur. Solving this system for atmospheric neutrinos (where we neglect mass-induced oscillations other than $\nu_\mu \to \nu_\tau$) results in the ν_μ survival probability [55]:

$$P_{\nu_\mu \to \nu_\mu} = \frac{1}{3} \quad + \quad \frac{1}{2}\Bigg[\frac{1}{4}e^{-LD_3}(1+\cos 2\theta)^2 + \frac{1}{12}e^{-LD_8}(1-3\cos 2\theta)^2 + e^{-\frac{L}{2}(D_6+D_7)}$$
$$\cdot \quad \sin^2 2\theta\Bigg(\cos\left[\frac{L}{2}\sqrt{\left(\frac{\Delta m^2}{E}\right)^2 - (D_6-D_7)^2}\right]$$
$$+ \quad \frac{\sin\left[\frac{L}{2}\sqrt{\left(\frac{\Delta m^2}{E}\right)^2 - (D_6-D_7)^2}\right](D_6-D_7)}{\sqrt{\left(\frac{\Delta m^2}{E}\right)^2 - (D_6-D_7)^2}}\Bigg)\Bigg] \,. \tag{3.16}$$

Note the limiting probability of $\frac{1}{3}$, representing full decoherence into an equal superposition of flavors. The D_i not appearing in eq. 3.16 affect decoherence between other flavors, but not the ν_μ survival probability.

We note that in eq. 3.16, we must impose the condition $\Delta m^2/E > |D_6 - D_7|$, but this is not an issue in the parameter space we explore in this analysis. If one wishes to ensure strong conditions such as complete positivity [57], there may be other inequalities that must be imposed (see e.g. the discussion in ref. [59]).

The energy dependence of the decoherence terms D_i depends on the underlying microscopic model. As with the VLI effects, we choose a generalized phenomenological approach where we suppose the D_i vary as some integral power of the energy,

$$D_i = D_i^* E^n, \; n \in [1,3] \,, \tag{3.17}$$

where E is the neutrino energy in GeV, and the units of the D_i^* are GeV^{-n+1}. The particularly interesting E^2 form is suggested by decoherence calculations in non-critical string theories involving recoiling D-brane geometries [60]. We show the $n = 2$ survival probability as a function of neutrino energy for maximal baselines in fig. 3.1.

An analysis of Super-Kamiokande in a two-flavor framework has resulted in an upper limit at the 90% CL of $D^* < 9.0 \times 10^{-28}$ GeV^{-1} for an E^2 model and all D_i^* equal [61]. ANTARES has reported sensitivity to various two-flavor decoherence scenarios as well, using a more general formulation [58]. Analyses of Super-Kamiokande, KamLAND, and K2K data [62, 63] have also set

strong limits on decoherence effects proportional to E^0 and E^{-1}. Because our higher energy range does not benefit us for such effects, we do not expect to be able to improve upon these limits, and we focus on effects with $n \geq 1$.

Unlike the VLI system, we have used a full three-flavor approach to the phenomenology of the QD system. There is no theoretical justification for doing so in one but not the other, but for the special case in which all decoherence parameters are equal, the choice is important. This is because in a three-flavor system, the limiting survival probability is 1/3, compared to 1/2 in a two-flavor system. Since heuristically the equality of decoherence parameters suggests that the interactions with space-time are flavor-agnostic, we feel that using a three-flavor description is more apt.

Chapter 4

Neutrino Detection

4.1 General Techniques

A major obstacle to overcome in the detection of the neutrino is its small cross section: while the neutrino-nucleon cross section rises with energy, at 1 TeV the interaction length is still 2.5 *million* kilometers of water [64]. Thus, any potential detector must encompass an enormous volume to achieve a reasonable event rate. Once an interaction does occur in or near the detector, we can detect the resulting charged particles by means of their radiation. A (relatively) cost-effective approach is to use natural bodies of water or transparent ice sheets as the target material, and then instrument this volume with photomultiplier tubes. While originally proposed in 1960 by K. Greisen and F. Reines [65, 66], large-scale detectors of this sort have only been in operation for the past decade or so.

Water or ice neutrino detectors typically consist of vertical cables (called "strings" or "lines") lowered either into deep water or into holes drilled in the ice. Photomultiplier tubes (PMTs) in pressure housings are attached to the cables, which supply power and communications. A charged-current neutrino interaction with the surrounding matter produces a charged lepton via the process

$$\nu_l(\overline{\nu}_l) + q \rightarrow l^-(l^+) + q' \ , \tag{4.1}$$

where q is a valence or sea quark in the medium, and q' is as appropriate for charge conservation. In the case of a muon neutrino, the resulting muon can travel a considerable distance within the medium.

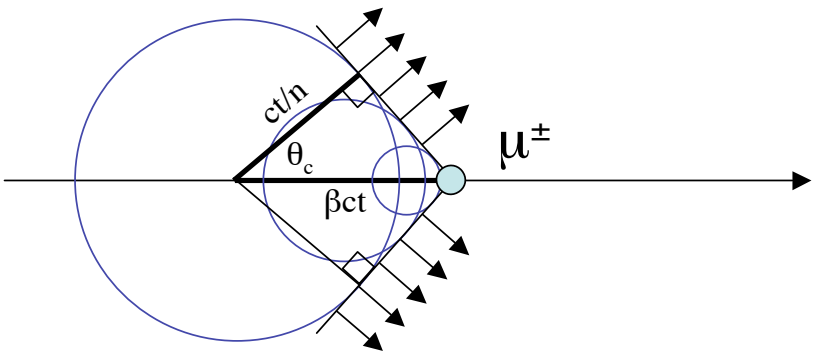

Figure 4.1: Formation of a Čerenkov cone by a relativistic muon moving through a medium.

4.2 Čerenkov Radiation

Because the relativistic muon produced in the neutrino interaction is traveling faster than the speed of light in the medium, it will radiate via the Čerenkov effect. A coherent "shock wave" of light forms at a characteristic angle θ_c depending on the index of refraction n of the medium, specifically,

$$\cos\theta_c = \frac{1}{n\beta} \ , \tag{4.2}$$

where $\beta = v/c$ is the velocity of the particle. For ice, where $n \approx 1.33$, the Čerenkov angle is about 41° for relativistic particles ($\beta \approx 1$). A full treatment differentiates between the phase and group indices of refraction, but this is a small correction (see e.g. [67]). Figure 4.1 presents a geometric derivation of the simpler form shown in eq. 4.2.

The number of Čerenkov photons emitted per unit track length as a function of wavelength λ is given by the Franck-Tamm formula [68]

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha}{\lambda^2}\left(1 - \frac{1}{\beta^2 n^2}\right) \ , \tag{4.3}$$

where α is the fine-structure constant. Because of the $1/\lambda^2$ dependence, the high-frequency photons dominate the emission, up to the ultraviolet cutoff imposed by the glass of the PMT pressure vessel

(about 365 nm [69]). Between this and the frequency at which is the ice is no longer transparent (about 500 nm; see section 5.2), we expect an emission of about 200 photons per centimeter [70].

4.3 Muon Energy Loss

Čerenkov radiation from the bare muon is not its dominant mode of energy loss. The rate of energy loss as a function of distance, dE/dx, can be parametrized as

$$-\frac{dE}{dx} = a(E) + b(E)\ E\ , \tag{4.4}$$

where $a(E)$ is the ionization energy loss given by the standard Bethe-Bloch formula (see e.g. [71]), and $b(E)$ is the sum of losses by e^+e^- pair production, bremsstrahlung, and photonuclear interactions. The energy losses from various contributions are shown in figure 4.2.

The ionization energy losses are continuous in nature, occurring smoothly along the muon track. However, at high energies, the losses by bremsstrahlung, pair production, and photonuclear interactions are not continuous but stochastic: rare events that result in large depositions of energy via particle and photon creation. The particles produced are highly relativistic, and if charged, they too will radiate via the Čerenkov effect. Furthermore, because they are kinematically constrained to the approximate direction of the muon, this emission will peak at the Čerenkov angle of the muon. The roughly conical Čerenkov emission of the bare muon is thus enhanced by the various energy losses described above [73].

4.4 Other Event Topologies

For charged-current ν_e and ν_τ interactions, or neutral-current interactions of any flavor, the event topology is less track-like than the muon case described above, and is instead more spherical or "cascade-like." For ν_e events, this is because of the short path length of the resulting electron or positron within the ice. For ν_τ events (except for those of very high energy), the resulting τ lepton will decay immediately, in most cases resulting in a hadronic shower. However, 17% of the time [20], the τ will decay via

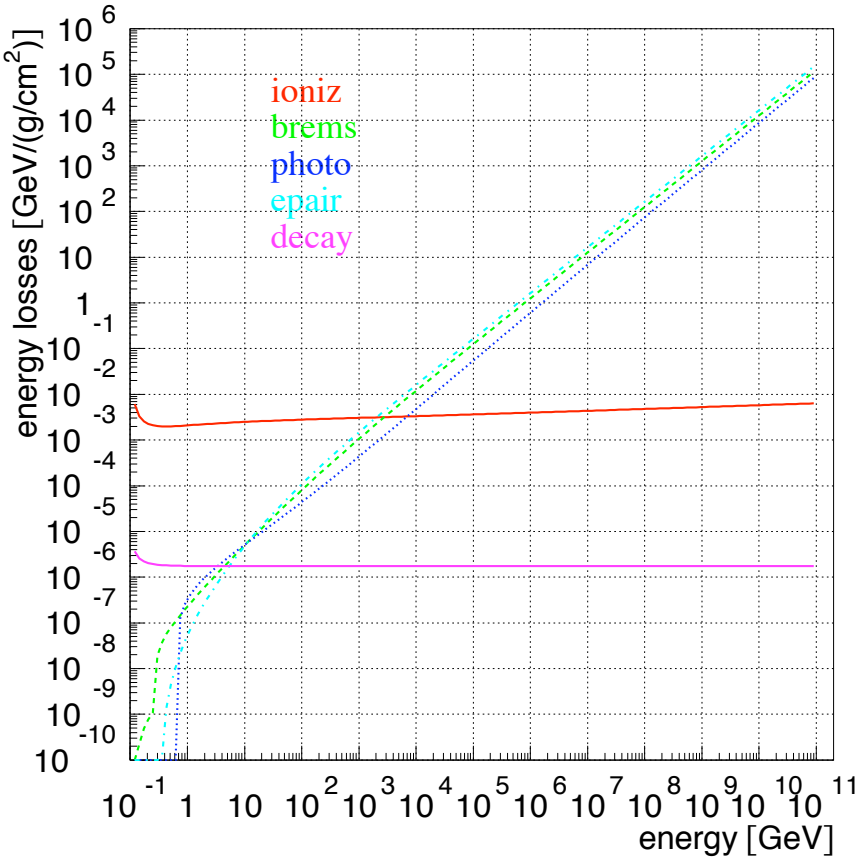

Figure 4.2: Average muon energy loss in ice as a function of energy, from [72].

$$\begin{aligned} \tau^+ &\rightarrow \mu^+ + \nu_\mu + \bar{\nu}_\tau \\ \tau^- &\rightarrow \mu^- + \bar{\nu}_\mu + \nu_\tau \ , \end{aligned} \tag{4.5}$$

possibly resulting in a detectable muon track (albeit of significantly lower energy than the original ν_τ). For a neutral-current event, there is no outgoing charged lepton, although there may be a hadronic shower from the collision.

Because cascade-like and track-like events have very different signatures and strategies for background rejection, one generally focuses on one or the other early in the analysis. We consider only ν_μ-induced muons in this analysis; other types of event will be removed by the data-filtering procedures which we describe in chapter 6.

4.5 Background

While we have described the means by which we might detect a neutrino-induced muon, the background to such a search is formidable. Even with kilometers of overburden, high-energy atmospheric muon bundles dominate over neutrino events by a factor of about 10^6. Selecting only "up-going" muons allows us to reject the large background of atmospheric muons, using the Earth as a filter to screen out everything but neutrinos (see fig. A.2). In practice, we must also use other observables indicating the quality of the muon directional reconstruction, in order to eliminate mis-reconstructed atmospheric muon events — a topic we will revisit in chapter 6.

Chapter 5

The AMANDA-II Detector

5.1 Overview

AMANDA, or the **A**ntarctic **M**uon **A**nd **N**eutrino **D**etector **A**rray, consists of 677 optical modules (OMs) on 19 vertical cables or "strings" frozen into the deep, clear ice near the geographic South Pole. Each OM consists of an 8" diameter Hamamatsu R5912-2 photomultiplier tube (PMT) housed in a glass pressure sphere. The AMANDA-II phase of the detector commenced in the year 2000, after nine outer strings were added. Fig. 5.1 shows the geometry of the detector, as well as the principal components of the OMs.

The bulk of the detector lies between 1550 and 2050 meters under the snow surface, where the Antarctic ice sheet is extremely clear. The 19 strings are arranged roughly in three concentric cylinders, the largest of which is approximately 200 meters in diameter. The OMs are connected to cables which supply power and transmit PMT signals to the surface. Multiple cabling technologies are used: coaxial, twisted-pair, and fiber optic. While most transmitted signals are analog, string 18 contains prototype digital optical modules (DOMs) which digitize the PMT signal before transmission.

5.2 Optical Properties of the Ice

Far from being a homogeneous medium, the ice at the South Pole consists of roughly horizontal layers of varying clarity. As the ice layers accumulated over geological time periods, varying amounts of atmospheric dust were trapped during the deposition, depending on the climatological conditions at the time. These "dust layers" strongly affect both the optical scattering and absorption lengths

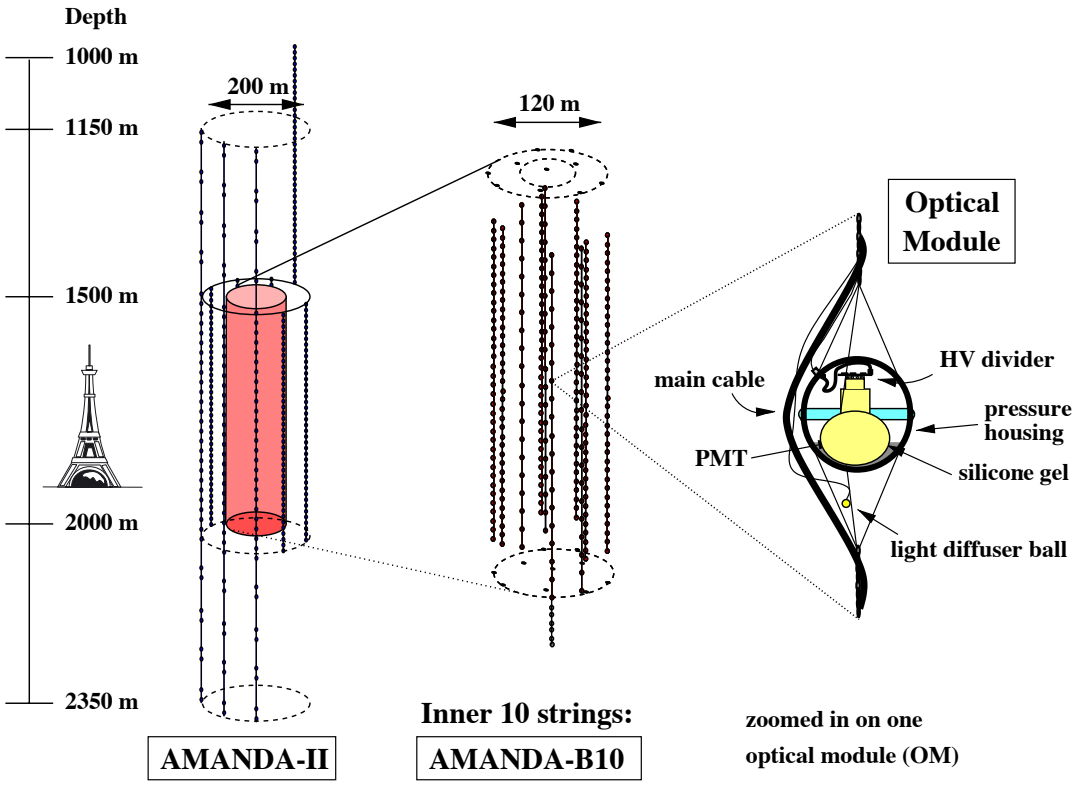

Figure 5.1: Diagram of the AMANDA-II detector with details of an optical module (from [74]). The Eiffel tower on the left illustrates the scale.

and must be taken into account for event reconstruction and simulation.

The scattering and absorption properties have been measured or inferred using a number of *in situ* light sources [75], resulting in a comprehensive model of the ice properties known as the "Millennium" ice model. Since the publication of that work, the effect of smearing between dust layers has been examined, resulting in an updated model of the ice known as "AHA." The effective scattering length in this model λ_s^{eff}, defined such that

$$\lambda_s^{\text{eff}} = \frac{\lambda_s}{1 - \langle \cos\theta_s \rangle} \ , \tag{5.1}$$

with an average scattering angle of $\langle \cos\theta_s \rangle \approx 0.95$, is shown along with the absorption length λ_a in fig. 5.2. The effective scattering length is approximately 20 m, whereas the absorption length (at 400 nm) is about 110 m.

5.3 Data Acquisition and Triggering

Cables from the deep ice are routed to surface electronics housed in the Martin A. Pomerantz Observatory (MAPO). PMT signals, broadened after transmission to the surface, are amplified in **Sw**edish **amp**lifiers (SWAMPs). A prompt output from the SWAMPs is fed to a discriminator, which in turn feeds the trigger logic and a time-to-digital converter (TDC). The TDC records the leading and falling edges when the signal crosses the discriminator level. Each edge pair forms a *hit*, of which the TDC can store eight at a time. The difference between the edges is referred to as the time-over-threshold, or TOT.

The main trigger requires 24 hit OMs within a sliding window of 2.5 μs. The hardware core of the trigger logic is formed by the **d**igital **m**ultiplicity **ad**der-**d**iscriminator (DMADD). When the trigger is satisfied, the trigger electronics open the gate to a peak-sensing analog-to-digital converter (ADC) which is fed by a delayed signal from the SWAMPs. The ADC gate remains open for 9.8 μs, and the peak amplitude during that window is assigned to *all* hits in that particular channel. 10 μs after the trigger, a stop signal is fed to the TDC. The trigger is also sent to a GPS clock which timestamps the event. Events are recorded to magnetic tape and then flown north during the austral summer season.

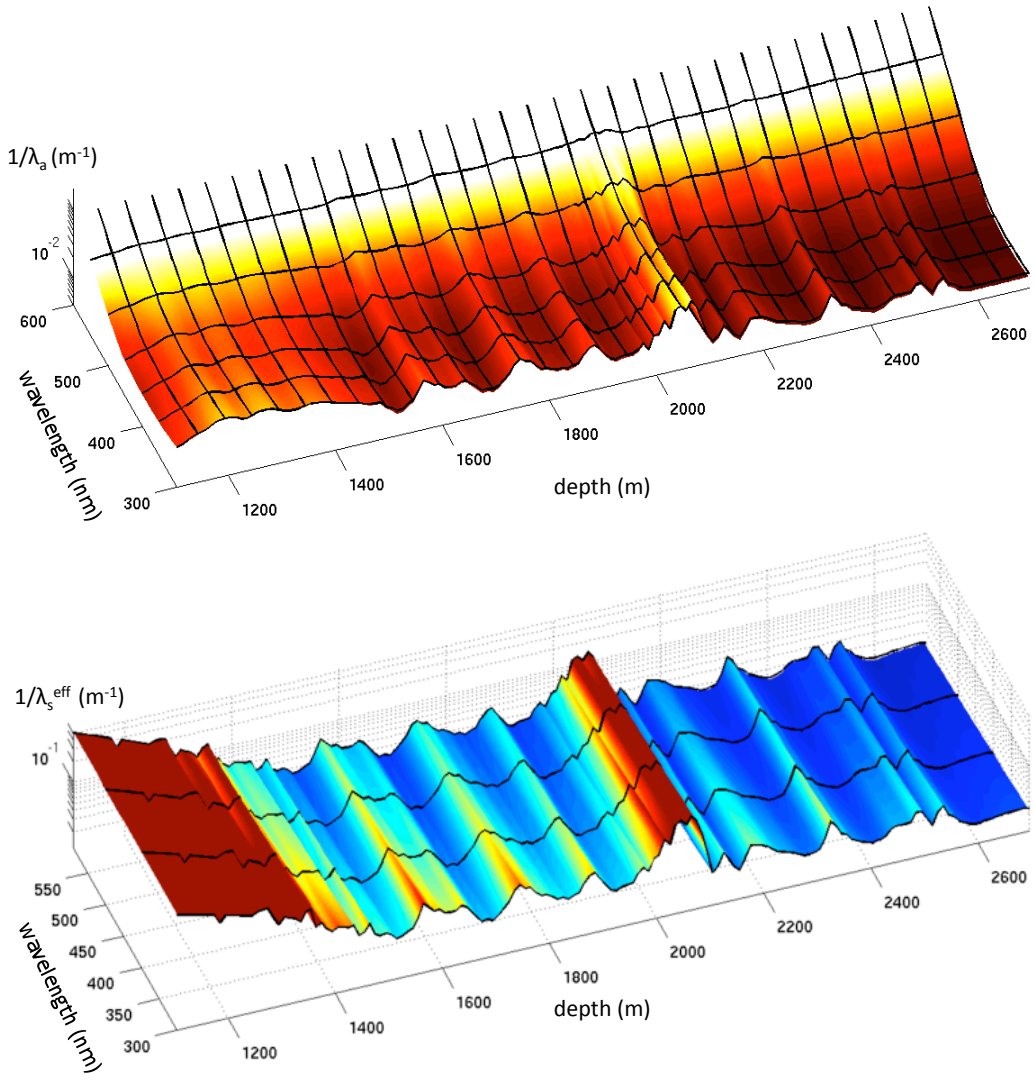

Figure 5.2: Inverse absorption and effective scattering lengths as a function of depth and wavelength in the AHA ice model. Note the large dust peak at a depth of roughly 2050 m, with three smaller dust peaks at shallower depths.

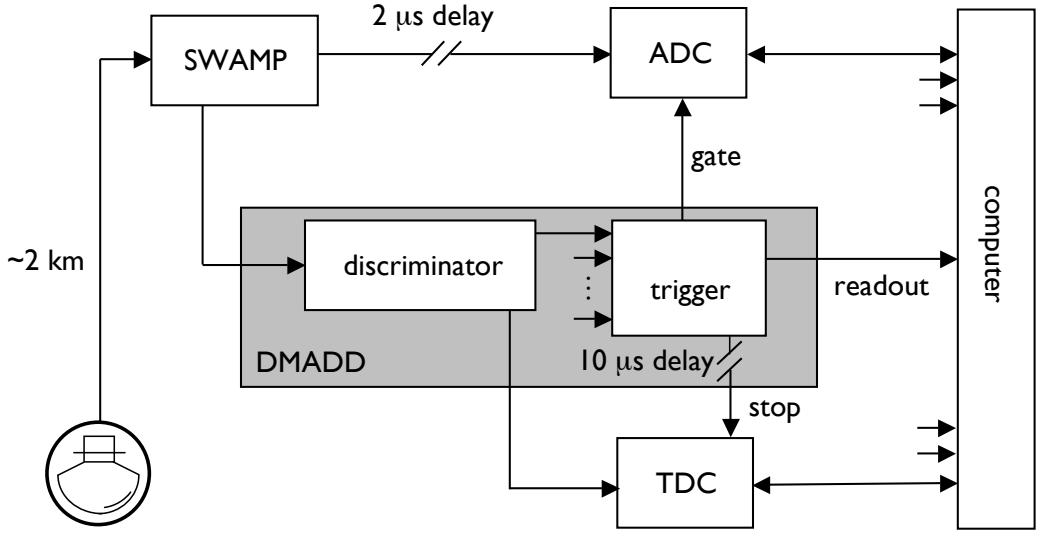

Figure 5.3: Principal components of the AMANDA μ-DAQ (adapted from [67]).

The data acquisition system (DAQ) described here is known as the muon DAQ, or μ-DAQ, and operated through 2006. A parallel DAQ that also records waveform information, the TWR-DAQ, has been fully operational since 2004. Only μ-DAQ data are used in this analysis. A simplified block diagram showing the principal components in shown in fig. 5.3.

5.4 Calibration

Calibration of cable time delays and the corrections for dispersion are performed with *in situ* laser sources. After calibration, the time resolution for the first 10 strings is $\sigma_t \approx 5$ ns (those with coaxial or twisted-pair cables), and the time resolution for the optical fiber strings is $\sigma_t \approx 3.5$ ns [74]. The time delay calibration is cross-checked using down-going muon data.

The amplitude calibration uses single photoelectrons (SPEs) from low-energy down-going muons as a "calibration source" of known charge. The uncalibrated amplitude distribution of SPEs is fit as the sum of an exponential and a Gaussian distribution, with the peak of the Gaussian portion representing one PE.

Chapter 6

Simulation and Data Selection

6.1 Simulation

In order to meaningfully compare our data with expectations from various signal hypotheses, we must have a detailed Monte Carlo (MC) simulation of the atmospheric neutrinos and the subsequent detector response. For the input atmospheric muon neutrino spectrum, we generate an isotropic power-law flux with the NUSIM neutrino simulator [76] and then reweight the events to standard flux predictions [16, 17]. As discussed in section 2.2.2, we have extended the predicted fluxes to the TeV energy range via the NEUTRINOFLUX package, which fits the low-energy region with the Gaisser parametrization [11] and then extrapolates above 700 GeV. We add standard oscillations and/or non-standard flavor changes by weighting the events with the muon neutrino survival probabilities in eqs. 2.13, 3.7, or 3.16.

Muon propagation and energy loss near and within the detector are simulated using MMC [72]. Photon propagation through the ice, including scattering and absorption, is modeled with PHOTONICS [77], incorporating the depth-dependent characteristic dust layers from the AHA ice model (see section 5.2). The detector simulation AMASIM [78] records the photon hits, and then identical filtering and reconstruction methods are performed on data and simulation. Cosmic ray background rejection is ensured at all but the highest quality levels by a parallel simulation chain fed with atmospheric muons from CORSIKA [79], although when reaching contamination levels of $O(1\%)$ — a rejection factor of 10^8 — computational limitations become prohibitive.

As we will discuss further in chapter 8, the absolute sensitivity of the OMs is one of the

larger systematic uncertainties. Determining the effect on our observables can only be achieved by rescaling the sensitivity within AMASIM and re-running the detector simulation. We have generated atmospheric neutrino simulation for 7 different optical module sensitivities, and for each set we reach an effective livetime (see appendix D) of approximately 60 years.

6.2 Filtering and Track Reconstruction

Filtering the large amount of raw AMANDA-II data from trigger level to neutrino level is an iterative procedure. First, known bad optical modules are removed, resulting in approximately 540 OMs for use in the analysis. Unstable or incomplete runs ("bad files") are identified and excluded. Hits caused by electrical crosstalk and isolated noise hits are also removed ("hit cleaning").

The initial data volume is so large that only fast, first-guess algorithms can be run on all events. These include the *direct walk* algorithm [74] and JAMS [80], both of which employ pattern-matching algorithms to reconstruct muon track directions. If the zenith angle is close to up-going (typically greater than 70° or 80°), the event is kept in the sample. This step is known as "level 1" or L1 filtering. "Level 2" and "level 3" filtering steps consist of more computationally intensive directional reconstructions, along with another zenith angle cut using the more accurate results.

The best angular resolution is achieved by likelihood-based reconstructions utilizing the timing information of the photon hits. The iterative *unbiased likelihood* (UL) reconstruction uses the timing of the first hit in an OM, and maximizes the likelihood

$$\mathsf{L} = \prod_i p(t_i|\vec{a}) \ , \tag{6.1}$$

where $\vec{a} = (x, y, z, \theta, \phi)$ is the track hypothesis and t_i is the *timing residual* for hit i. The timing residual is the difference between the expected photon arrival time based only on geometry and the actual arrival time, which in general is delayed by scattering in the ice. A parametrization of the probability distribution function describing the time residuals of the first photon hits is given by the Pandel function[1] [81], and this is convoluted with a Gaussian to include PMT jitter. A high-quality sample event is shown in fig. 6.1 along with the track from the UL reconstruction.

[1]For this reason, the UL reconstruction is also commonly referred to as the "Pandel" reconstruction (even though other reconstructions also use the Pandel p.d.f.). We use the terms interchangeably.

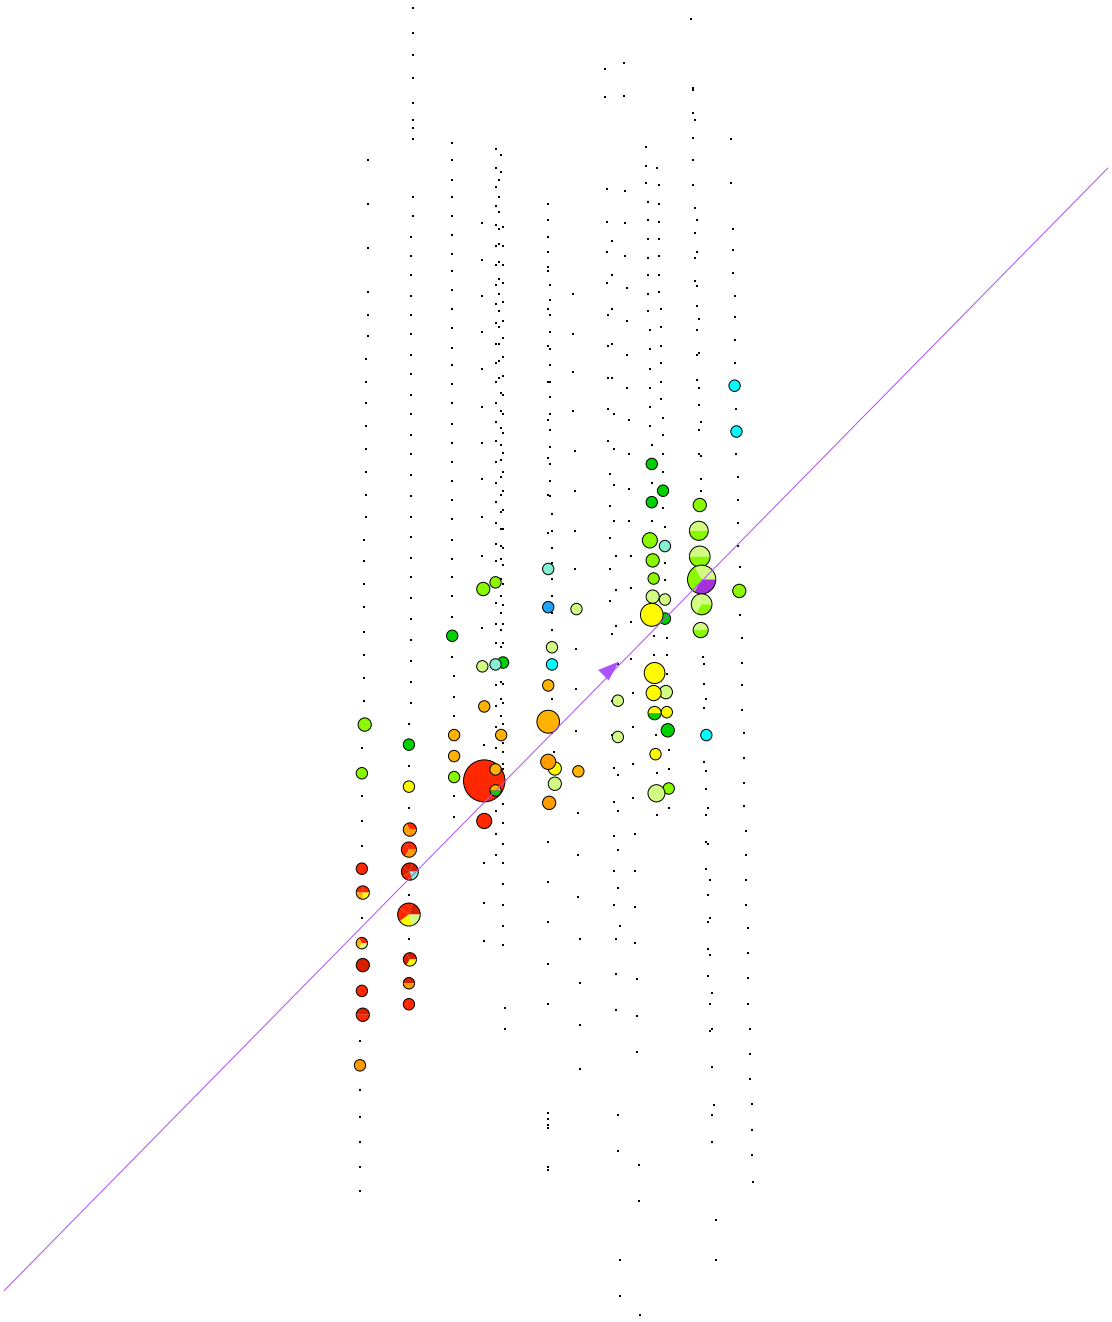

Figure 6.1: Sample AMANDA-II event from 2001, with number of OMs hit $N_{ch} = 77$. Colors indicate the timing of the hits, with red being earliest. The size of the circles indicate the amplitude of the PMT signal. The line is the reconstructed track obtained from the unbiased likelihood or "Pandel" fit ($\sigma = 1.2°$).

Because we are interested in rejecting atmospheric muons, it is useful to compare the UL likelihood with a reconstruction using the prior hypothesis that the event is down-going. This *Bayesian likelihood* (BL) reconstruction weights the likelihood with the prior probability density function (PDF) $P_\mu(\theta)$, a polynomial fit to the zenith-angle dependence of the muon flux at the depth of the AMANDA-II detector. The likelihood to maximize therefore becomes

$$\mathsf{L}_{\text{Bayes}} = \prod_i p(t_i|\vec{a})\ P_\mu(\theta)\ . \tag{6.2}$$

The log-likelihood ratio of the UL fit to the BL fit is then a test statistic which we can use as a quality parameter.

6.3 Quality Cuts

After initial filtering and reconstruction (after level 3), atmospheric neutrino events are still dominated by mis-reconstructed atmospheric muons (down-going muons that are incorrectly reconstructed as up-going). In order to remove (or "cut") these events, we must use several variables indicating the quality of the track reconstruction.

6.3.1 Point-source Cuts

As a starting point for these quality cuts, we use the criteria developed for the 2000-2004 AMANDA-II point source analysis [82]. These cuts are applicable to an atmospheric neutrino analysis primarily because they are not optimized for a high-energy extraterrestrial neutrino flux, and so their efficiency for lower-energy atmospheric neutrinos is still quite good. These cuts are shown in table 6.1.

6.3.1.1 Likelihood Ratio

As described in section 6.2, the log likelihood ratio $\log \mathsf{L}_{\text{UL}}/\mathsf{L}_{\text{BL}}$ tests the relative probability of the unbiased "Pandel" fit (which, because of the zenith-angle cuts, indicates the track is up-going or close to it), and the Bayesian down-going fit. The larger this ratio, the less likely the event is to be down-going. Including a dependence on the zenith angle is necessary because this test gets less powerful near the horizon, where the up-going and down-going tracks may only be separated by a

Table 6.1: Initial quality cuts, originally designed for the 2000-2004 point-source analysis. $x = \sin\delta_{UL}$, where δ is the reconstructed declination, and $\Theta()$ is the unit step function. Quality variables are described in the text.

Description	**Criterion (to keep)**
Likelihood ratio	$\log \mathsf{L}_{\mathrm{UL}}/\mathsf{L}_{\mathrm{BL}} > 34 + 25(1-\Theta(x-0.15))\cdot(x-0.15)$
Smoothness	$\lvert S_{\mathrm{Phit,UL}}\rvert < 0.36$ and $S_{\mathrm{Phit,UL}} \neq 0$
Paraboloid error	$\sigma < 3.2^\circ - 4^\circ \cdot \Theta(x-0.75)\cdot(x-0.75)$ and $\sigma \neq 0$
Flariness	$F_{\text{TOT-short}} + F_{B10} + F_{\overline{B10}} < 10$
Stability period	2000: day $\in [47, 309]$ 2001: day $\in [44, 293]$ 2002: day $\in [43, 323]$ 2003: day $\in [43, 315]$ 2004: day $\in [43, 309]$ 2005-06: included in file selection
File size	2000-04: Runs/file > 6
Bayesian bug fix	2005-06: $\theta_{\mathrm{BL}} < 90^\circ$

small angle.

6.3.1.2 Smoothness

The *smoothness* of a track hypothesis is a topological parameter varying from -1 to 1 that measures the evenness of the hits along the track. Positive values of the smoothness indicate hits concentrated at the beginning of the track, while negative values indicate hits concentrated at the end. Small absolute value of the smoothness indicates that the hits are more evenly distributed, supporting the fit hypothesis.

The particular implementation of the smoothness calculation that we use, the P_{hit} smoothness or S_{Phit}, only considers direct (unscattered) hits within a 50 m cylinder around the UL track. It then compares the number of hits in the cylinder to the number expected given a minimally ionizing muon (see ref. [80] for more detail). In table 6.1, we also explicitly exclude events with exactly zero smoothness; in an early implementation, this result indicated *no* direct hits within the cylinder.

6.3.1.3 Paraboloid Error

The *paraboloid error*, or angular resolution parameter, is an estimate of the 1σ error on the direction of the UL fit, and poorly reconstructed tracks will tend to have higher values of this error. This parameter is obtained by fitting a paraboloid to the likelihood space near the best-fit minimum

[83]. Using the fit to approximate the 1σ confidence level results in an error ellipse with major and minor axes σ_1 and σ_2, and we form with these a single error parameter σ by taking the geometric mean

$$\sigma = \sqrt{\sigma_1 \sigma_2} \ . \tag{6.3}$$

The quality cut on the paraboloid error is tightened near the horizon, where background contamination worsens. Pathological results, such as zero or negative values of the ellipse axes, must be excluded.

6.3.1.4 Flariness and Stability Period

In periods of windy or stormy weather at the South Pole, electrical induction on the surface cables (especially twisted-pair cables) can actually result in enough spurious "hits" to trigger the detector. These types of fake events are known as *flares*, and to enable removal of these events, a number of characteristic *flare variables* have been developed [84]. We use three of these flare variables: $F_{\text{TOT-short}}$, the number of hits in twisted-pair channels with TOT shorter than expected; F_{B10}, the ratio of hits in strings 1-4 (coaxial cabling) to strings 5-10 (twisted-pair cabling); and $F_{\overline{B10}}$, the ratio of hits in strings 11-19 with optical cabling to those with twisted-pair cabling. The flare indicators are normalized as logarithms of probabilities, so we use their sum as a combined flare indicator.

Finally, since configuration changes often take place during the austral summer season, we exclude these periods with a *stability period* cut based on the day of the year of the event. This is only necessary for data from 2000-2004, as for 2005 and 2006 this was performed as part of the file selection before the initial data filtering.

6.3.2 Purity Cuts

After applying the cuts described in the previous section, we can examine the purity of the sample using a cut-tightening ratio procedure, and then try to isolate any remaining background events. First, we uniformly tighten or loosen all the cut parameters by scaling them linearly with a "cut strength" parameter. We multiply the observable x by either α or $1/\alpha$ depending on the cut:

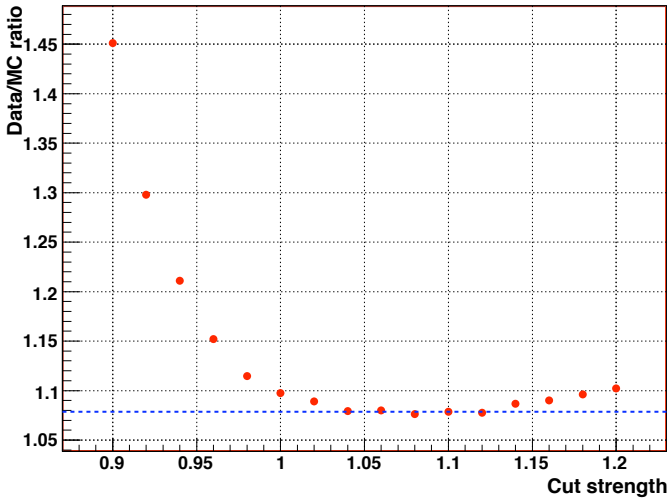

Figure 6.2: Data/MC ratio as a function of cut strength (1 = point-source cuts). The dashed line indicates the average of the flat region.

$$\begin{aligned} (x_i < X_0) &\rightarrow (\alpha x_i < X_0) \ , \\ (x_i > X_0) &\rightarrow \left(\frac{x_i}{\alpha} > X_0\right) \ \forall \ i \ . \end{aligned} \tag{6.4}$$

Figure 6.2 shows the result of this procedure. As the cuts are tightened, the background is eventually eliminated, leaving only a systematic normalization shift in the simulation. After performing this procedure using the point-source cuts, we find that there is a background contamination of 2-3% (the difference between the ratio at the nominal cuts, cut strength=1, and the ratio in the flat region). As we will see, however, this procedure is not foolproof; it only eliminates background at the tails of the quality parameters we are already using.

We have isolated several parameters which are of use to reduce this background even further:

- Ψ, the space angle between the Pandel/UL and JAMS fits;
- N_{dirC}, the number of direct hits (class C, with $t_{res} \in [-15, 75]$ ns); and
- L_{dirB}, the maximum separation along the track of direct hits (class B, with timing residual $t_{res} \in [-15, 25]$ ns).

In each of these variables, there are tails in the data that are not matched by atmospheric neutrino simulation but are consistent with mis-reconstructed muons. This is shown in fig. 6.3, along with the region we have chosen to eliminate. These purity cuts are also shown in table 6.2.

Table 6.2: Purity cuts designed to remove the remaining background contamination after applying the quality cuts shown in table 6.1.

Description	**Criterion (to keep)**
Up-going	$\cos\theta_{\mathrm{UL}} < 0$
Space angle	$\Psi_{\mathrm{UL,JAMS}} < 25^{\circ}$
Direct hits	$N_{\mathrm{dirC}} > 6$
Direct length	$L_{\mathrm{dirB}} > 60$ m

After reapplying the same cut-tightening ratio procedure described above, we estimate that these cuts achieve a purity of greater than 99%. This reduces the data from 6099 to 5686 events (-7%) while reducing the atmospheric neutrino MC prediction by 2%. However, this turns out to be an overestimation of the purity.

6.3.3 High-N_{ch} Excess and Additional Purity Cuts

The selection criteria described above, as well as the analysis procedure described in chapter 7, were designed blindly. Specifically, our observables (the zenith angle and number of OMs hit, N_{ch}; see section 7.1) were hidden until the cuts and analysis procedures were finalized. However, after unblinding, we examined the N_{ch} and $\cos\theta_{\mathrm{UL}}$ distributions and found an unexpected 1.5% excess in the ($60 < N_{ch} < 120$) region (see fig. 6.4). This is slightly higher than the estimated 0.5% background contamination, and more importantly, it is not distributed evenly across the observable space (which is how we model background contamination in the systematic errors). In this section, we discuss the impact of the excess, and how we have chosen to address it.

An analysis of the events in the high-N_{ch} region suggests that the excess (about 85 events compared to Monte Carlo normalized to the low-N_{ch} region) consists of mis-reconstructed muons. Salient observations about the excess include:

- the excess is evenly spread across all years;
- events scanned in the event viewer are not track-like (perhaps muon bundles passing outside

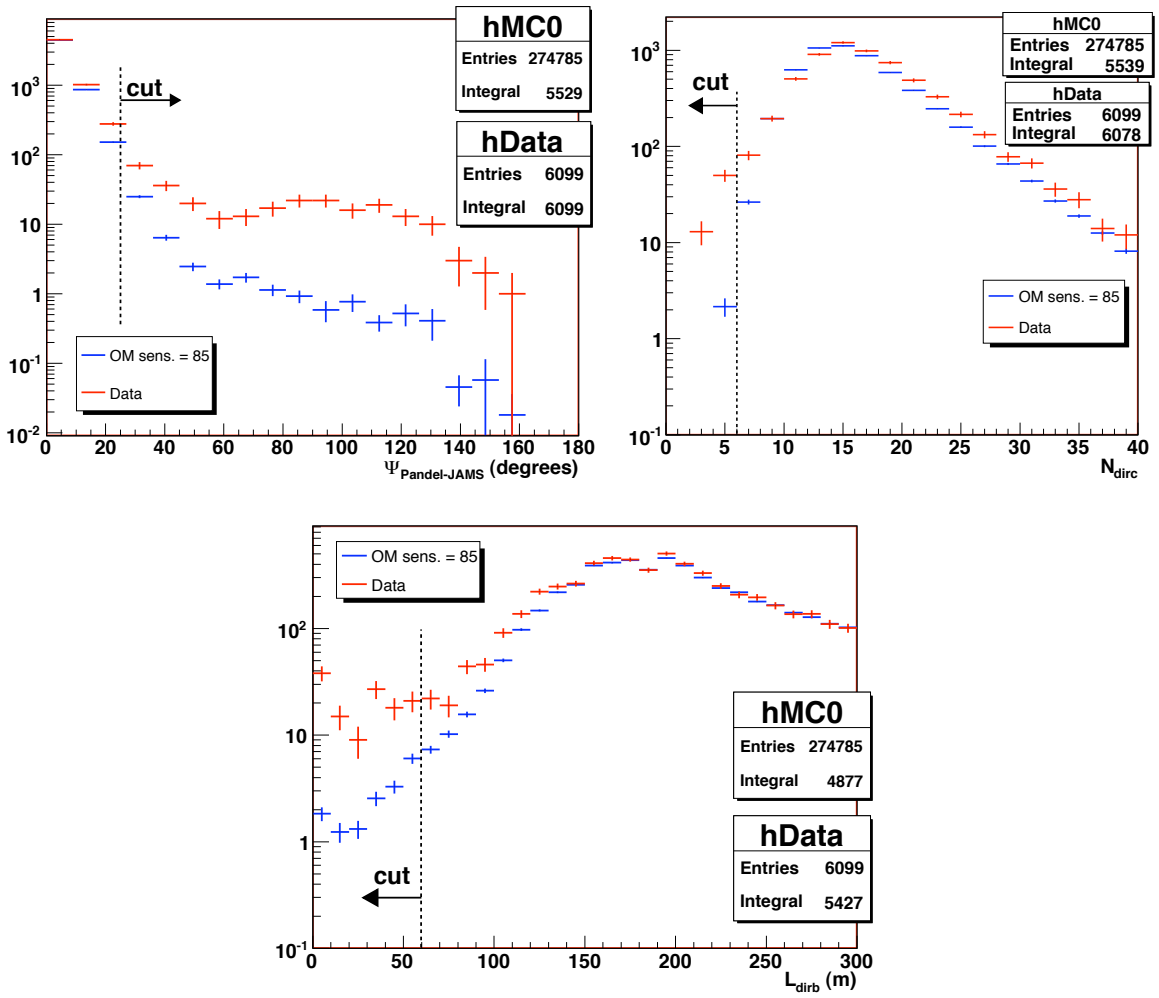

Figure 6.3: Data (red points) and atmospheric neutrino simulation (blue points) exhibiting background contamination in several quality variables after the point-source cuts. Purity cuts designed to reduce this background are shown graphically with vertical lines.

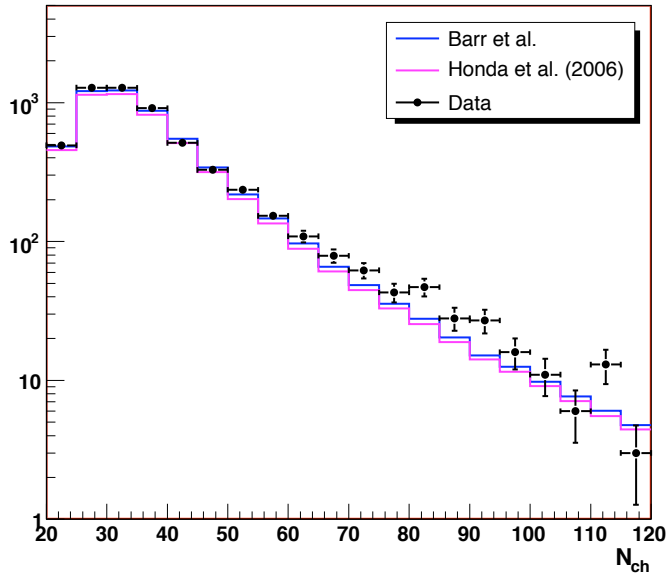

Figure 6.4: Unblinded N_{ch} distribution for 2000-2006 with original purity cuts, showing excess at $N_{ch} > 60$ relative to atmospheric neutrino predictions.

the detector) and appear preferentially in the clean ice regions, but we did not notice other distinguishing topologies (such as association with certain strings or OMs);

- these poor events tend to have worse up- to down-going likelihood ratios; and
- perhaps most convincing, a significant fraction of the excess appears to have higher paraboloid error and JAMS/Pandel space angle difference.

To illustrate the last point, in fig. 6.5 we show the distribution of paraboloid error and JAMS/Pandel space angle for high-N_{ch} events. The excess is concentrated at poor values of both, which is what we would expect for mis-reconstructed background. Also, an excess is not characteristic of any of our signal hypotheses. Therefore, we have chosen to isolate and remove the excess.

In order to isolate the excess, we use another point from the list above: the excess is concentrated at poor up-to-down likelihood ratio (see fig. 6.6). We can first roughly isolate the population using their likelihood ratio, and only then apply any cuts to paraboloid error and space angle difference. Since the likelihood ratio ($\mathrm{LR} = \log \mathsf{L}_{\mathrm{UL}}/\mathsf{L}_{\mathrm{BL}}$) is highly dependent on the zenith angle, instead of using a constant value we follow the median LR as a function of $\cos\theta_{\mathrm{UL}}$, as derived from MC. This

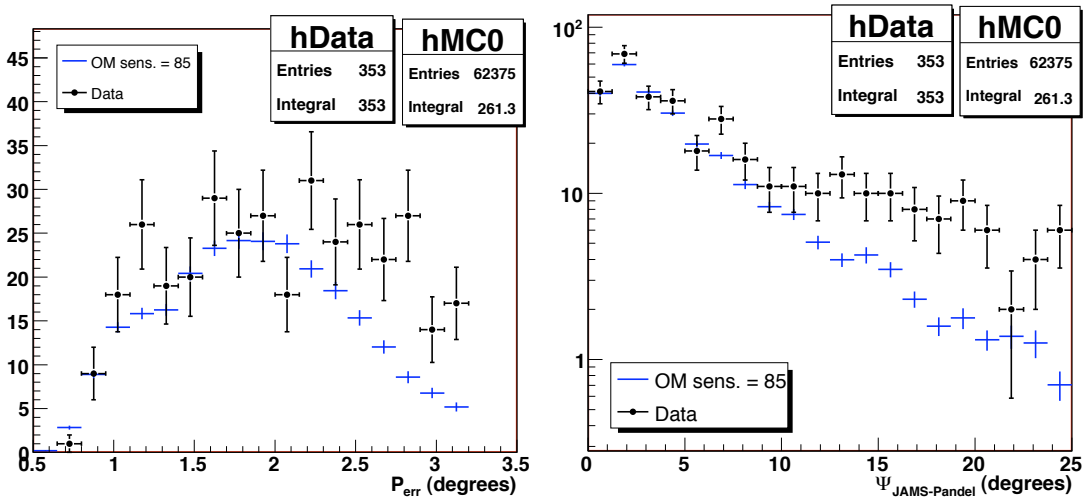

Figure 6.5: Paraboloid error (left) and JAMS/Pandel space angle (right) for events with $N_{ch} > 65$, after applying the original purity cuts. Note that the excess is concentrated in poor regions of these two quality parameters.

dependence, which we refer to as $\mathrm{LR}_{\mathrm{med}}(\theta)$, is shown graphically in fig. 6.7 and explicitly formulated in table 6.3.

To determine how to tighten the cuts as a function of N_{ch}, we examine two-dimensional plots of paraboloid error and space angle vs. N_{ch}, only for the events below $\mathrm{LR}_{\mathrm{med}}(\theta)$. A cut is shown superimposed on the distributions for atmospheric MC, data, and the ratio of data to atmospheric MC in figures 6.8 and 6.9. These contours define our level 2 purity cuts and are listed in table 6.3.

Table 6.3: Level 2 purity cuts defined after unblinding to remove high-N_{ch} background contamination. The combined criterion keeps events that exceed the median likelihood ratio **OR** pass both the space angle and paraboloid error conditions.

Description	**Criterion (to keep)**
Median LR	$(\cos\theta_{\mathrm{UL}} < -0.7) \wedge (\mathrm{LR} > -20/0.3 \cdot (\cos\theta_{\mathrm{UL}} + 0.7) + 52) \vee$ $(\cos\theta_{\mathrm{UL}} \geq -0.7) \wedge (\cos\theta_{\mathrm{UL}} < -0.4) \wedge (\mathrm{LR} > 52) \vee$ $(\cos\theta_{\mathrm{UL}} \geq -0.4) \wedge (\mathrm{LR} > -20/0.4 \cdot (\cos\theta_{\mathrm{UL}} + 0.4) + 52)$
Paraboloid error	$\sigma < (-1.1^\circ/30 \cdot (N_{ch} - 50) + 3.2^\circ) \vee$ $(N_{ch} > 80) \wedge (\sigma < 2.1^\circ)$
Space angle	$\Psi < (-15^\circ/30 \cdot (N_{ch} - 50) + 25^\circ) \vee$ $(N_{ch} > 80) \ \wedge \ (\Psi < 10^\circ)$

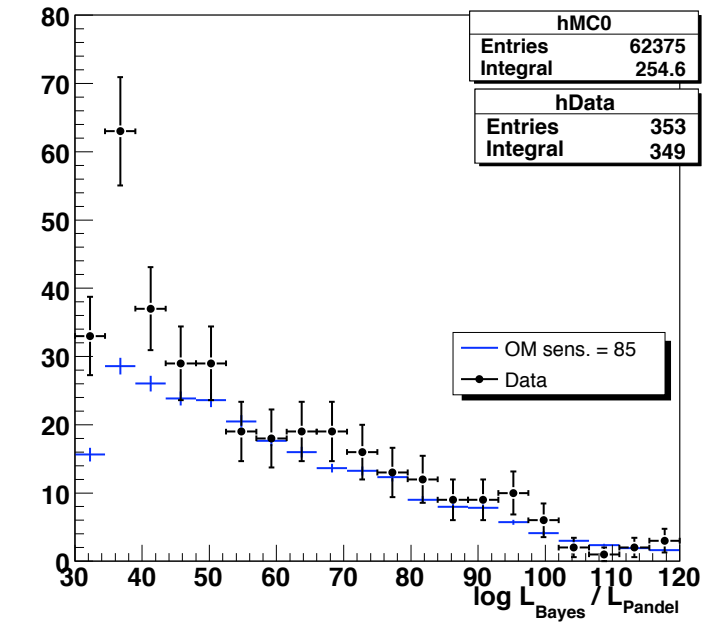

Figure 6.6: Up- to down-going likelihood ratio for high-N_{ch} events, showing excess at low values (more likely to be down-going).

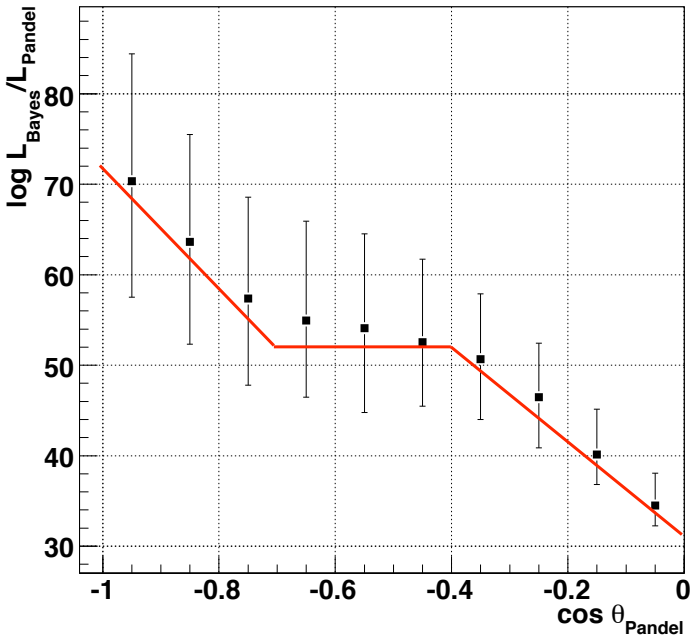

Figure 6.7: Median profile of Bayes/Pandel likelihood ratio as a function of zenith angle, for atmospheric neutrino MC events with $N_{ch} > 50$. Error bars give the 25%-75% range. The red line indicates our parametrization, denoted $\mathrm{LR}_{\mathrm{med}}(\theta)$.

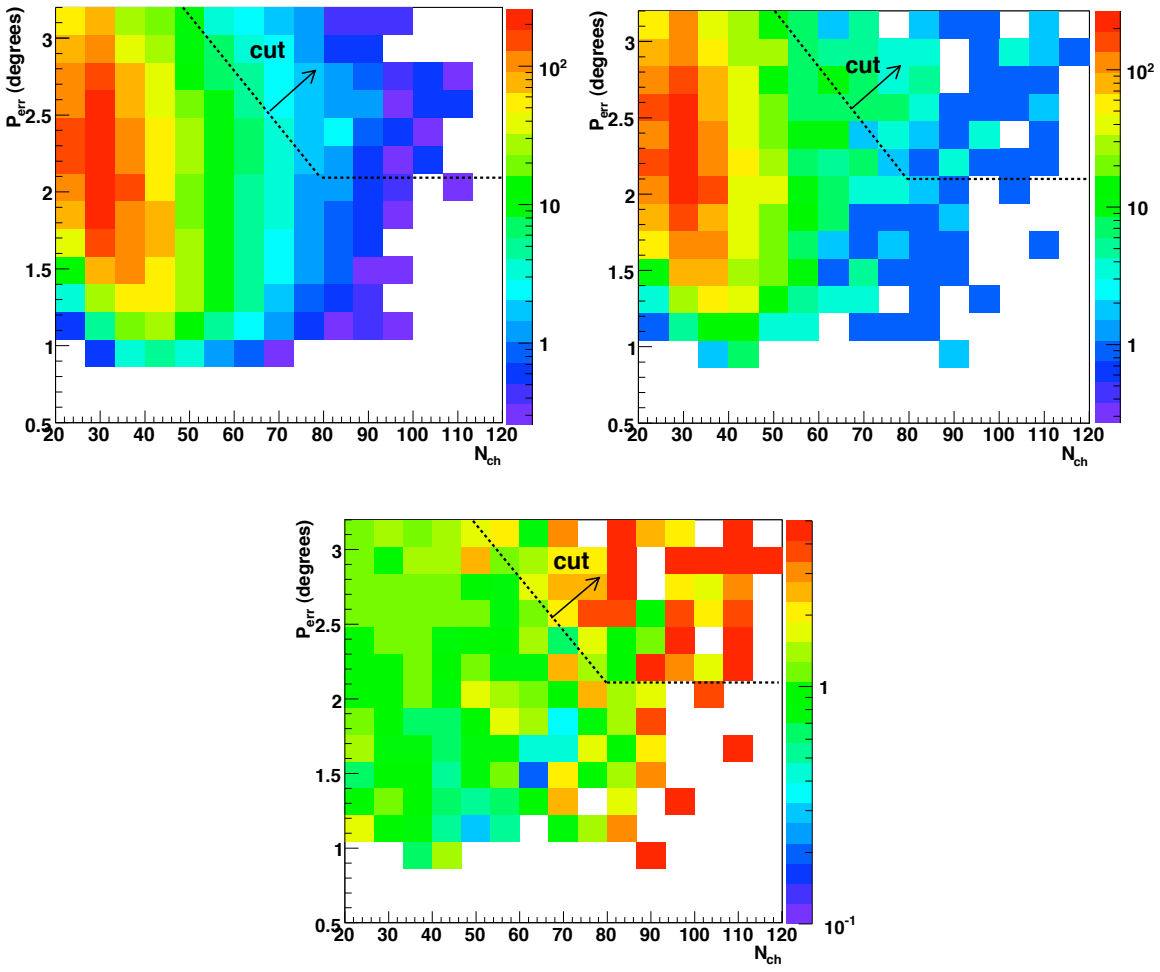

Figure 6.8: Paraboloid error vs. N_{ch} for atmospheric neutrino MC (upper left), data (upper right), and the ratio of data to MC. The 2-dimensional purity cut is shown with the dotted line on each plot.

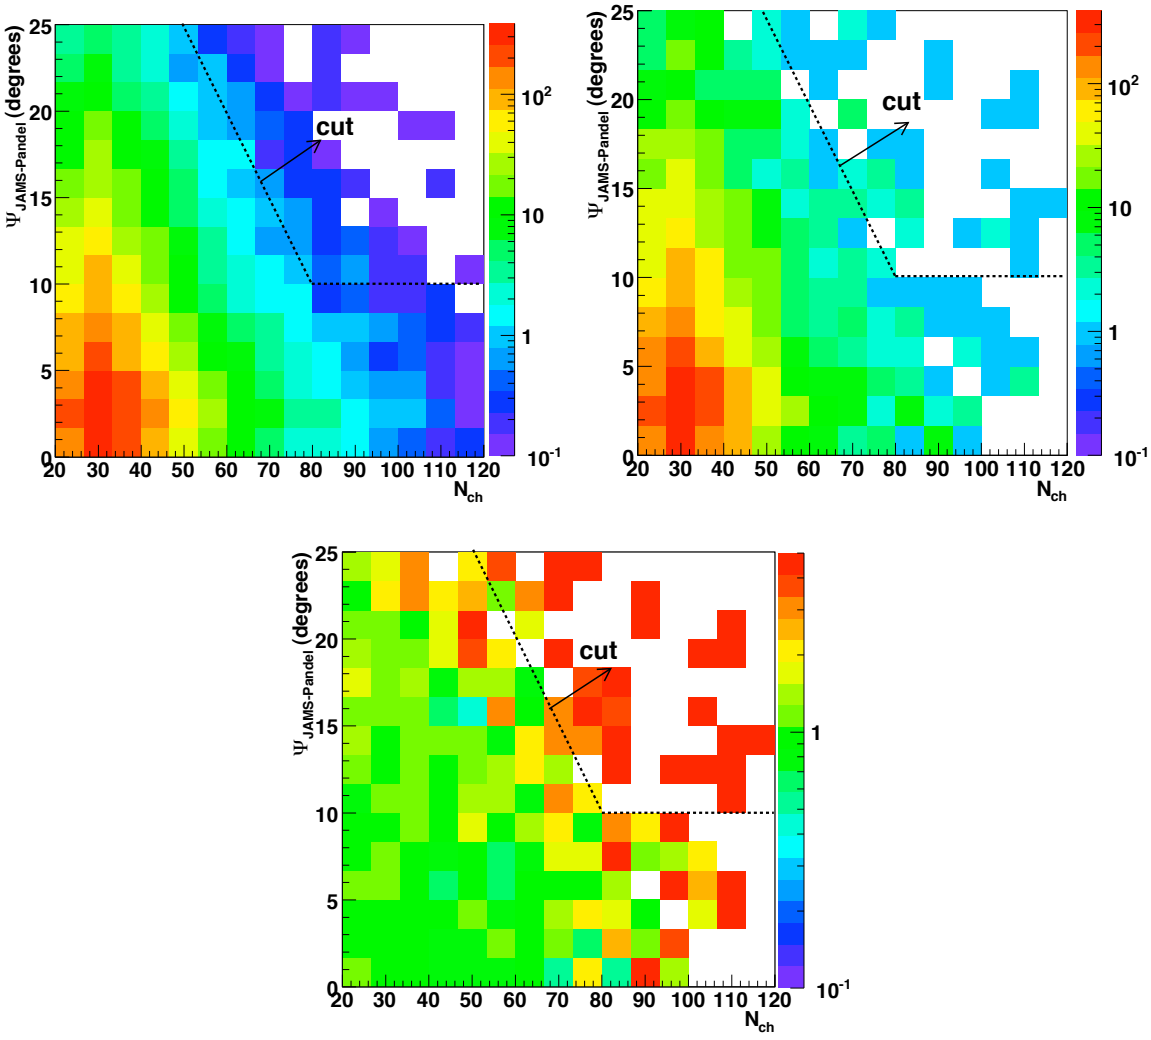

Figure 6.9: JAMS/Pandel space angle vs. N_{ch} for atmospheric neutrino MC (upper left), data (upper right), and the ratio of data to MC. The 2-dimensional purity cut is shown with the dotted line on each plot.

Table 6.4: Summary of livetime and events remaining after each filtering step for 2000 to 2006. The livetime of each year corresponds to the filtered data sample, whereas the raw event totals are for the entire year, including unstable periods.

	2000	**2001**	**2002**	**2003**	**2004**	**2005**	**2006**
Livetime (d)	197	193	204	213	194	199	187
Raw events	1.37 B	2.00 B	1.91 B	1.86 B	1.72 B	2.06 B	2.00 B
L1	45.4 M	81.8 M	68.3 M	65.3 M	60.8 M	184 M	177 M
L2	5.50 M	6.87 M	7.59 M	8.02 M	7.47 M	—	—
L3	1.63 M	1.90 M	2.10 M	2.22 M	2.09 M	5.21 M	4.89 M
Point src.	560	799	976	1034	966	897	934
Purity	516	750	908	966	901	810	835
Purity L2	504	730	884	953	883	780	810

Table 6.5: Total 2000-2006 livetime and events remaining after each filtering step.

	Total
Livetime (d)	1387
Raw events	12.9 B
L1	683 M
L3	20.0 M
Point src.	6166
Purity	5686
Purity L2	**5544**

6.4 Final Data Sample

After the level 2 purity cuts, we are left with 5544 candidate neutrino events below the horizon. In table 6.4, we show the livetime and various numbers of events remaining after each filtering and cut level, with 7-year totals shown in table 6.5. The livetime of the detector accounts for both excluded data-taking periods as well as inherent deadtime due to the DAQ itself. We also note that the larger number of filtered events in 2005 and 2006 is due to additional filter streams added.

We also show the distributions of our primary quality variables at the final cut level in fig. 6.10. The efficiency of these selection criteria for simulated atmospheric neutrinos that trigger the detector is 24%.

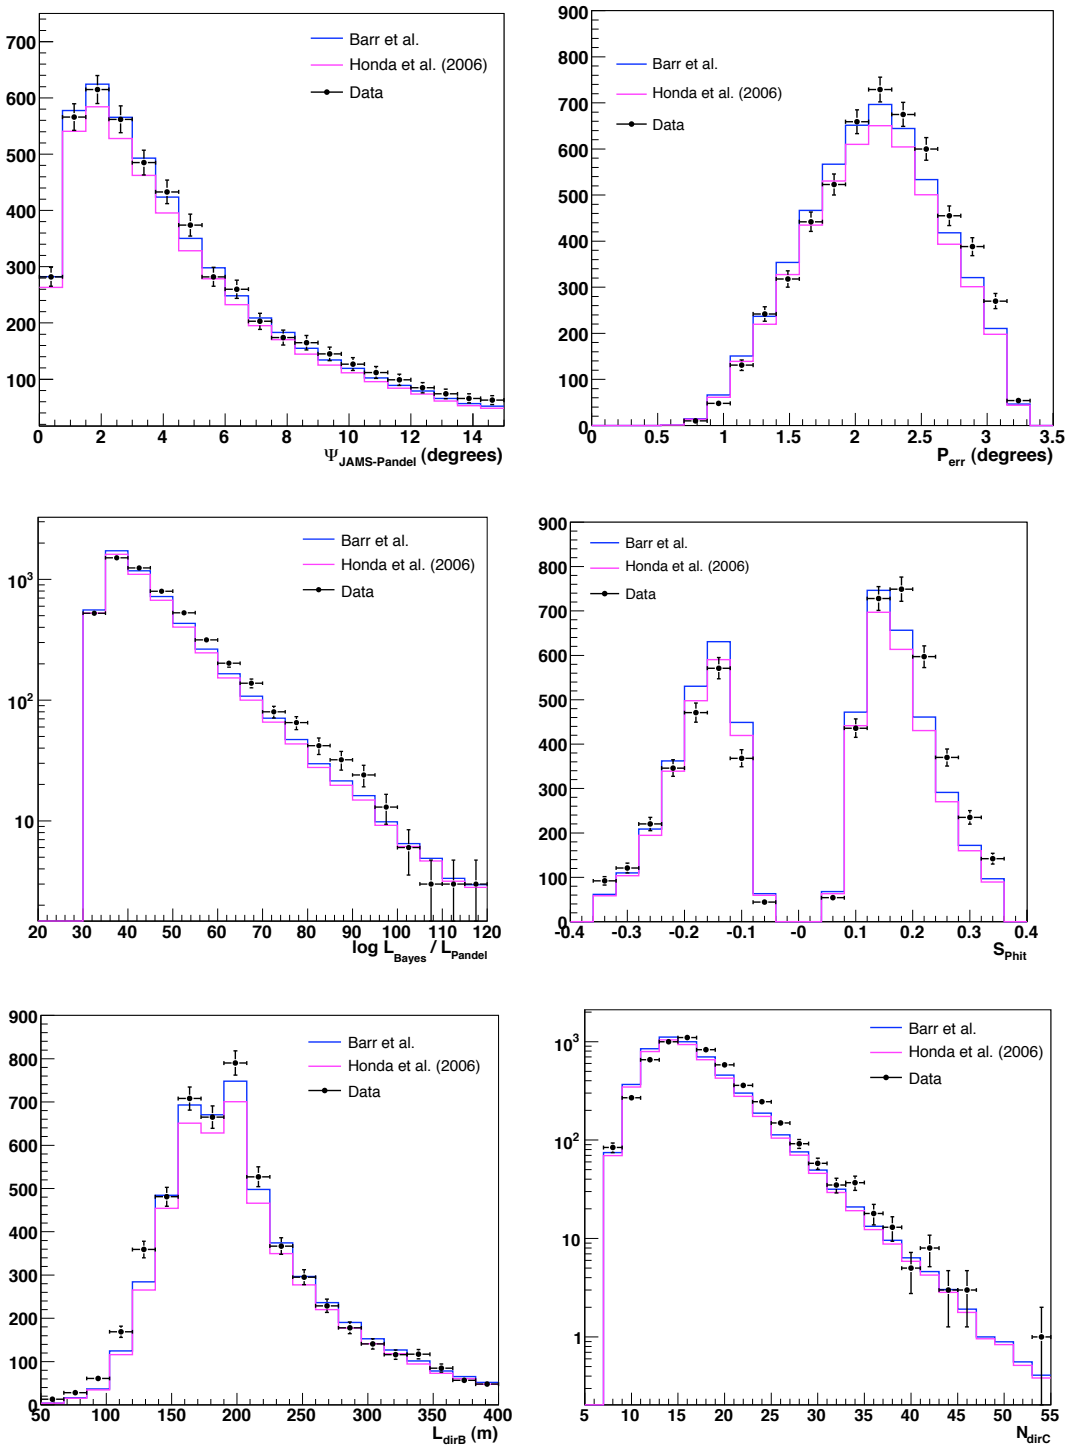

Figure 6.10: Track quality variables for data and atmospheric neutrino simulation after all purity cuts.

6.5 Effective Area

The effective area of a detector is the area $A_{\text{eff}}(E, \theta, \phi)$ of a corresponding ideal detector with 100% detection efficiency. For neutrino telescopes, because of the low interaction probability, the neutrino effective area A^{ν}_{eff} is much smaller than the physical cross-section of the detector. The effective area encapsulates all efficiencies in a particular analysis; this includes not only the efficiency of quality cuts, but also physical effects like Earth absorption and oscillations (we can consider the Earth as part of the detector). For a discussion of calculation techniques, see appendix B. The neutrino effective area for various zenith angle ranges at the final cut level is shown in fig. 6.11, along with the ratio of effective area of $\bar{\nu}_\mu$ to ν_μ.

6.6 A Note on Blindness

The original purity cuts designed for this analysis and described in section 6.3.2 were designed in a blind fashion. By this we mean that the observables for our analysis, N_{ch} and the reconstructed zenith angle, were removed from the files and not used to design the cuts, in order to limit the chance of biasing the results and forcing agreement with one hypothesis or the other.

Such a procedure works well when background can easily be determined from the data itself: for example, in a point-source search where the background is determined by looking off-source. However, for an analysis in which eliminating background is crucial, and simulating the final 1% of that background is not feasible, we find this blindness procedure of limited usefulness.

First, the point-source cuts were designed with the zenith angle unblinded, and so by using these cuts we are arguably not blind in this variable to begin with. Furthermore, blinding the N_{ch} distribution simply prevented us from characterizing and eliminating the background, and we eventually had to alter the cuts in a non-blind way. In cases such as this where the background cannot be measured from the data, it is crucially important to understand the background events and the detector response, and we argue that by careful justification of any non-blind cuts, we can still achieve meaningful results.

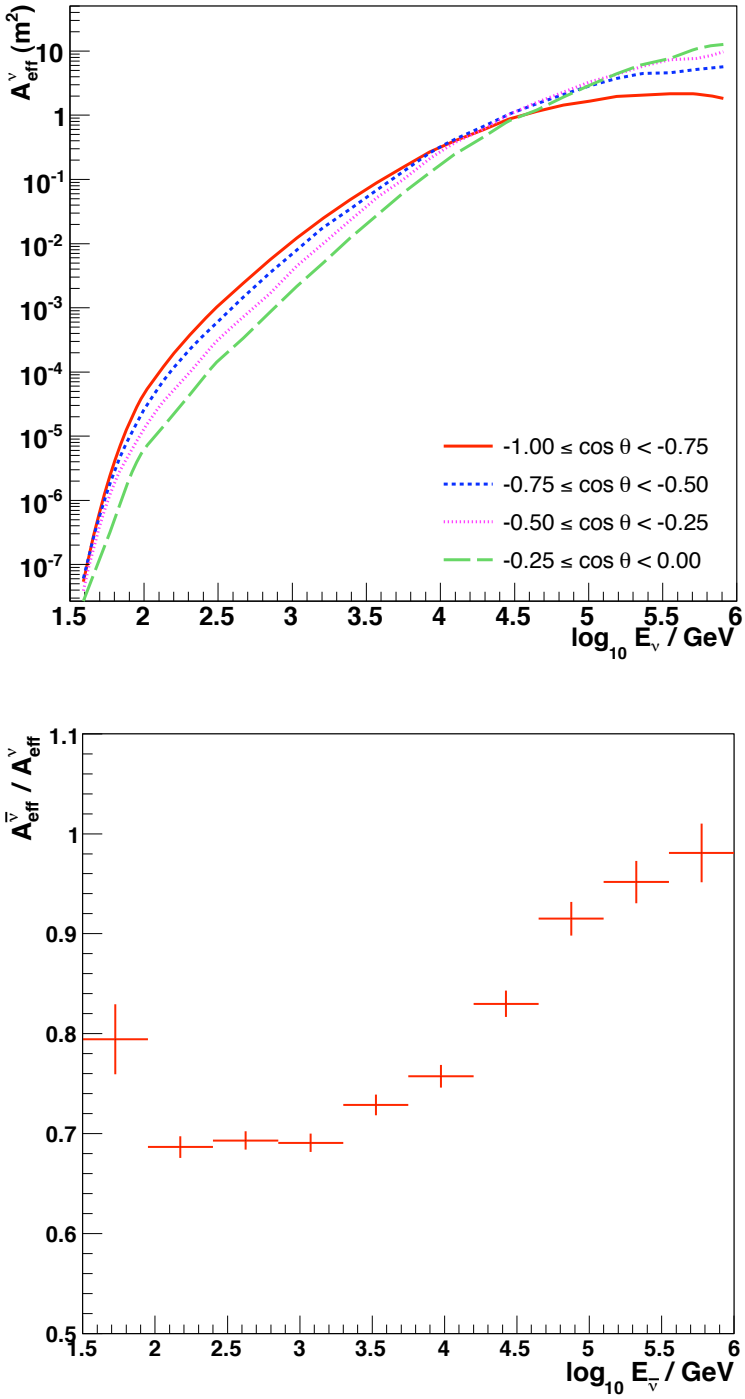

Figure 6.11: Neutrino effective area as a function of neutrino energy. Top: effective area for ν_μ for various zenith angle ranges. Bottom: ratio of angle-averaged effective area for $\bar{\nu}_\mu$ to ν_μ.

Chapter 7

Analysis Methodology

7.1 Observables

As described in chapter 3, the signature of a flavor-changing new physics effect such as VLI or QD is a deficit of ν_μ events at the highest energies and longest baselines (i.e., near the vertical direction). For our directional observable, we use the cosine of the reconstructed zenith angle as given by the UL fit, $\cos\theta_{\rm UL}$ (with -1 being the vertical up-going direction). As reconstructing the neutrino energy from the energy loss of the through-going muon is difficult, we use instead an energy-correlated observable, the number of OMs (or channels) hit, N_{ch}.

The simulated energy response to the Barr *et al.* atmospheric neutrino flux [16] (without any new physics) is shown in fig. 7.1. For this flux, the simulated median energy of the final event sample is 640 GeV, and the 5%-95% range is 105 GeV to 8.9 TeV. Fig. 7.2 shows the median neutrino energy for a given event N_{ch}. Fig. 7.3 shows the simulated effects of QD and VLI on both the zenith angle and N_{ch} distributions: a deficit of events at high N_{ch} and at more vertical directions. Because the N_{ch} energy estimation is approximate, the VLI oscillation minima are smeared out, and the two effects look similar in the observables. Furthermore, the observable minima are not exactly in the vertical direction. This is because the N_{ch}-energy relationship varies with zenith angle, since the detector is taller than it is wide, and events off the vertical actually have a higher median energy.

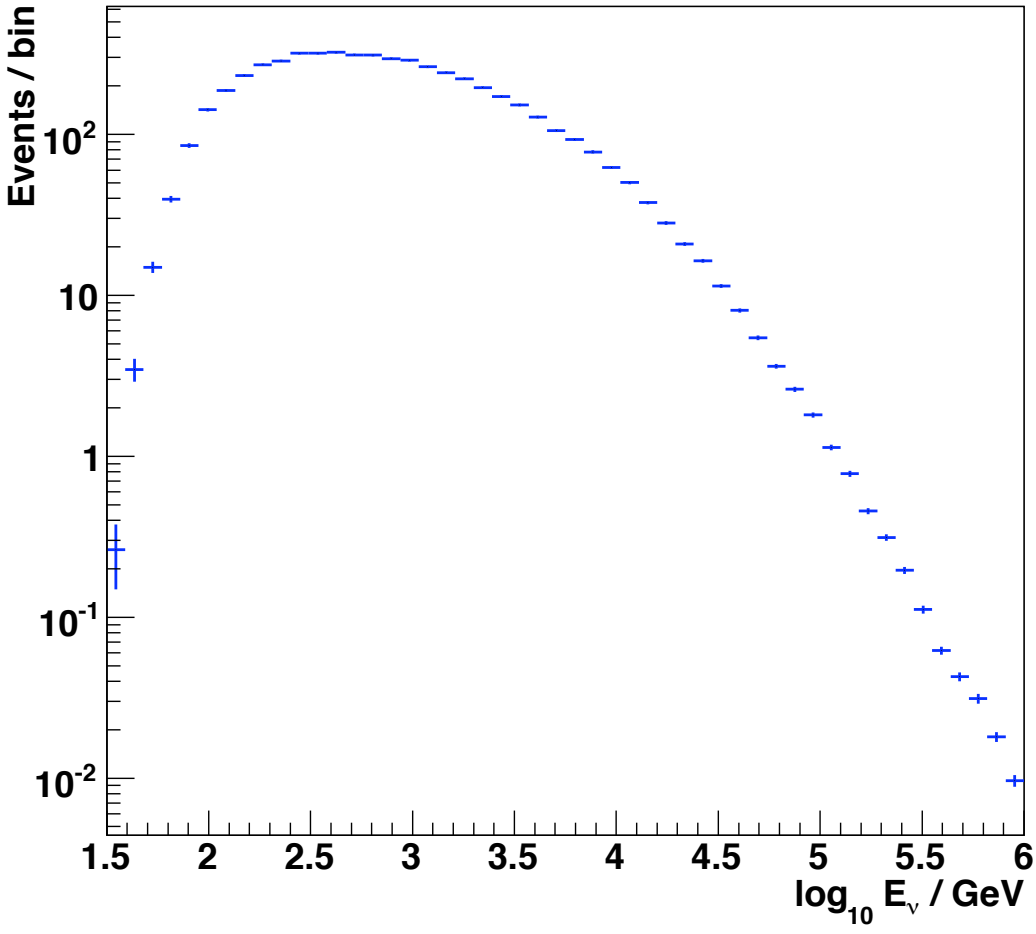

Figure 7.1: Simulated $\nu_\mu + \bar{\nu}_\mu$ energy distribution of the final event sample, assuming the Barr *et al.* input spectrum.

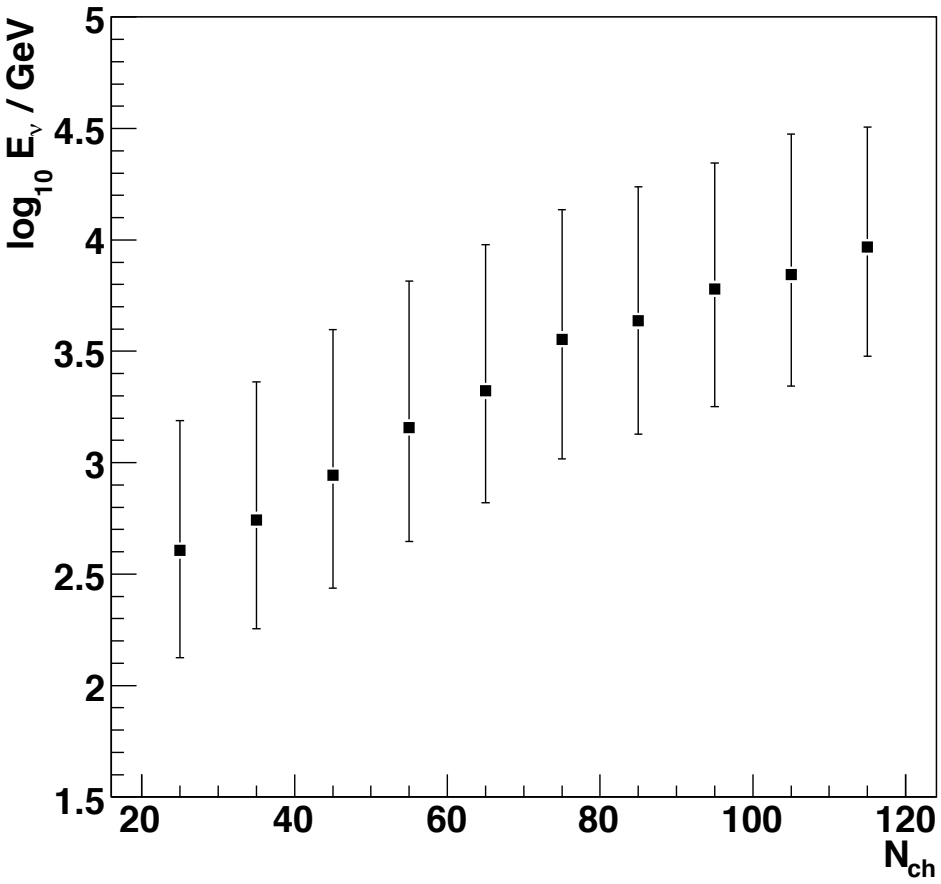

Figure 7.2: Simulated profile of median neutrino energy versus number of OMs hit (N_{ch}) over all zenith angles below the horizon ($\pm 1\sigma$ error bars).

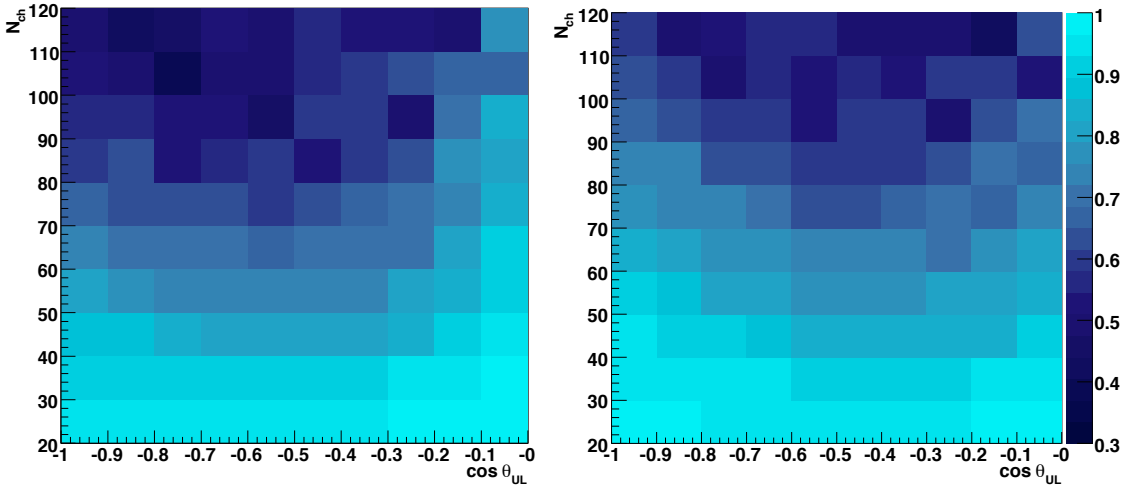

Figure 7.3: Simulated effects of VLI (left, with $n = 1$, $\sin 2\xi = 1$, and $\Delta\delta = 10^{-26}$) and QD (right, with $n = 2$ and $D^* = 10^{-30}$ GeV^{-1}) on the zenith angle and N_{ch} distributions, shown as the ratio to conventional oscillation predictions.

7.2 Statistical Methods

The goal of our analysis is to quantify whether our data are well-described by certain classes of hypotheses; namely, violation of Lorentz invariance, quantum decoherence, or the Standard Model. Each of the general scenarios above — VLI, for example — can be further expanded into a set of specific hypotheses, each with different parameters. We need a methodology that will allow us to test each point in this parameter space, and then define a region of this space which is allowed or excluded at a certain confidence level. We would also like to include the effect of systematic errors into this methodology.

In this chapter, we describe a frequentist method of defining central confidence intervals that incorporates systematic errors. This method, the *profile construction* method, is an extension by G. Feldman of the frequentist approach described in his paper with R. Cousins [85]. Originally described (albeit very tersely) in [86], it has only recently been applied to physics analyses.

7.2.1 Likelihood Ratio

We first define a test statistic to compare our observables x for various hypotheses, characterized by physics parameters θ_r. For a binned distribution, a natural choice arises from the Poisson probability (or likelihood)

$$P(x|\theta_r) = \prod_{i=1}^{N} e^{-\mu_i} \frac{\mu_i^{n_i}}{n_i!} \ , \tag{7.1}$$

where we form a product over the N bins of our observable(s) x, and in each bin of the data we see n_i counts on an expected μ_i for the hypothesis we are testing with parameters θ_r. At this point, it is conventional to switch to the negative logarithm (the *log likelihood*):

$$\mathcal{L}(\theta_r) = -2 \ln P = 2 \sum_{i=1}^{N} \left(\mu_i - n_i \ln \mu_i + \ln n_i! \right) \ . \tag{7.2}$$

We will come back to the additional factor of 2.

To compare the probabilities of two hypotheses, H_1 and H_2, of generating our observed data, we take the *likelihood ratio* (or, working with the logarithm, the difference):

$$\mathcal{L}(\theta_{r1}) - \mathcal{L}(\theta_{r2}) = 2 \sum_{i=1}^{N} \left(\mu_{1,i} - \mu_{2,i} + n_i \ln \frac{\mu_{i,2}}{\mu_{i,1}} \right) \ , \tag{7.3}$$

where hypothesis H_1 with parameters θ_{r1} gives us an expected count $\mu_{1,i}$, and hypothesis H_2 with parameters θ_{r2} gives us an expected count $\mu_{2,i}$, and again we have observed n_i counts in a given bin. Using this, our test statistic compares the hypothesis at a point θ_r to the hypothesis that fits the data the best. Specifically, in the physics parameter space θ_r, the test statistic is the difference of the log likelihood at this point to the best-fit hypothesis with parameters $\hat{\theta}_r$ ($\mathcal{L}$ is minimized[1] by $\hat{\theta}_r$):

$$\Delta\mathcal{L}(\theta_r) = \mathcal{L}(\theta_r) - \mathcal{L}(\hat{\theta}_r) \ . \tag{7.4}$$

The additional factor of 2 added in equation 7.2 arises because in the Gaussian regime, $\Delta\mathcal{L}$ so defined approaches a χ^2 distribution with degrees of freedom equal to the dimension of θ_r (Wilks' Theorem).

[1] By minimizing the negative log likelihood, we maximize the probability.

7.2.2 Confidence Intervals

At this point, we wish to examine all the physically allowed hypotheses by iterating over the space θ_r, and determine which are allowed given our observation x. It is not uncommon to use Wilks' Theorem and define confidence intervals using a χ^2 distribution. Specifically, one calculates $\Delta\mathcal{L}$ at every point θ_r, and for a given confidence level (CL) α, the allowed region is the set

$$\{\theta_r\}_\alpha = \{\theta_r \mid \Delta\mathcal{L}(\theta_r) < \chi^2(\alpha, \dim\, \theta_r)\} \ . \tag{7.5}$$

For two parameters and a 90% confidence level, we would allow the region where $\Delta\mathcal{L} < 4.61$. This is known as the *global scan* method.

As demonstrated in [85], the global scan method has several disadvantages when the likelihood varies in a complicated way over the parameter space. In particular, $\Delta\mathcal{L}$ can deviate from the simple χ^2 distribution by a significant amount if, for example, one of the parameters is extended into a region that has little effect on the observables. In this case, the effective dimensionality of θ_r is reduced and the χ^2 used has too many degrees of freedom. In this case, we prefer a frequentist approach to define the confidence intervals that takes this and other issues into account to achieve proper coverage.

Specifically, at each point in the parameter space θ_r, we perform a number of Monte Carlo experiments where we sample from the parent distribution $\{x \mid \theta_r\}$ and then calculate the likelihood ratio $\Delta\mathcal{L}_i$ for the experiment. The sampling to generate the MC "data" for an experiment can be achieved a number of ways. We choose to first select the total number of events N from a Poisson distribution with μ equal to the integral of the parent distribution, and then sample N times from the parent observable distribution (as a probability density function) to find the observables for each MC event.

The set of $\{\Delta\mathcal{L}_i\}$ from the MC experiments allows us to see how our test statistic behaves under statistical variations only. To define our confidence intervals at CL α, we find the critical value $\Delta\mathcal{L}_{\text{crit}}$ such that

$$\left(\int_0^{\Delta\mathcal{L}_{\text{crit}}} \Delta\mathcal{L}_i\right) \Big/ \left(\int_0^{\infty} \Delta\mathcal{L}_i\right) = \alpha \ , \tag{7.6}$$

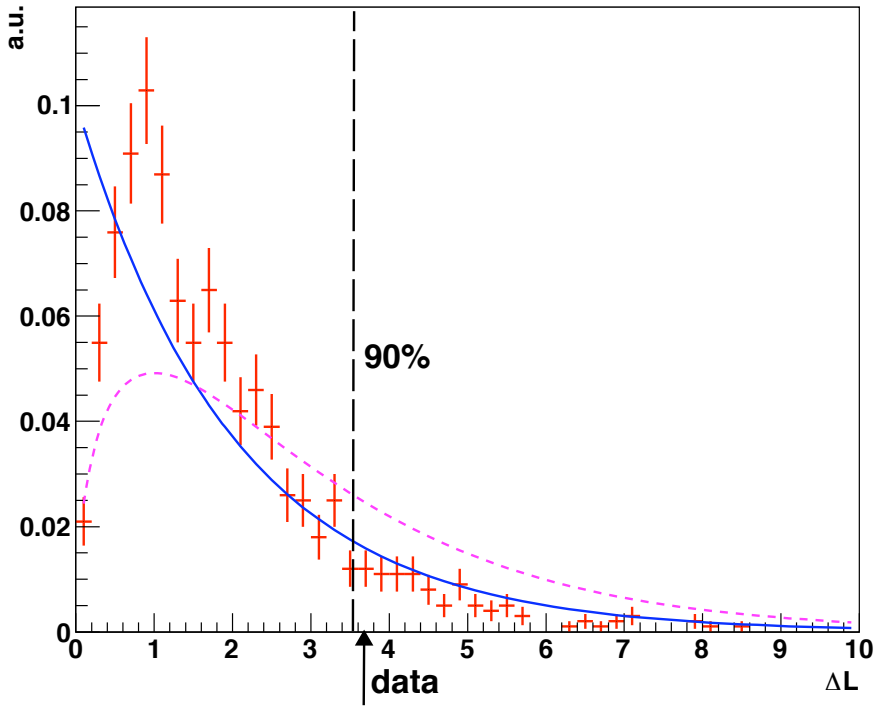

Figure 7.4: Simulated likelihood ratio distribution (red points) for a point in VLI parameter space, from 1000 MC experiments. Shown for comparison are χ^2 distributions for two (solid line) and three (dashed line) degrees of freedom. The vertical line marks the 90% critical value, and the arrow marks the likelihood ratio for the data, indicating that this hypothesis is excluded at the 90% CL.

and our acceptance region is the set $\{\theta_r\}$ where $\Delta\mathcal{L}_{\rm data}(\theta_r) < \Delta\mathcal{L}_{\rm crit}(\theta_r)$. As an illustration, fig. 7.4 shows a simulated $\Delta\mathcal{L}$ distribution along with the 90% critical value. By employing the $\Delta\mathcal{L}$ distribution to determine the confidence level, we have used the likelihood ratio as an *ordering principle* to sort the possibilities into increasing statistical significance. We also point out that the exclusion region at CL α is simply the complement of this set, as acceptance / exclusion is defined by which side of the critical value one is on.

7.3 Incorporating Systematic Errors

Unfortunately, the above procedure does not incorporate any kind of systematic errors. In statistical terms, a systematic error can be treated as a *nuisance parameter*: a parameter that one must know to calculate the expected signal, but the value of which is not important to the result.

The likelihood depends now on both physics parameters θ_r and nuisance parameters θ_s, but one needs to "project out" any confidence intervals into only the θ_r space.

The key to this procedure is to use an approximation for the likelihood ratio that, in a sense, uses the worst-case values for the nuisance parameters θ_s — the values which make the data fit the hypothesis the best at that point θ_r. Mathematically, we find the best values for θ_s in both the numerator and the denominator of the likelihood ratio

$$\Delta\mathcal{L}_p(\theta_r) = \mathcal{L}(\theta_r, \hat{\hat{\theta}}_s) - \mathcal{L}(\hat{\theta}_r, \hat{\theta}_s) \ , \tag{7.7}$$

where we have globally minimized the second term, and we have conditionally minimized the first term, keeping θ_r fixed but varying the nuisance parameters to find $\hat{\hat{\theta}}_s$. This statistic is called the *profile likelihood*.

The profile likelihood is used in combination with the χ^2 approximation in the "MINOS" method in the MINUIT suite [87] and is also explored in some detail by Rolke *et al.* in [88, 89]. To extend the Feldman-Cousins frequentist construction to the profile likelihood, we follow the method suggested by Feldman [90]: we perform Monte Carlo experiments as before, but instead of iterating through the entire (θ_r, θ_s) space, at each point in the physics parameter space θ_r, we fix θ_s to its best-fit value from the *data*, $\hat{\hat{\theta}}_s$. Then we recalculate the profile likelihood for the experiment as defined in equation 7.7. As before, this gives us a set of likelihood ratios $\{\Delta\mathcal{L}_{p,i}\}$ with which we can define the critical value for a CL α that depends only on θ_r.

To summarize, we describe the procedure step-by-step:

1. The test statistic / ordering principle is the profile likelihood $\Delta\mathcal{L}_p$ as defined in eq. 7.7.

2. The profile likelihood for the data is calculated at each point θ_r, with the numerator being a conditional minimum at $(\theta_r, \hat{\hat{\theta}}_s)$ and the denominator the global minimum at some $(\hat{\theta}_r, \hat{\theta}_s)$.

3. For each point θ_r, we perform a number of Monte Carlo experiments in which we sample from the parent distribution $\{x \mid \theta_r, \hat{\hat{\theta}}_{s,\mathrm{data}}\}$, then we recalculate the profile likelihood for each experiment.

4. For a CL α, at each point we find the critical value $\Delta\mathcal{L}_{p,\mathrm{crit}}(\theta_r)$ using eq. 7.6, and this point is in the allowed region if $\Delta\mathcal{L}_{p,\mathrm{data}}(\theta_r) < \Delta\mathcal{L}_{p,\mathrm{crit}}(\theta_r)$.

7.4 Discussion

We note that the problem of incorporating systematic errors into confidence intervals is still an area of active research. For a survey of recent approaches, including hybrid Bayesian-frequentist methods not discussed here, see [91]. Another approach which uses random variations in the MC experiments to model systematic errors is presented in [92][2]. Two fully frequentist constructions (not using the profile likelihood approximation) have been employed in test cases by G. Punzi [93] and K. Cranmer [94], but there is not a general consensus on an ordering principle. For further information, we refer the reader to the discussion by Cranmer in [95][3].

7.5 Complications

7.5.1 Computational Requirements

The primary drawback of the profile construction method is its computational requirements. For a given experiment, locating the likelihood minimum in a multi-dimensional parameter space is time-consuming. Furthermore, to define the confidence regions, we need hundreds or thousands of Monte Carlo experiments *at each point* in the parameter space. We have employed a few tricks to make the problem manageable, which we describe below. These may or may not be applicable in other implementations.

1. **Pre-generation of histograms.** The method requires a histogram of the observables at each point in parameter space (that is, for every value of the physics and nuisance parameters). Generating such a weighted Monte Carlo distribution can be time-consuming, so pre-computing the histograms and saving/loading the binned results is faster. This computation can also be performed in parallel on a cluster.

[2]This reference also provides an exceptionally clear description of the Feldman-Cousins method.
[3]We also note this as the origin of the term "profile construction" to describe this method.

2. **Parameter space gridding.** Pre-computing the observable histograms requires one to operate on a discrete grid. In practice, the physics parameters should be binned somewhat finely, and the nuisance parameters can be binned more coarsely. Our grids contained a total of approximately 5×10^4 points for each class of hypothesis ($n = 1$ VLI, for example).

3. **Likelihood minimization.** The downside to using a grid is that most minimization algorithms require continuous parameters. Because of this, and to avoid any issue with local minima, we chose simply to perform an exhaustive search for the minimum when required.

4. **Normalization nuisance parameter.** Minimization of a normalization nuisance parameter (see chapter 8 for specifics) can be handled differently, because it can be easily varied via histogram scaling. The likelihood can be automatically minimized in this dimension by normalization of the total number of events[4].

5. **Histogram interpolation.** One of our nuisance parameters, the OM sensitivity, is not variable via reweighting; changing it requires re-simulating neutrino Monte Carlo from the detector simulation stage onward. If we need histograms with OM sensitivities between those generated, we linearly interpolate between the observable histograms of higher and lower sensitivity as an approximation.

6. **MC Experiments and confidence levels.** As we perform the Monte Carlo experiments, we estimate not only the confidence level of the data but also the error on that confidence level. If that confidence level is far enough away from the confidence levels we care about (generally 90% to 99%), we abort the MC experiments early and record the approximate confidence level. Furthermore, recording the actual confidence level at a certain point (instead of just a yes/no of whether the point is excluded) allows offline contour interpolation finer than the initial binning.

Several improvements are possible as computational power increases and/or specific implementation details change. First, if observable histogram generation is fast enough, the grid approach can be abandoned and a minimization algorithm such as MINUIT [87] can be employed. Alternatively,

[4]This can be easily proved by differentiating the Poisson log-likelihood with respect to an overall scaling factor on the μ_i.

an adaptive grid algorithm could be used to avoid unnecessary time spent in uninteresting portions of the parameter space.

7.5.2 Zero Dimensions

One advantage of the Feldman-Cousins method is that, unlike the χ^2 approximation, the effective dimensionality of the parameter space is not fixed. Specifically, if one enters a region of the parameter space where varying a parameter doesn't affect the observables very much, the critical value of the likelihood ratio will tend toward a χ^2 distribution with fewer degrees of freedom.

This results in a mild pathology, however, in the case where our physics parameter space has only one dimension to begin with. Suppose one is parametrizing new physics with the logarithm of a single parameter, $\log_{10} \Delta\delta$. Then as this gets smaller and smaller, eventually there may be no effect on the observables at all, and the $\Delta\mathcal{L}$ distribution representing statistical variations will approach a δ-function (a χ^2 distribution with zero degrees of freedom). This may not be a problem if the data is perfectly described by the simulation; however, any small differences may mean that a relatively small $\Delta\mathcal{L}_{\text{data}}$ in this region of the parameter space may be artificially blown up into a huge significance (see fig. 7.5 for an example). We view this more as a coordinate singularity brought on by infinitely stretching out the region between no new physics ($\Delta\delta = 0$) and the $\Delta\delta$ to which we are sensitive. Our solution was to avoid this region of the parameter space in the one-dimensional case, but we encourage further studies of the effect.

7.6 Binning and Final Event Count

In general, finer binning provides higher sensitivity with a likelihood analysis. Indeed, we find a monotonic increase in simulated sensitivity to VLI effects while increasing the number of bins in $\cos\theta_{\text{UL}}$ and N_{ch} (see fig. 7.6). However, because the further gains in sensitivity are minimal with binning finer than 10×10, we limit ourselves to this size to avoid any systematic artifacts caused by binning, say, finer than our angular resolution. We also limit the N_{ch} range for the analysis to $20 \leq N_{ch} < 120$ in order to avoid regions with very poor statistics, limiting the possibility that a few remaining high-energy background events might pollute the analysis. This reduces the number of candidate neutrino events in the analysis region to 5511.

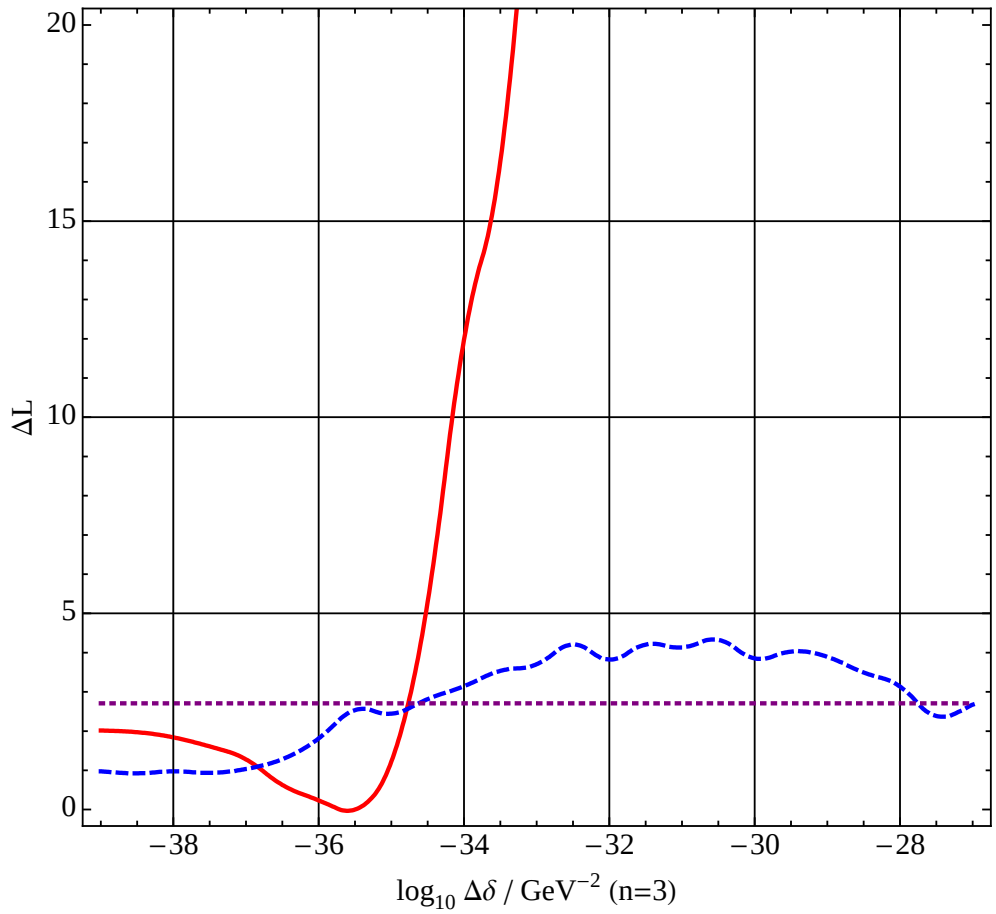

Figure 7.5: Likelihood ratio $\Delta\mathcal{L}$ for data (red solid line), shown with the 90% CL $\Delta\mathcal{L}_{\rm crit}$ (blue dashed line) and the 90% CL for a χ^2 distribution with one degree of freedom (purple dotted line), as a function of the VLI new physics parameter in the $n = 3$ case. The rightmost intersection of the red solid and blue dashed lines is the 90% CL upper limit, with all larger values of $\log_{10}\Delta\delta$ excluded. The leftmost intersection is a result of the collapse of the dimensionality of the space as discussed in the text. The point where $\Delta\mathcal{L}_{\rm data} = 0$ is the best-fit point.

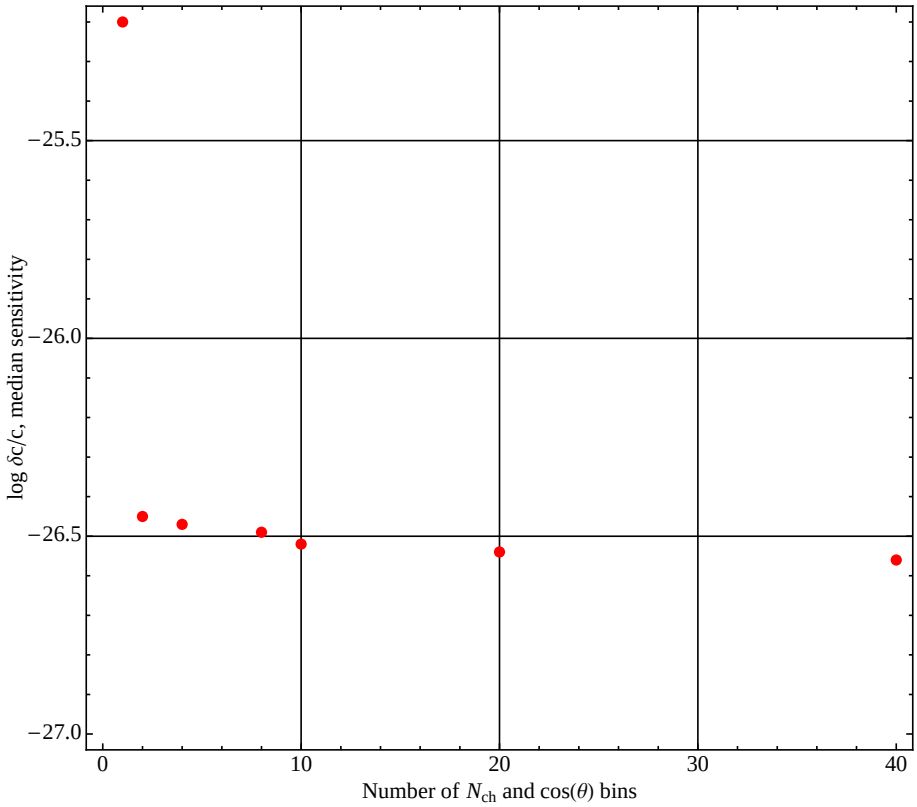

Figure 7.6: VLI sensitivity (using χ^2 approximation) versus number of bins in N_{ch} and $\cos\theta_{\rm UL}$.

Chapter 8

Systematic Errors

Systematic errors represent uncertainties in quantities necessary to predict our observables given a certain hypothesis. These can be physics-related (for example, the absolute normalization of the flux of atmospheric neutrinos) or detector-related (the sensitivity of the AMANDA-II optical modules). Here we quantify a number of these systematic errors and discuss how they are incorporated into this analysis.

8.1 General Approach

Each systematic error, or nuisance parameter, added to the likelihood test statistic increases the dimensionality of the space we must search for the minimum; therefore, to add systematic errors, we group them by their effect on the $(\cos\theta_{\mathrm{UL}}, N_{ch})$ distribution. We define the following four classes of errors: 1) *normalization* errors, affecting only the total event count; 2) *slope* errors, affecting the energy spectrum of the neutrino events and thus the N_{ch} distribution; 3) *tilt* errors, affecting the $\cos\theta_{\mathrm{UL}}$ distribution; and 4) *OM sensitivity* errors, which affect the probability of photon detection and change both the $\cos\theta_{\mathrm{UL}}$ and N_{ch} distribution. These errors are incorporated into the simulation as follows:

- Normalization errors are incorporated via a uniform weight $1 \pm \sqrt{(\alpha_1^2 + \alpha_2^2)}$;
- slope errors are incorporated via an energy-dependent event weight $(E/E_{\mathrm{median}})^{\Delta\gamma}$, where E_{median} is the median neutrino energy at the final cut level, 640 GeV (see fig. 8.1);

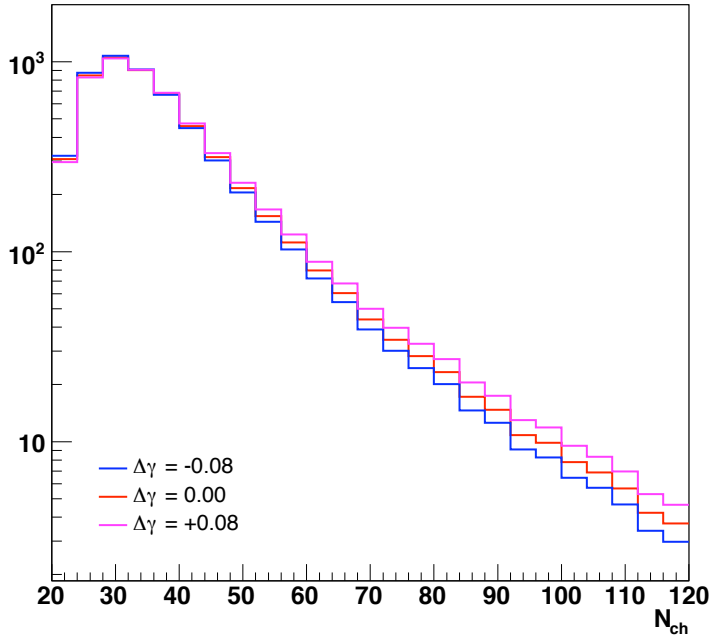

Figure 8.1: Simulated effect of slope uncertainty $\Delta\gamma$ on the N_{ch} distribution.

- tilt errors are incorporated by linearly tilting the $\cos\theta_{\rm UL}$ distribution via a factor $1+2\kappa(\cos\theta_{\rm UL}+\frac{1}{2})$ (see fig. 8.2); and
- OM sensitivity errors are incorporated by regenerating atmospheric neutrino simulation while globally changing the sensitivity of all OMs in the detector simulation from the nominal value by a factor $1+\epsilon$ (see fig. 8.3).

As we discuss later, we split the normalization error into two components, α_1 and α_2, to facilitate the determination of the conventional atmospheric flux. The slope error is normalized at the median energy to isolate slope changes from a change in normalization. The tilt term linearly tilts the $\cos\theta_{\rm UL}$ distribution around $\cos\theta_{\rm UL} = -0.5$, with the magnitude of κ corresponding to the percent change at $\cos\theta_{\rm UL} = 0$.

Table 8.1 summarizes sources of systematic error and the class of each error. The total normalization errors α_1 and α_2 are obtained by adding the individual normalization errors in quadrature, while the tilt κ and slope change $\Delta\gamma$ are added linearly. Asymmetric error totals are conservatively assumed to be symmetric, using whichever deviation from the nominal is largest. Each

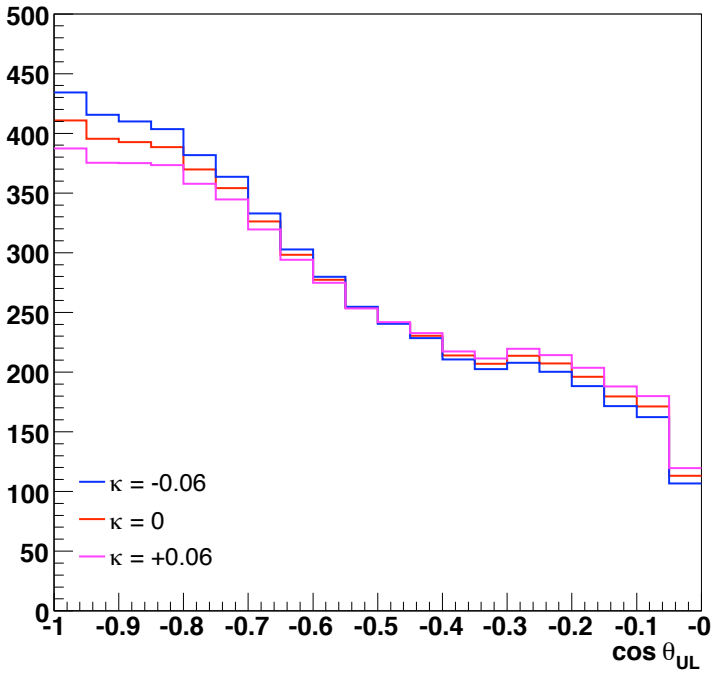

Figure 8.2: Simulated effect of tilt uncertainty κ on the $\cos\theta_{\rm UL}$ distribution.

class of error maps to one dimension in the likelihood space, so for example in the VLI case, $\mathcal{L}(\theta_r, \theta_s) = \mathcal{L}(\Delta\delta, \sin 2\xi, \alpha, \Delta\gamma, \kappa, \epsilon)$. During minimization, each nuisance parameter is allowed to vary freely within the range allowed around its nominal value, with each point in the likelihood space giving a specific prediction for the observables, N_{ch} and $\cos\theta_{\rm UL}$. In most cases, the nominal value of a nuisance parameter corresponds to the predictions of the Barr *et al.* flux, with best-known inputs to the detector simulation chain. We describe each of the individual errors shown in table 8.1 in the following sections.

8.2 Sources of Systematic Error

8.2.1 Atmospheric Neutrino Flux Normalization

One of the largest sources of systematic error is the overall normalization of the atmospheric neutrino flux. While the total $\nu_\mu + \bar{\nu}_\mu$ simulated event rate for recent models [16, 17] only differs by $\pm 7\%$, this masks significantly larger differences in the individual ν_μ and $\bar{\nu}_\mu$ rates. We take the latter

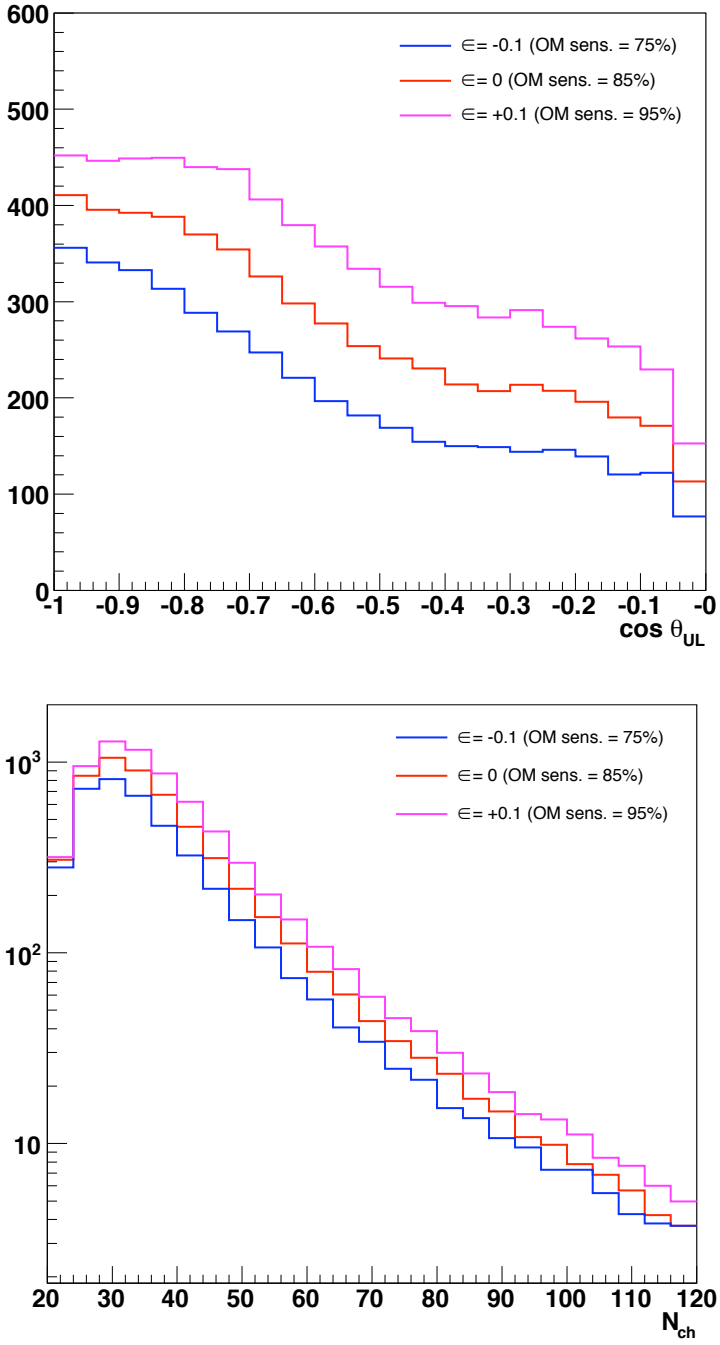

Figure 8.3: Simulated effect of OM sensitivity uncertainty ϵ on $\cos\theta_{\rm UL}$ and N_{ch} distributions.

Table 8.1: Systematic errors in the atmospheric muon neutrino flux, separated by effect on the observables $\cos\theta_{\rm UL}$ and N_{ch}. See the text for more detail on each source of error.

Error	Class	Magnitude
Atm. $\nu_\mu + \bar{\nu}_\mu$ flux	α_1	$\pm 18\%$
Neutrino interaction	α_2	$\pm 8\%$
Reconstruction bias	α_2	-4%
ν_τ-induced muons	α_2	$+2\%$
Background contamination	α_2	$+1\%$
Charmed meson contribution	α_2	$+1\%$
Timing residual uncertainty	α_2	$\pm 2\%$
Muon energy loss	α_2	$\pm 1\%$
Primary CR slope (H, He)	$\Delta\gamma$	± 0.03
Charmed meson contribution	$\Delta\gamma$	$+0.05$
Pion/kaon ratio	κ	$+0.01/-0.03$
Charmed meson contribution	κ	-0.03
OM sensitivity, ice	ϵ	$\pm 10\%$

difference of $\pm 18\%$ to be more representative of the true uncertainties in the models. This is also in line with the total uncertainty in the flux estimated in ref. [17].

8.2.2 Neutrino Interaction

To estimate the error in our simulation of neutrino interactions (from the cross section, scattering angle, parton distribution functions, etc.), we compare our NUSIM atmospheric neutrino simulation with a sample generated with ANIS [96]. ANIS uses the more modern CTEQ5 neutrino-nucleon cross sections and parton distribution functions [97], compared to MRS [98] in NUSIM, and it also accurately simulates the neutrino-muon scattering angle. We find an 8% difference in the normalization for an atmospheric neutrino spectrum (ANIS produces fewer events). There is a small difference in the median energy of the events (3% lower in ANIS), but we do not find an appreciable difference in the shapes of the observable distributions.

8.2.3 Reconstruction Bias

We characterize our uncertainty in our reconstruction quality parameters ("reconstruction bias") by investigating how systematic disagreements between data and simulation affect the number of events surviving to the final cut level. In particular, the smoothness and paraboloid error variables have systematically lower values in simulation than in data. To quantify how this affects our event

sample, we first use a Kolmogorov test to find the scaling factor for these variables that reduces the disagreement the most. Specifically, the scaling factor ζ approximately corrects the simulated value of the quality variable q_{MC} to the observed value q_{data}:

$$q_{\mathrm{data}} \approx \zeta \; q_{\mathrm{MC}} \; . \tag{8.1}$$

Note that this is not a normalization change, but a scaling of the individual observable values. These scaling factors are shown in table 8.2.

Table 8.2: Systematic scaling factors for smoothness and paraboloid error.

Parameter	**Scaling Factor**
$\lvert S_{\mathrm{Phit}} \rvert$	1.09
P_{err}	1.025

We then reapply the cuts to the simulation using the scaled values of these two parameters. This decreases the final number of events in the sample by 4%, and we use this as the estimated uncertainty due to systematic shifts in our quality parameters. We note this is smaller than the estimate of 9% in ref. [82], because our agreement between data and simulation in the paraboloid error is better (possibly because we use PHOTONICS to simulate the photon propagation).

8.2.4 Tau-neutrino-induced Muons

Normally, ν_τ-induced muons are negligible when considering atmospheric neutrinos, since mass-induced oscillations are such a small effect at our energies. However, in the new physics scenarios we consider here, up to 50% of our ν_μ flux can oscillate to ν_τ, which can then interact, generate a τ lepton, and then decay to a muon. To estimate this flux, we generate a sample of tau neutrinos with ANIS. Then, we weight each event with a ν_μ atmospheric weight times its oscillation probability, for the extreme case described above. The final event rate is only 2% of the total.

The reason the flux is so small is that we are significantly penalized by the steep power-law energy spectrum, considering that the muon takes only a fraction of the tau energy. We are also penalized by the branching ratio of $\tau \to \mu\nu_\mu\nu_\tau$, which is only 17% [20].

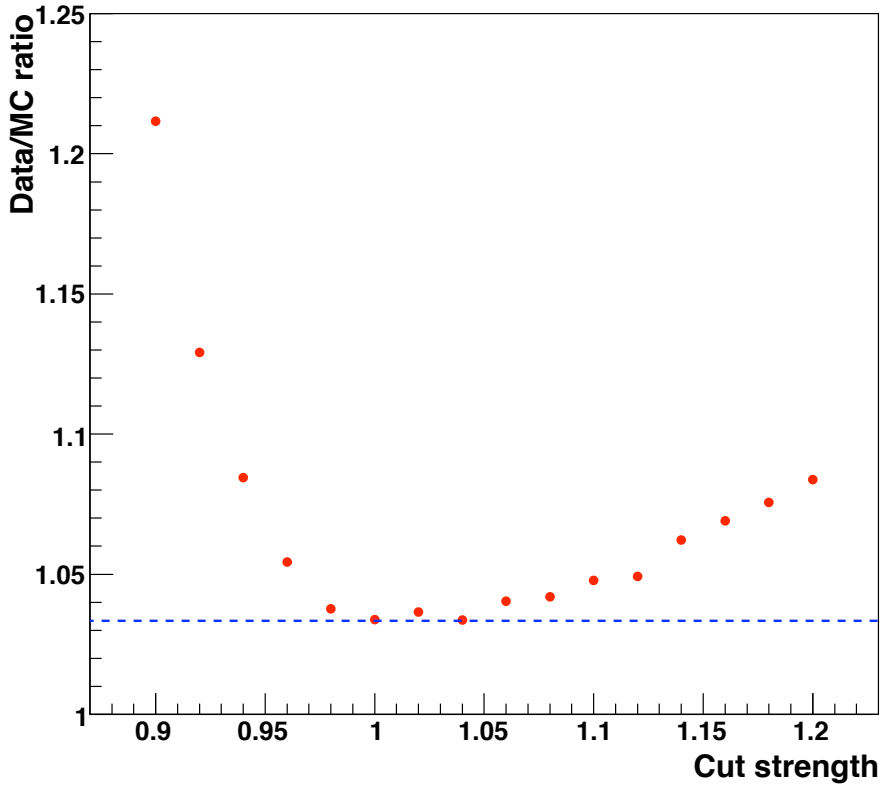

Figure 8.4: Ratio of data to atmospheric neutrino MC as a function of cut strength (see eq. 6.4 for definition), after application of L2 purity cuts. The dashed line represents the minimum of the "flat" region. There is no significant difference between the minimum and the ratio at the analysis final cut level (cut strength = 1.0).

8.2.5 Background Contamination

As described in section 6.3.2, we estimate the background contamination by tightening the quality cuts until the ratio of data to simulation flattens out. After application of the L2 purity cuts, we recheck the background contamination with this method and still estimate it to be less than 1% (see fig. 8.4).

Although frequently used, this method is far from ideal. In addition to the drawbacks mentioned in chapter 6, the data/MC ratio does not reliably have the shape shown in fig. 8.4, with a clear region of constant ratio. In fig. 8.4, only the initial point-source cuts were scaled as a function of cut strength, and the purity cuts were not. Scaling the purity cuts as well presents several problems (e.g., how to "scale" a two-dimensional cut), and frequently the data/MC ratio does not flatten.

A better approach would be to model the background contamination in the observables and then include the contamination level as a nuisance parameter, but this requires a more sophisticated understanding of the characteristics of mis-reconstructed muons that survive to high quality levels.

8.2.6 Timing Residual Uncertainty

The uncertainty in the timing of the optical modules is less than 5 ns, as determined by YAG laser pulses. We take the effect on the normalization of the final event sample to be 2%, from ref. [82].

8.2.7 Muon Energy Loss

The uncertainty in the muon energy loss from various sources is rather small at TeV energies, and we use the estimate in ref. [82] of a 1% effect in the absolute event rate.

8.2.8 Primary Cosmic Ray Slope

Some uncertainty remains in the primary cosmic ray (CR) spectral index [99]. If this is small, we can model this by just changing the spectral slope by some amount $\Delta\gamma$. The uncertainty in the slope of the proton component is small (0.01), but the uncertainty in the Helium component is much larger (0.07). To find how much this uncertainty changes the total flux, we approximate a change in spectral index of a secondary component as follows:

$$E^{-\gamma} + fE^{-\gamma+\Delta\gamma} \approx (1+f)E^{-\gamma}E^{\frac{f}{1+f}\Delta\gamma} = (1+f)E^{-\gamma+\frac{f}{1+f}\Delta\gamma} \tag{8.2}$$

That is, we note that to first order, a change in spectral index of $\Delta\gamma$ in a secondary component is scaled by approximately $f/(1+f)$, the fraction of the total flux for that component. Since f_{He} is at most 30% in our energy range, we set the uncertainty in the primary cosmic ray spectral index to $\Delta\gamma = 0.01 + 0.3 \cdot 0.07 = 0.03$.

8.2.9 Charmed Meson Contribution

The Barr *et al.* and Honda *et al.* atmospheric neutrino flux predictions only include ν_μ from charged pion and kaon decay, but at high enough energies, a charmed meson (e.g. $D^\pm$) can be produced. These decay almost immediately[1] and can produce a high-energy contribution to the

[1]This is why the charmed meson component of the flux is also referred to as "prompt."

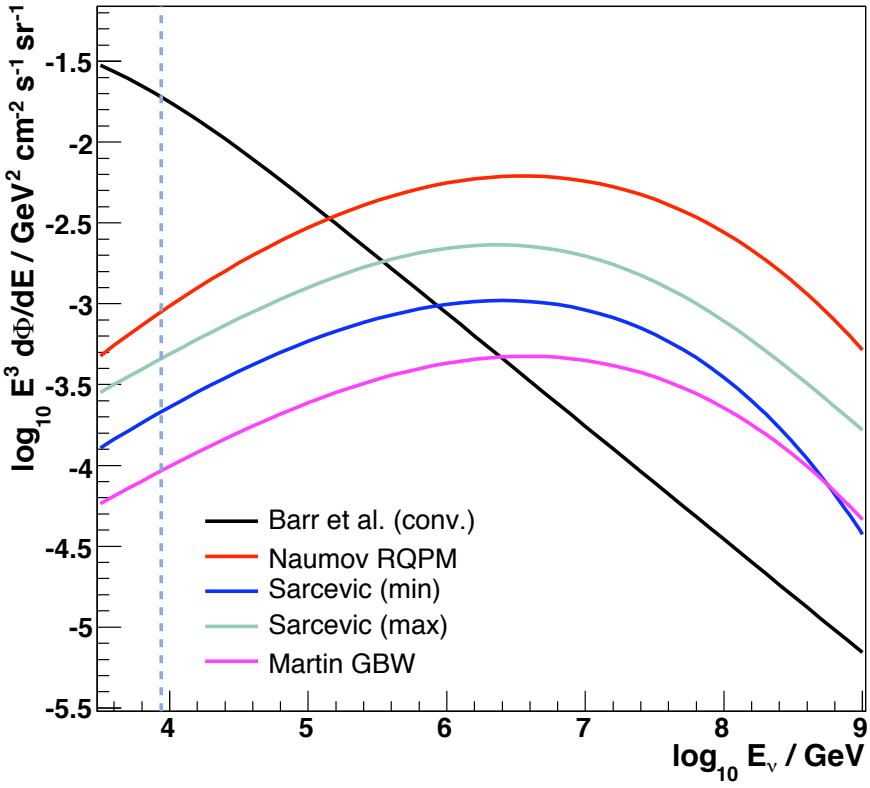

Figure 8.5: Predicted atmospheric prompt neutrino fluxes, averaged over zenith angle and multiplied by E^3 to enhance features, compared to the Barr *et al.* conventional flux. The vertical dotted line marks the 95% point of the energy range of our data sample.

atmospheric neutrino flux. The cross sections that are relevant for charm production, however, are still uncertain, leading to huge uncertainties in the normalization of this flux (see [100] for a review).

In our simulation, we neglect any charm component. To estimate the systematic uncertainty by neglecting this component, we compare the predicted flux when adding the Naumov RQPM ("Recombination Quark Parton Model") flux [101]. The RQPM model is a non-perturbative approach, and is a conservative choice as the predicted flux is quite large. More recent perturbative-QCD calculations predict maximum fluxes that are quite a bit smaller (see e.g. [102, 103]). Even in the Naumov RQPM model, however, the predicted flux is almost negligible in the energy range of this analysis, as shown in fig. 8.5.

We find the difference in normalization to be 1% at our final cut level. We also incorporate the effect on the N_{ch} distribution by modeling the charm contribution as a change in slope of the energy

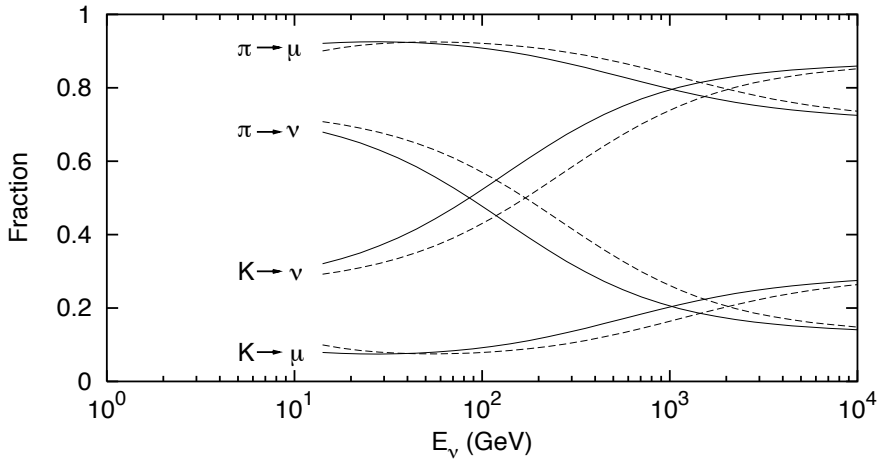

Figure 8.6: Relative contribution of pions and kaons to atmospheric muons and muon neutrinos, as a function of energy (solid = vertical, dashed = 60°, from [104]).

spectrum. We find that changing the spectral index by +0.05 matches the increase at high N_{ch} caused by the Naumov flux. Changing the spectral index actually distorts the zenith angle spectrum by a tiny amount (a tilt of $\kappa = -0.03$), but we saw no observable difference when just adding the Naumov flux, so to be conservative we "correct" for this by adding in a tilt uncertainty along with the slope term described above.

8.2.10 Rock Density

The uncertainty in the density of bedrock under the polar ice is approximately 10% [82]. To model the effect of this, we modify both ANIS and MMC, increasing the density of rock by 10% in both, and compare to an unmodified ANIS sample. We find a negligible difference in atmospheric event rates of $< 0.1\%$. We note that increases in interaction probability due to increased density are offset by decreased muon range.

8.2.11 Pion/Kaon Ratio

The relative contribution to the atmospheric neutrino flux from pions and kaons is both energy and zenith-angle dependent. This general dependence is shown in fig. 8.6. However, the cross section for the production of kaons is still relatively uncertain, and thus the exact value of the π/K ratio contributes to our systematic errors.

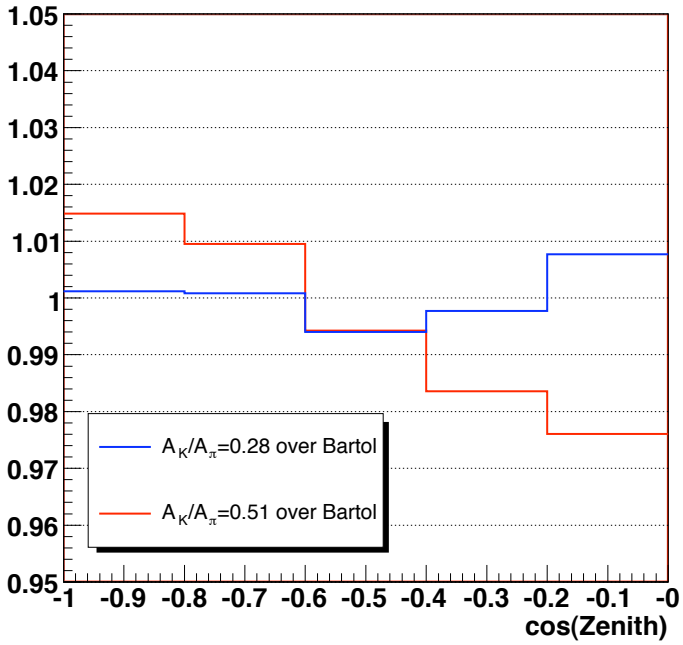

Figure 8.7: Ratio of $\cos\theta_{\rm UL}$ distributions for extreme values of the π/K ratio to the standard Barr *et al.* prediction.

To determine the effect on the zenith angle distribution from the uncertainty in the pion/kaon ratio, we first implement an atmospheric flux weighting scheme that directly uses the Gaisser parametrization (eq. 2.5), with coefficients fitted to reproduce the Barr *et al.* flux. We then compare the $\cos\theta_{\rm UL}$ distributions using extreme values for A_K/A_π of 0.28 and 0.51, as derived from the $Z_{N\pi}$ and Z_{NK} uncertainty tabulated in ref. [105]. The effect is small and can be approximated by a linear tilt κ in the $\cos\theta_{\rm UL}$ distribution of $+0.01/-0.03$ (see fig. 8.7).

8.2.12 Optical Module Sensitivity and Ice

A significant source of error is the uncertainty in the absolute sensitivity of the optical modules. This has a large effect on both the overall detector event rate (a decrease of 1% in sensitivity results in a decrease of 2.5% in event rate) and the shape of the zenith angle and N_{ch} distributions, as shown in fig. 8.3. We quantify this uncertainty by comparing the trigger rate of down-going muons in 2005 with simulation predictions given various OM sensitivities, including the uncertainty of hadronic interactions, by using CORSIKA air shower simulations with the SIBYLL 2.1 [106], EPOS 1.60 [107], and QGSJET-II-03 [108] interaction models. For the purposes of this study, we assume that the

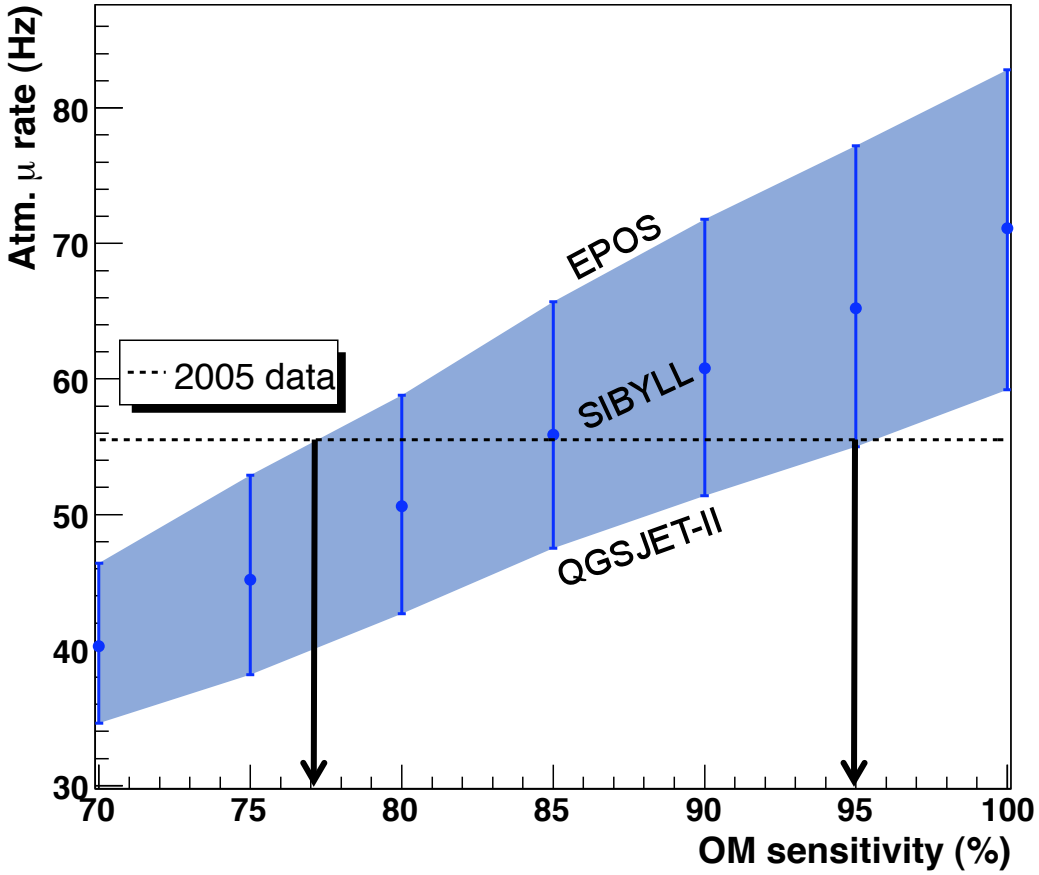

Figure 8.8: AMANDA-II atmospheric muon rate vs. OM sensitivity for 2005 data and various simulated hadronic interaction models. The arrows indicate the range of OM sensitivities compatible with data using this estimation procedure.

uncertainty in the hadronic interaction model is the primary uncertainty in the atmospheric muon rate. There is also uncertainty in the normalization of the cosmic ray flux itself, but this is included in the atmospheric neutrino flux uncertainty, and including it again here would be "double-counting."

We find that we can constrain the optical module sensitivity to within $+10\%/-7\%$, around a central value of 85% (see fig. 8.8). The range of uncertainty in the sensitivity is consistent with ref. [82]. That analysis used the shape of the atmospheric neutrino zenith angle distribution to constrain the OM sensitivity to 100%+3%-10%. However, we cannot use this approach since this distribution was blinded to us.

Table 8.3: Relative simulated atmospheric neutrino rates for various ice models, compared to the rate with PTD and MAM.

Ice Model	**Relative atm. ν_μ rate**
Millennium	+39%
AHA	+23%
AHA (85% OMs)	-8%

The difference in central values of the OM sensitivity is due to the differences in ice model used; we use the PHOTONICS package with the AHA ice model, while the simulation in ref. [82] used PTD with the MAM ice model. While the MAM ice model was tuned to muon data, and reproduces both muon and neutrino rates fairly well, all ice models released with PHOTONICS produce muon and neutrino rates that are significantly higher. It is currently unclear whether this is a problem with the ice model or with the PHOTONICS software, but we can "fix" this discrepancy by changing the nominal OM sensitivity to 85% when using the AHA ice model. We note that this also brings the expected neutrino rate in line with the MAM ice model, as shown in table 8.3.

This entanglement between ice and optical module sensitivity is rather insidious, because to first order, changing the ice model and OM sensitivity has a similar effect on the N_{ch} and $\cos\theta_{\mathrm{UL}}$ distributions. Using timing information (such as timing residuals), it is possible to disentangle the two. To our observables, however, the effects are similar, so we model both as a single source of error. Changing the OM sensitivity by $\pm 10\%$ covers the range in uncertainty in ice models and is also in line with our muon rate analysis (see fig. 8.9), so we use this as the final error estimate for this class of uncertainty.

8.3 Final Analysis Parameters

We make a few more simplifications to reduce the dimensionality of the likelihood space. First, we note that the phase η in the VLI survival probability (eq. 3.7) is only relevant if the VLI effects are large enough to overlap in energy with conventional oscillations (i.e., below 100 GeV). Since our neutrino sample is largely outside this range, we set $\cos\eta = 0$ for this search. This means we can also limit the VLI mixing angle to the range $0 \leq \sin 2\xi \leq 1$. Second, in the QD case, we vary the decoherence parameters D_i^* in pairs (D_3^*, D_8^*) and (D_6^*, D_7^*). If we set D_3^* and D_8^* to zero, $\frac{1}{2}$ of ν_μ

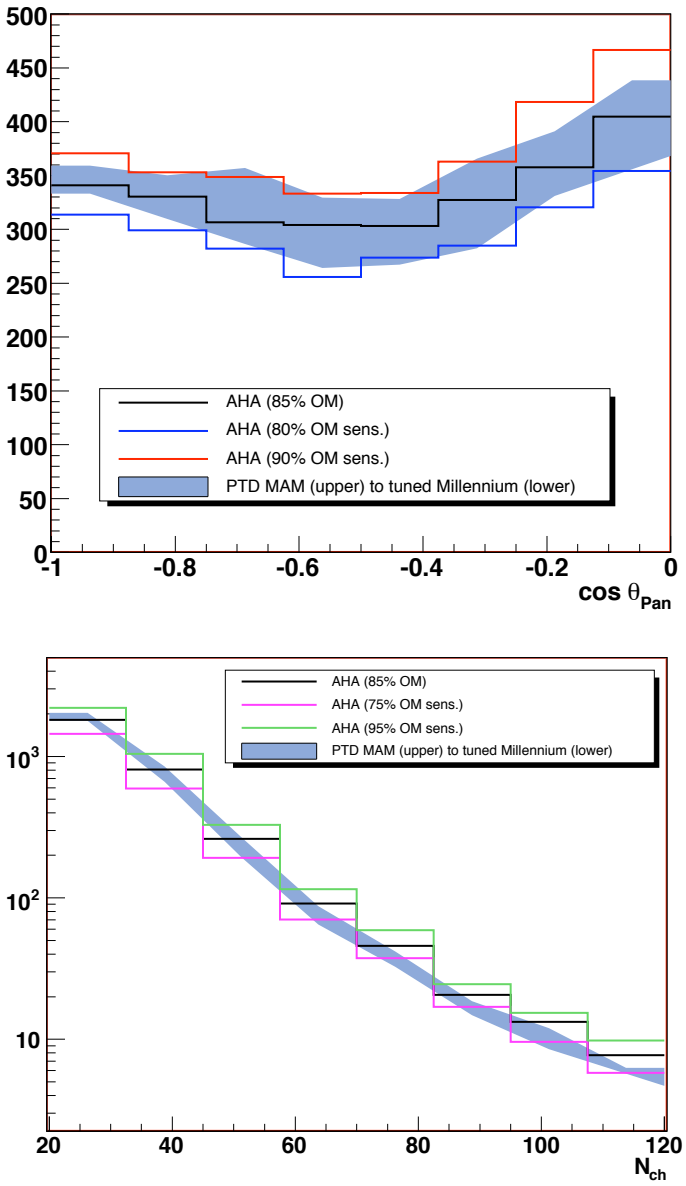

Figure 8.9: Simulated effect of OM sensitivity on $\cos\theta_{\rm UL}$ and N_{ch} distributions, compared to the spread from different ice models (2005 L3 simulation).

Table 8.4: Physics parameters and nuisance parameters used in each of the likelihood analyses (VLI, QD, and conventional).

Analysis	Physics parameters	Nuisance parameters
VLI	$\Delta\delta,\ \sin 2\xi$	$\alpha_1,\ \alpha_2,\ \Delta\gamma,\ \kappa,\ \epsilon$
QD	$D^*_{3,8},\ D^*_{6,7}$	$\alpha_1,\ \alpha_2,\ \Delta\gamma,\ \kappa,\ \epsilon$
Conv.	$\alpha_1,\ \Delta\gamma$	$\alpha_2,\ \kappa,\ \epsilon$

remain after decoherence; with D^*_6 and D^*_7 set to zero, $\frac{5}{6}$ remain; and with all D^*_i equal and nonzero, $\frac{1}{3}$ remain after decoherence.

Finally, in the absence of new physics, we can use the same methodology to determine the conventional atmospheric neutrino flux. In this case, the nuisance parameters α_1 (the uncertainty in the atmospheric neutrino flux normalization) and $\Delta\gamma$ (the change in spectral slope relative to the input model) become our physics parameters. The determination of an input energy spectrum by using a set of model curves with a limited number of parameters is commonly known as *forward-folding* (see e.g. ref. [109]).

Table 8.4 summarizes the likelihood parameters used for the VLI, QD, and conventional analyses.

Chapter 9

Results

9.1 Final Zenith Angle and N_{ch} Distributions

After performing the likelihood analysis described in chapter 7 on the $(\cos\theta_{\rm UL}, N_{ch})$ distribution, we find no evidence for VLI-induced oscillations or quantum decoherence, and the data are consistent with expectations from atmospheric flux models. The reconstructed zenith angle and N_{ch} distributions compared to standard atmospheric neutrino models are shown in fig. 9.1, projected into one dimension and rebinned.

9.2 Likelihood Ratio and Best-fit Point

The profile likelihood ratio of the data to the best-fit point over the parameter space, along with the critical value $\Delta\mathcal{L}_{\rm crit}$ at a confidence interval, are the two fundamental results of the profile construction method described in chapter 7. As an example, we show in fig. 9.2(a) the likelihood ratio as a function of the $n = 1$ VLI parameters $\log_{10}\Delta\delta$ and $\sin^2 2\xi$ (we have switched the mixing angle parameter from $\sin 2\xi$ to $\sin^2 2\xi$ to enhance the region of interest in the parameter space). Regions with high values of the likelihood ratio are easily excluded. The boundary between the allowed and excluded regions at a certain confidence level is given by the intersection of this surface with the critical surface for that confidence level $\Delta\mathcal{L}_{\rm crit}$, shown in fig. 9.2(b). We can see that the critical value varies quite substantially across the parameter space from 4.61, the 90% CL χ^2 value for two degrees of freedom. Instead of finding the intersection of these two surfaces, in practice it is easier (and likely more accurate) to compute the actual CL at a given grid point and then interpolate to

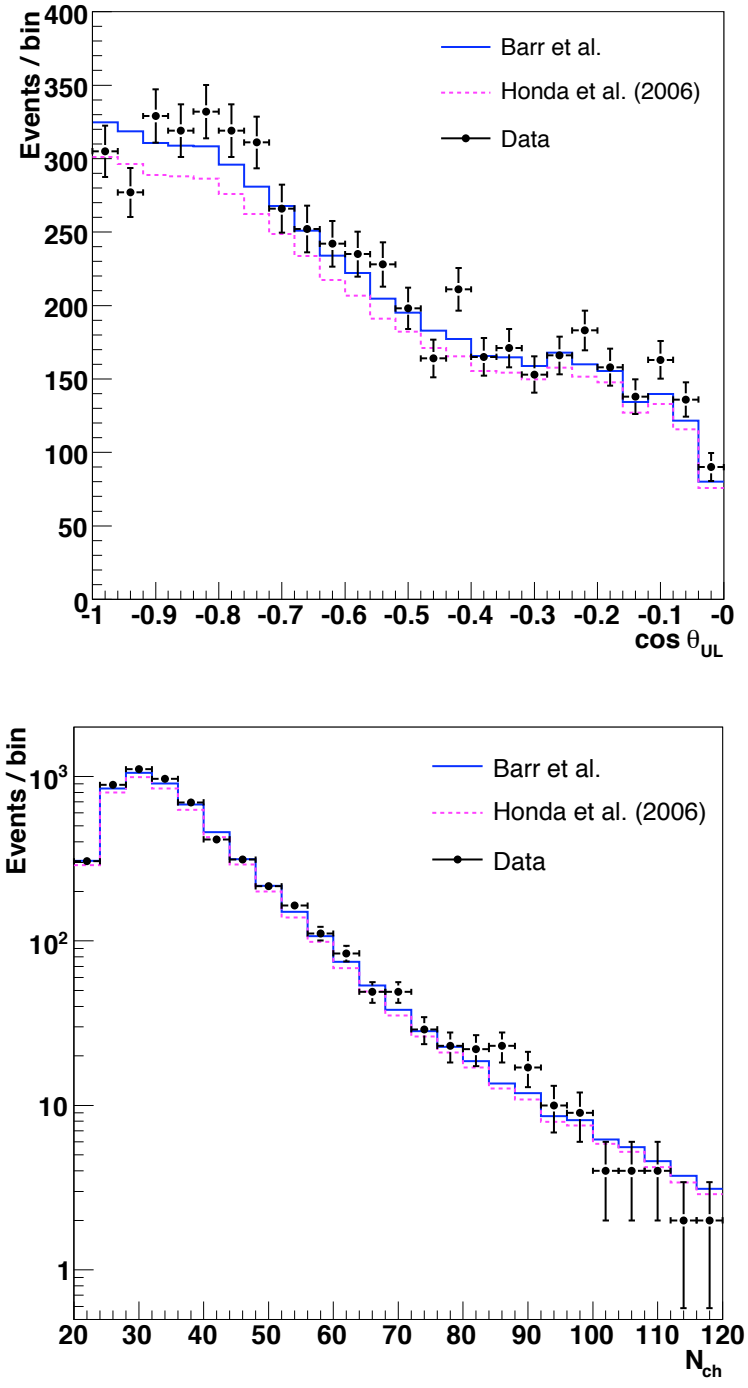

Figure 9.1: Zenith angle and N_{ch} distributions of the 5511 candidate atmospheric neutrino events in the final sample, compared with Barr *et al.* [16] and Honda *et al.* [17] predictions (statistical error bars).

find the desired contours.

The best-fit point shown in fig. 9.2(a) has physics parameters ($\log_{10}\Delta\delta = -25.5,\ \sin^2 2\xi = 0.4$) and nuisance parameters ($1+\alpha = 1.14,\ \Delta\gamma = 0.08,\ \kappa = -0.03,\ \epsilon = 0.85$). The minimum occurs at a rather large value of the VLI parameter, but with a small mixing angle, and by construction the nuisance parameters have been adjusted from their central values to make the hypothesis fit the data. A high normalization and harder spectrum compensate for the loss of events due to VLI. The fact that the best-fit point does not correspond to Standard Model physics (the lower left-hand corner of fig. 9.2(a)) is likely due to statistical fluctuations or small remaining disagreements between data and simulation. As we will see in the next section, the difference is not statistically significant.

9.3 Upper Limits on Quantum Gravity Parameters

In all new physics scenarios, the data are consistent with no new physics at the 90% CL. We show here the allowed regions in the various parameter spaces, and we set upper limits on the quantum gravity parameters.

9.3.1 Violation of Lorentz Invariance

The 90% CL upper limits on the VLI parameter $\Delta\delta$ for oscillations of various energy dependencies, with maximal mixing ($\sin 2\xi = 1$) and phase $\cos\eta = 0$, are presented in table 9.1. Allowed regions at 90%, 95%, and 99% confidence levels in the $\log_{10}\Delta\delta$-$\sin^2 2\xi$ plane for the $n = 1$ hypothesis are shown in fig. 9.3. The upper limit at maximal mixing of $\Delta\delta \leq 2.8 \times 10^{-27}$ is competitive with that from a combined Super-Kamiokande and K2K analysis [43].

In the $n = 1$ case, recall that the VLI parameter $\Delta\delta$ corresponds to the splitting in velocity eigenstates $\Delta c/c$. Observations of ultra-high energy cosmic rays constrain VLI velocity splitting in other particle sectors, with the upper limit on proton-photon velocity splitting of $(c_p - c)/c < 10^{-23}$ [41]. While we probe a rather specific manifestation of VLI in the neutrino sector, our limits are orders of magnitude better than those obtained with other tests.

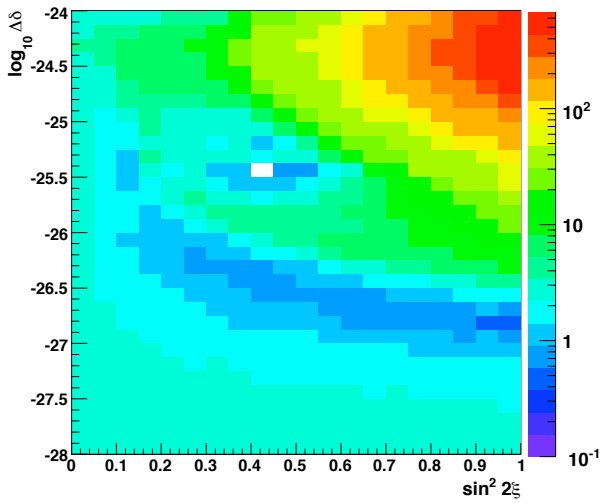

(a) Profile likelihood ratio $\Delta\mathcal{L}$ for the 2000-06 data sample. The white cell indicates the best-fit grid point, with $\Delta\mathcal{L} = 0$.

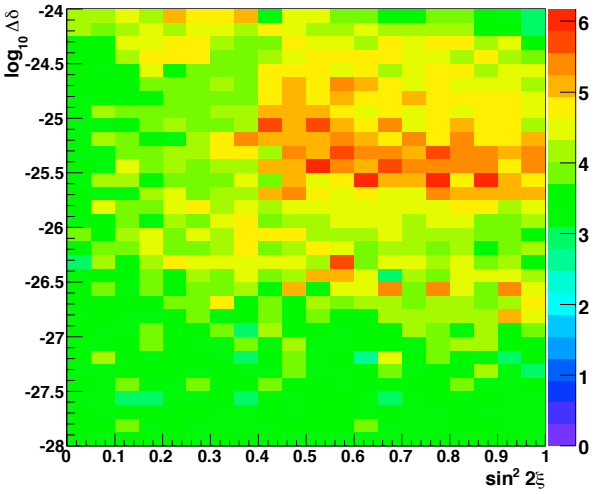

(b) Critical value of the likelihood ratio $\Delta\mathcal{L}_{\text{crit}}$ at the 90% CL as determined by MC experiments. Values in the cells far away from intersecting the surface in fig. 9.2(a) are approximate, as MC experiments are terminated early.

Figure 9.2: Profile likelihood ratios $\Delta\mathcal{L}$ and $\Delta\mathcal{L}_{\text{crit}}$ over the $n = 1$ VLI parameter space. Note that the vertical scales differ substantially.

Table 9.1: 90% CL upper limits from this analysis on VLI and QD effects proportional to E^n. VLI upper limits are for the case of maximal mixing ($\sin 2\xi = 1$), and QD upper limits are for the case of $D_3^* = D_8^* = D_6^* = D_7^*$.

n	VLI ($\Delta\delta$)	QD (D^*)	Units
1	2.8×10^{-27}	1.2×10^{-27}	–
2	2.7×10^{-31}	1.3×10^{-31}	GeV^{-1}
3	1.9×10^{-35}	6.3×10^{-36}	GeV^{-2}

9.3.2 Quantum Decoherence

The 90% CL upper limits on the decoherence parameters D_i^* given various energy dependencies are also shown in table 9.1. Allowed regions at 90%, 95%, and 99% confidence levels in the $\log_{10} D_{3,8}^*$-$\log_{10} D_{6,7}^*$ plane for the $n = 2$ case are shown in fig. 9.4. The 90% CL upper limit from this analysis with all D_i^* equal for the $n = 2$ case, $D^* \leq 1.3 \times 10^{-31}$ GeV^{-1}, extends the previous best limit from Super-Kamiokande by nearly four orders of magnitude. Because of the strong E^2 energy dependence, AMANDA-II's extended energy reach allows for much improved limits.

9.4 Determination of the Atmospheric Neutrino Flux

9.4.1 Result Spectrum

In the absence of evidence for violation of Lorentz invariance or quantum decoherence, we interpret the atmospheric neutrino flux in the context of Standard Model physics only. We use the likelihood analysis to perform a two-parameter forward-folding of the atmospheric neutrino flux to determine the normalization and any change in spectral index relative to existing models. As described in section 8.3, we test hypotheses of the form

$$\Phi(E, \theta, \phi) = (1 + \alpha_1)\ \Phi_{\mathrm{ref}}(E, \theta, \phi) \left(\frac{E}{E_{\mathrm{median}}} \right)^{\Delta\gamma} , \tag{9.1}$$

where $\Phi_{\mathrm{ref}}(E, \theta, \phi)$ is the Barr *et al.* or Honda *et al.* flux.

The allowed regions in the $(1 + \alpha_1)$-$\Delta\gamma$ parameter space are shown in fig. 9.5. We translate this result into a range of fluxes by forming the envelope of the set of curves allowed on the 90%

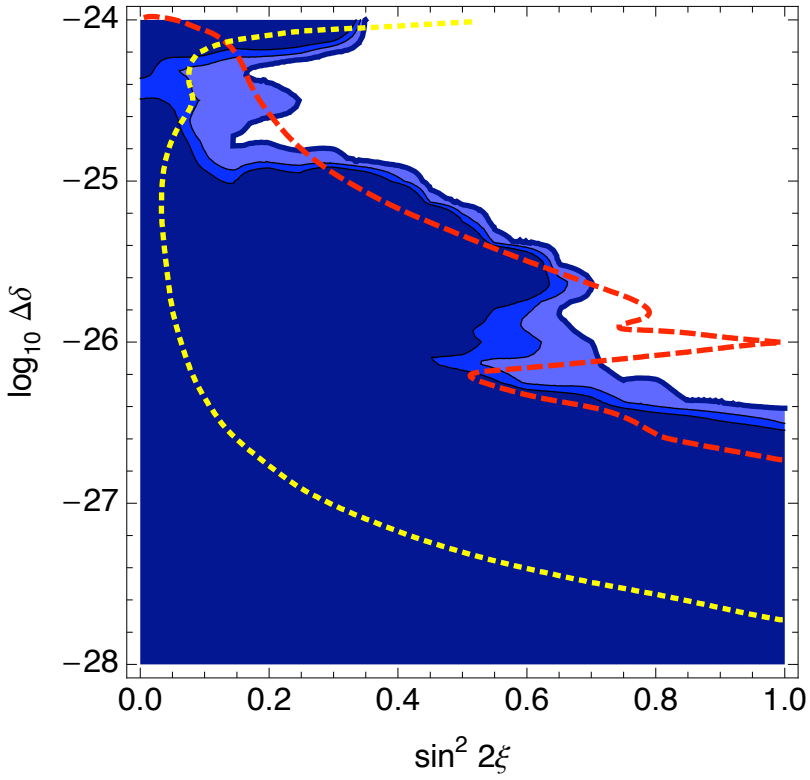

Figure 9.3: 90%, 95%, and 99% CL allowed regions (from darkest to lightest) for VLI-induced oscillation effects with $n = 1$. Also shown are the Super-Kamiokande + K2K 90% contour [43] (dashed line), and the projected IceCube 10-year 90% sensitivity [110] (dotted line).

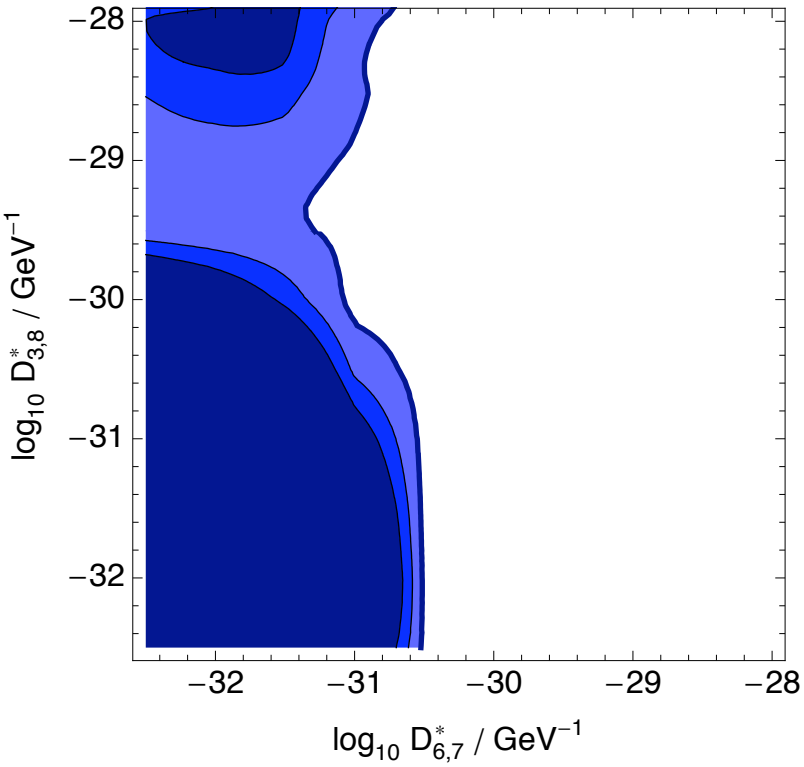

Figure 9.4: 90%, 95%, and 99% CL allowed regions (from darkest to lightest) for QD effects with $n = 2$.

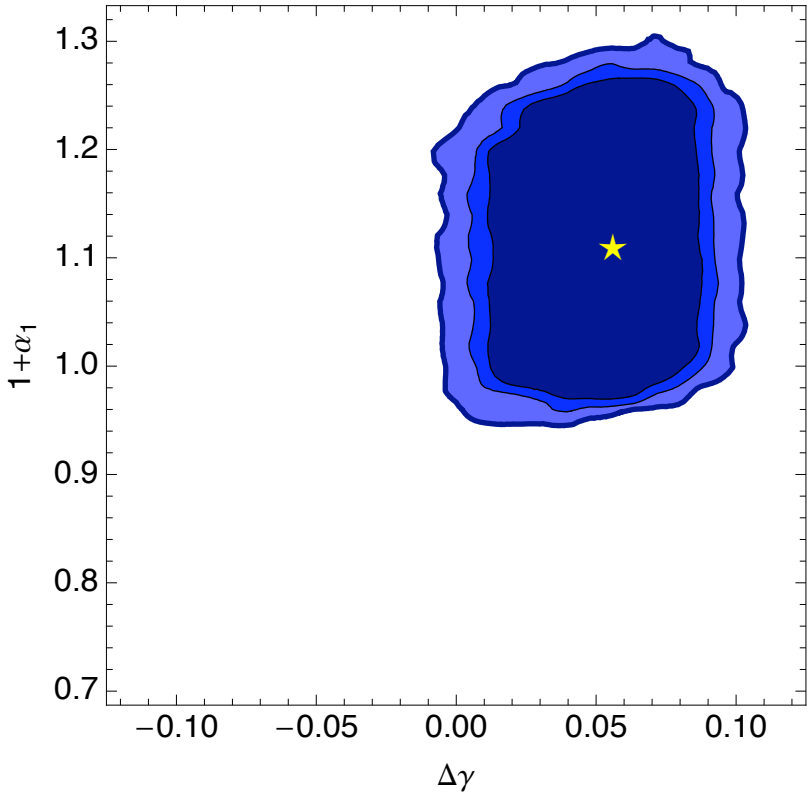

Figure 9.5: 90%, 95%, and 99% allowed regions (from darkest to lightest) for the normalization $(1 + \alpha_1)$ and change in spectral index $(\Delta\gamma)$ of the conventional atmospheric neutrino flux, relative to Barr *et al.* [16]. The star marks the central best-fit point.

contour line in fig. 9.5 (see fig. 9.6 for an illustration)[1]. This band of allowed energy spectra is shown in fig. 9.7 and compared to results obtained with Super-Kamiokande data [111].

The central best-fit point is also shown in figs. 9.5 and 9.7. However, because of the degeneracy between the normalization parameter α_1 and the systematic error α_2, the best-fit point actually spans a range of normalizations corresponding to the uncertainty α_2. Specifically, we find the best-fit spectra to be

$$\Phi_{\text{best-fit}} = (1.1 \pm 0.1) \left(\frac{E}{640 \text{ GeV}} \right)^{0.056} \cdot \Phi_{\text{Barr}} \tag{9.2}$$

for the energy range 120 GeV to 7.8 TeV (for a discussion of this range, see section 9.4.2). Note that

[1]Technically, the band should be constructed from the envelope of curves from the entire 90% allowed region, but in this case one can easily show that the boundary suffices.

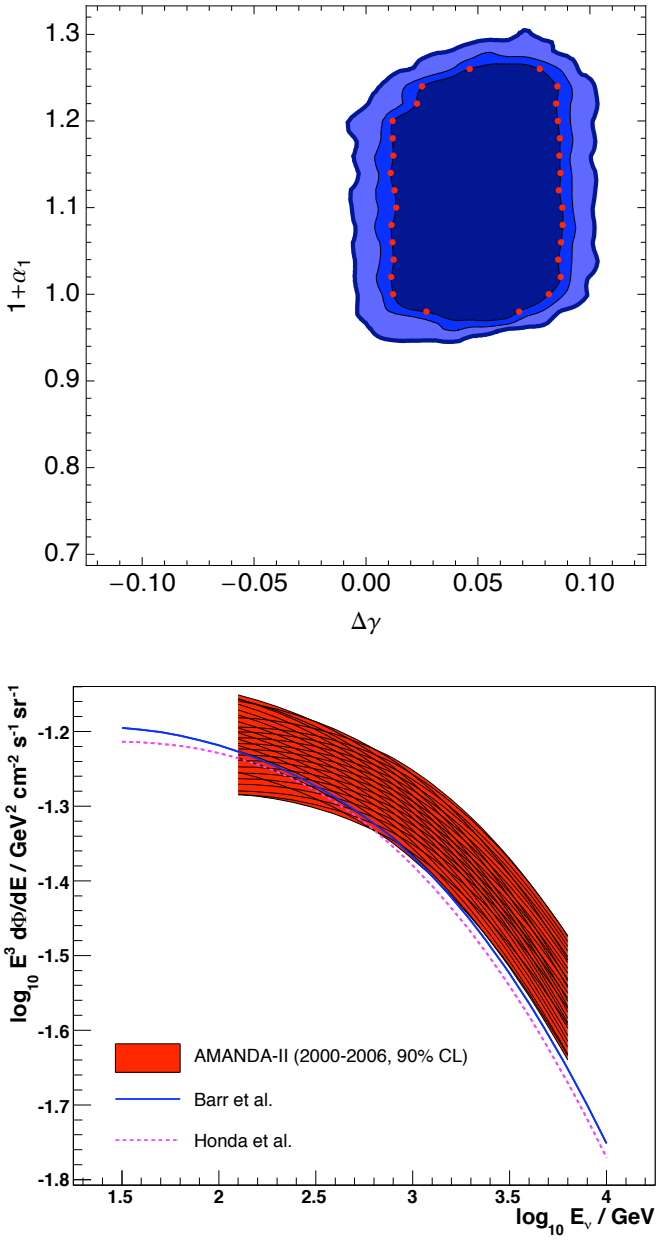

Figure 9.6: The allowed range of the atmospheric neutrino flux is determined by extracting the 90% CL contour from the parameter space (top; red points) and forming the envelope of this set of spectra (bottom; black curves and red band).

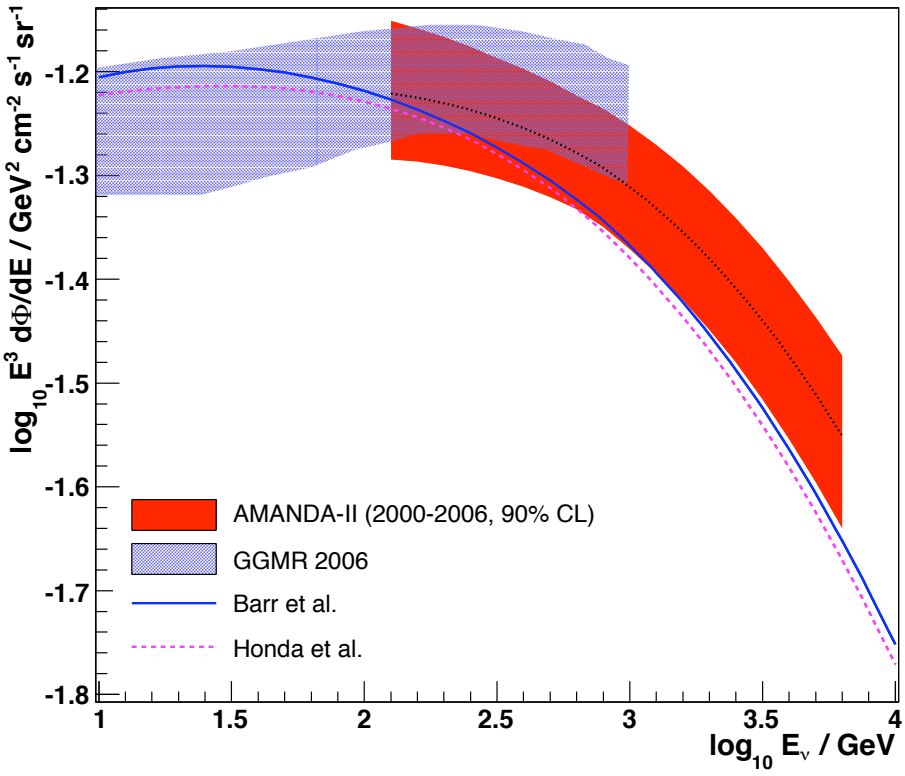

Figure 9.7: Angle-averaged $\nu_\mu + \bar{\nu}_\mu$ atmospheric neutrino flux (solid band, 90% CL from the forward-folding analysis), multiplied by E^3 to enhance features. The dotted line shows the central best-fit curve. Also shown is a previous result by González-García *et al.* using Super-Kamiokande data [111], as well as Barr *et al.* [16] and Honda *et al.* [17] predictions. All fluxes are shown without oscillations.

$\Phi_{\text{best-fit}}$ does not represent the entire allowed band at any CL, but only the set of best-fit spectra with $\Delta\mathcal{L} = 0$. There is no degeneracy on the $\Delta\gamma$ parameter, so there is only a single best-fit value for the change in spectral slope. Not including the normalization, the best-fit nuisance parameters for the minima are ($\kappa = 0.03, \epsilon = 0.82$). We note that the best-fit OM sensitivity of 82% is close to the nominal value of 85% found in our systematic error study of section 8.2.12.

9.4.2 Valid Energy Range of Result

We define the valid energy range of the resulting flux band as the intersection of the 5%-95% regions of the allowed set of spectra, as determined by the simulated neutrino energy distributions at the final cut level. We also marginalize over the OM sensitivity, which affects the energy distribution. This results in an energy range of 120 GeV to 7.8 TeV for this result.

We note this procedure is different (and more conservative) than that used to define the energy range covered by an unfolding analysis, which is often determined by the reconstructed energy of the highest-energy event. Since we bypass any energy reconstruction (it is unnecessary in an observable-based likelihood analysis), we have no such recourse. As an alternative, using the median energy of simulated events in the highest N_{ch} bin (110-120) would result in a similar energy range as the above method, with a slightly higher cutoff of 9.5 TeV. Using the 95% point of the simulated energy distribution of events in this highest bin would extend the energy range to 76 TeV, but there are not necessarily any data events at this energy.

9.4.3 Dependence on Flux Model

Because there are no significant differences in the shape of the atmospheric neutrino spectra of the Barr *et al.* and Honda *et al.* models, the results of the conventional analysis (as expressed as a primary flux) should be independent of the input model. To check this, we compare the flux obtained using the Honda *et al.* model as the primary model in place of the Barr *et al.* model which was used above.

The allowed regions in normalization and spectral slope parameter space are shown in fig. 9.8. The regions are similar, with a global offset in the normalization due to the different model normalizations. The best-fit slope point is identical. The resulting flux band calculated from the 90%

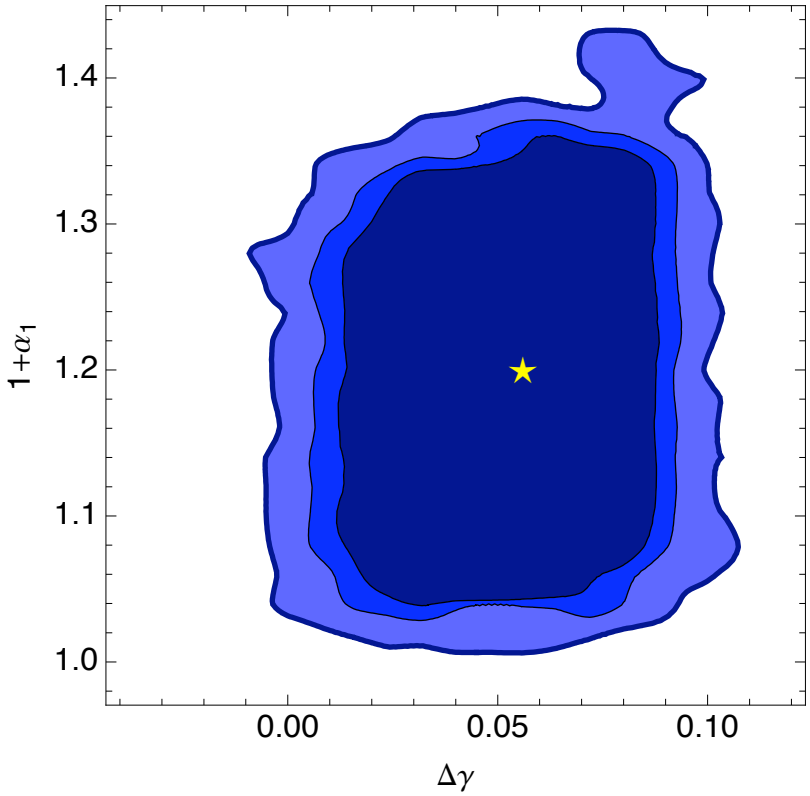

Figure 9.8: 90%, 95%, and 99% allowed regions (from darkest to lightest) for the normalization $(1 + \alpha_1)$ and change in spectral index $(\Delta\gamma)$ of the conventional atmospheric neutrino flux, relative to Honda *et al.* [17]. The star marks the central best-fit point.

allowed region is shown in fig. 9.9, compared to the band obtained with Barr *et al.* The maximum difference in the flux boundary is 6%. Compared to the vertical width of the band of 30% to 50%, the dependence on flux model is a subdominant effect.

9.5 Comparison with Other Results

The atmospheric neutrino spectrum determined with this analysis is compatible with an analysis of Super-Kamiokande data [111], and extends that measurement by nearly an order of magnitude in energy. Our data suggest an atmospheric neutrino spectrum with a slightly harder spectral slope and higher normalization that either the Barr *et al.* or Honda *et al.* model. We also compare our results to an unfolding of the Fréjus data [112] and to an unfolding of four years of AMANDA-II data [113] in fig. 9.10. Except for the Fréjus results, the fluxes are shown without any oscillation effects.

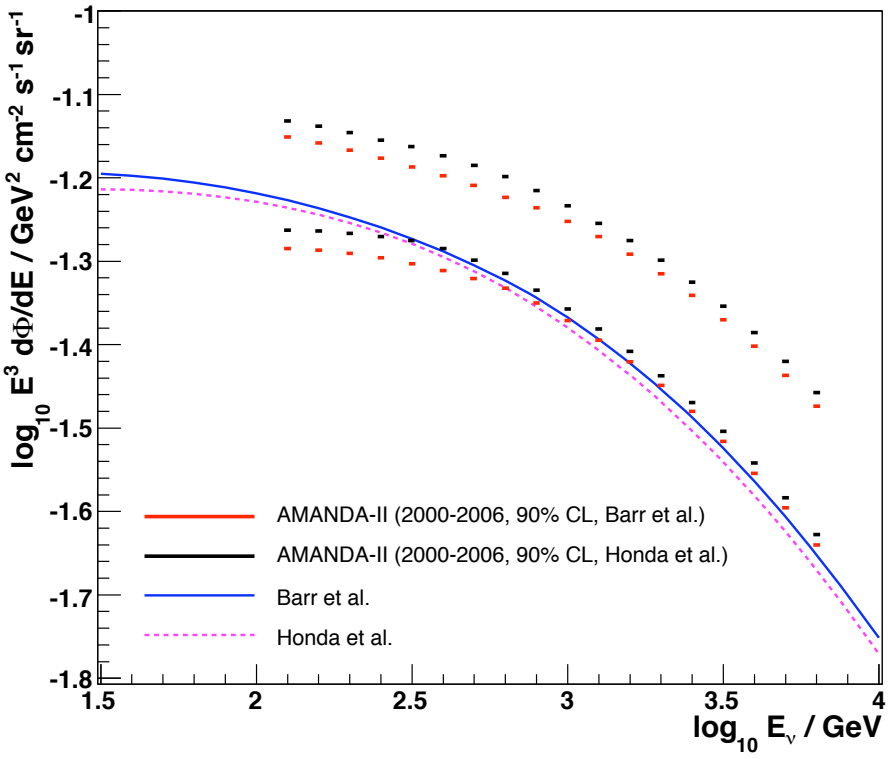

Figure 9.9: Angle-averaged $\nu_\mu + \bar{\nu}_\mu$ atmospheric neutrino fluxes (90% CL allowed ranges), multiplied by E^3, from the forward-folding method applied to both Barr *et al.* and Honda *et al.* models.

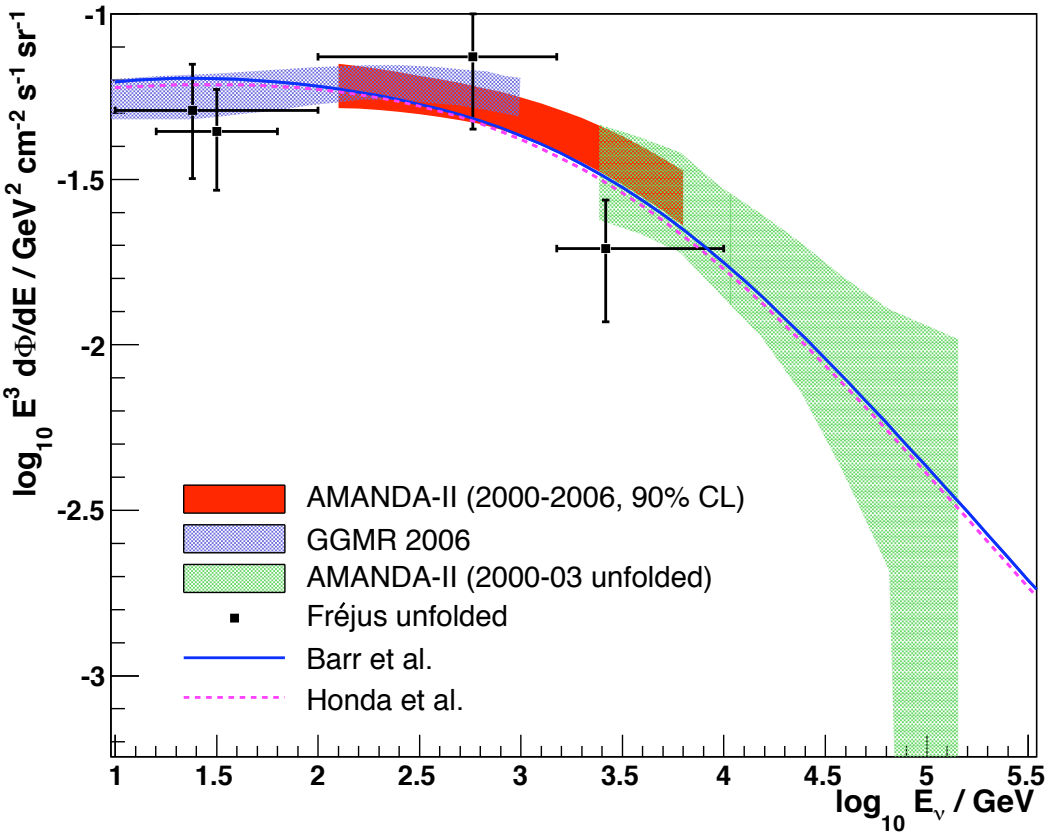

Figure 9.10: Angle-averaged $\nu_\mu + \bar{\nu}_\mu$ atmospheric neutrino flux (solid red band, 90% CL from the forward-folding analysis), multiplied by E^3, compared to results from Super-Kamiokande (blue band), Fréjus (black data points), and a four-year AMANDA-II unfolding analysis (green band).

Chapter 10

Conclusions and Outlook

10.1 Summary

We have set stringent upper limits on both Lorentz violation and quantum decoherence effects in the neutrino sector, with a VLI upper limit at the 90% CL of $\Delta\delta = \Delta c/c < 2.8 \times 10^{-27}$ for VLI oscillations proportional to E, and a QD upper limit at the 90% CL of $D^* < 1.3 \times 10^{-31}$ GeV^{-1} for decoherence effects proportional to E^2. We have also set upper limits on VLI and QD effects with different energy dependencies. Finally, we have determined the atmospheric neutrino spectrum in the energy range from 120 GeV to 7.8 TeV and find a best-fit result that is slightly higher in normalization and with a harder spectral slope than either the Barr *et al.* or Honda *et al.* model. This result is consistent with Super-Kamiokande data and extends that measurement by nearly an order of magnitude in energy.

10.2 Discussion

For an interpretation of the VLI and QD upper limits, we consider natural expectations for the values of such parameters. Given effects proportional to E^2 and E^3, one can argue via dimensional analysis that the new physics parameter should contain a power of the Planck mass $M_{\rm Pl}$ or $M_{\rm Pl}^2$, respectively [114]. For example, given the decoherence parameters D, we may expect

$$\begin{aligned} D &= D^* E_\nu^n \\ &= d^* \frac{E_\nu^n}{M_{\rm Pl}^{n-1}} \end{aligned} \tag{10.1}$$

for $n \geq 2$, and d^* is a dimensionless quantity of $O(1)$ by naturalness. From the limits in table 9.1, we find $d^* < 1.6 \times 10^{-12}$ $(n = 2)$ and $d^* < 910$ $(n = 3)$. For the $n = 2$ case, the decoherence parameter is far below the natural expectation, suggesting either a stronger suppression than described, or that we have indeed probed beyond the Planck scale and found no decoherence of this type.

10.3 Outlook

10.3.1 IceCube

While the AMANDA-II data acquisition system used in this analysis ceased taking data at the end of 2006, the next-generation, km^3-scale IceCube detector is under construction, with completion expected in 2011. IceCube improves upon AMANDA-II in a number of respects. First, its larger size will allow the detection of fainter neutrino sources (see diagram in fig. 10.1). The larger spacing between the strings increases the energy threshold (the exact value depends upon the trigger configuration, which is flexible), but the longer strings reach into the clearest ice below a large dust layer at 2050 m.

Second, IceCube uses digital optical modules (DOMs) in place of AMANDA's analog OMs (prototype DOMs are in use on AMANDA-II's string 18). The DOMs digitize the full PMT waveform in the ice, then transmit it digitally to the surface. The waveform provides substantially more information about the direction and distance to a particle track than the hit times recorded by the AMANDA μ-DAQ. Transmitting data digitally also alleviates any issues with cable dispersion and electrical crosstalk, and modern communication protocol techniques such as error correction can be employed.

In addition to the in-ice array, a surface air shower array ("IceTop") detects cosmic ray air showers. An IceTop station consists of a pair of tanks above each in-ice string, and each tank houses

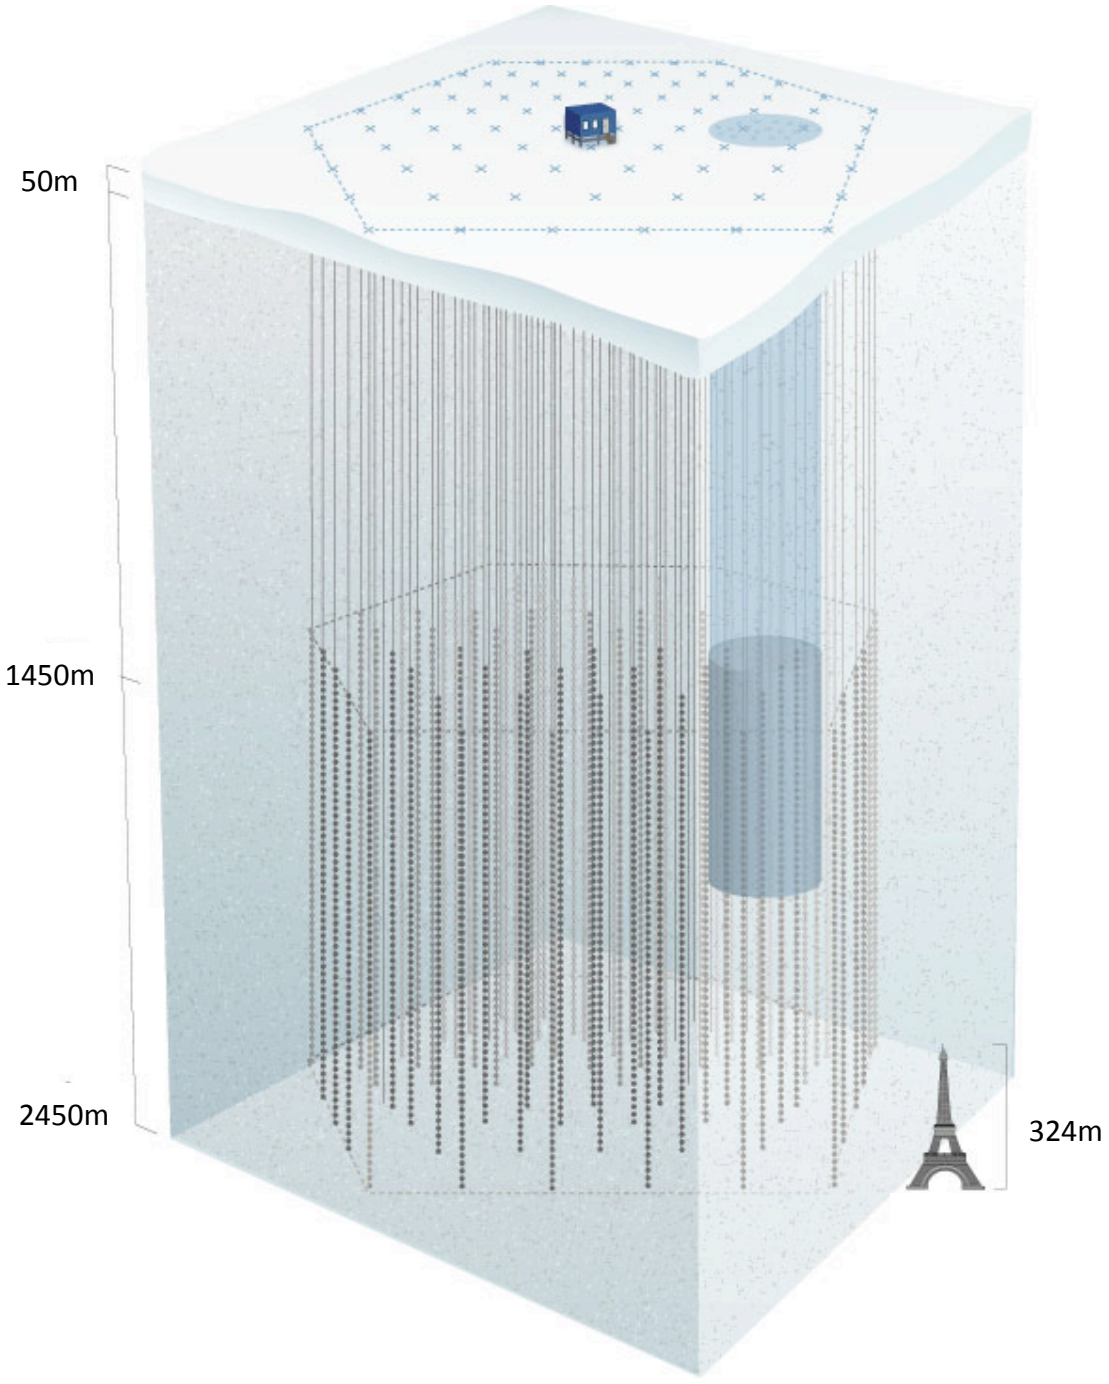

Figure 10.1: Diagram of the the next-generation, km^3-scale IceCube neutrino detector [115]. The darker cylinder marks the extent of the AMANDA-II detector. The Eiffel tower is shown for scale.

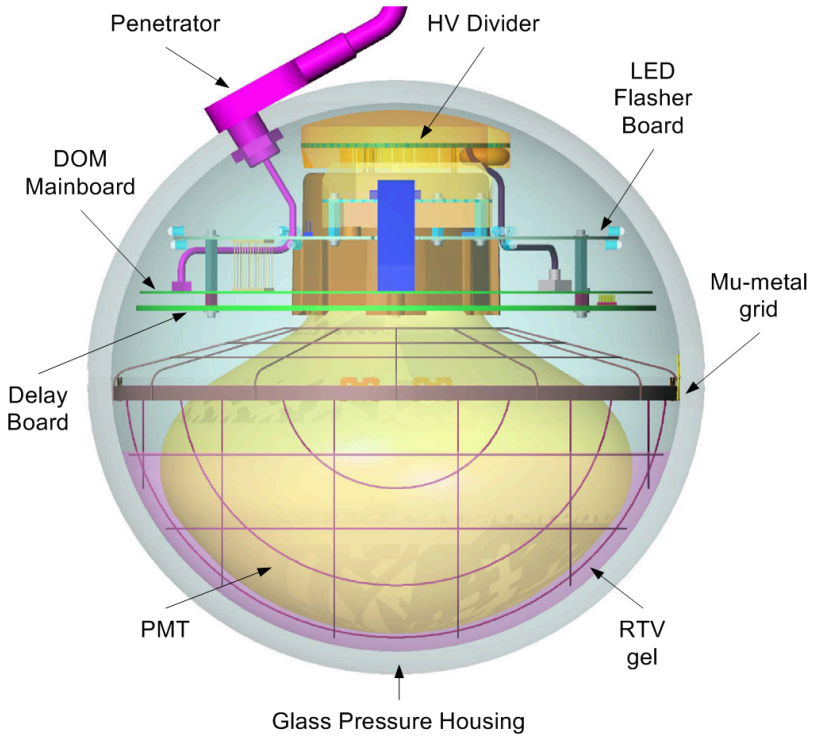

Figure 10.2: Diagram of an IceCube digital optical module (DOM) [115].

two DOMs. The tanks are filled with water that, by means of a circulation and de-gassing system, freezes into clear, bubble-free ice. In addition to fundamental cosmic ray composition studies, IceTop can be used as a veto to assist with atmospheric muon rejection. Figure 10.3 shows a large air shower event that has triggered numerous stations on the surface and a muon bundle that has penetrated into the deep ice.

10.3.2 Sensitivity Using Atmospheric Neutrinos

IceCube has the potential to improve greatly upon the quantum gravity limits obtained with AMANDA-II, as increased statistics of atmospheric neutrinos at the highest energies probe smaller deviations from the Standard Model. In particular, IceCube should be sensitive to $n = 1$ VLI effects an order of magnitude smaller than the limits from this analysis ([110]; see also fig. 9.3). We note that we have also only tested one particular manifestation of VLI in the neutrino sector. A search of the atmospheric neutrino data for a unexpected directional dependence (for example, in right ascension) could probe other VLI effects, such as a universal directional asymmetry (see e.g. [38]).

10.3.3 Astrophysical Tests of Quantum Gravity

Once high-energy astrophysical neutrinos are detected, analysis of the flavor ratio at Earth can probe VLI, QD, and CPT violation [114, 116]. Another technique is to probe VLI via the potential time delays between photons and neutrinos from gamma-ray bursts (GRBs). Given the cosmological distances traversed, this delay could range from 1 μs to 1 year, depending on the power of suppression by $M_{\rm Pl}$ [117]. Detection of high-energy neutrinos from multiple GRBs at different redshifts would allow either confirmation of the delay hypothesis or allow limits below current levels by several orders of magnitude [118]. Such a search is complicated by the low expected flux levels from individual GRBs, as well as uncertainty of any intrinsic $\gamma - \nu$ delay due to production mechanisms in the source (for a further discussion, see [119]). Other probes of Planck-scale physics may be possible as well, but ultimately this will depend on the characteristics of the neutrino sources detected.

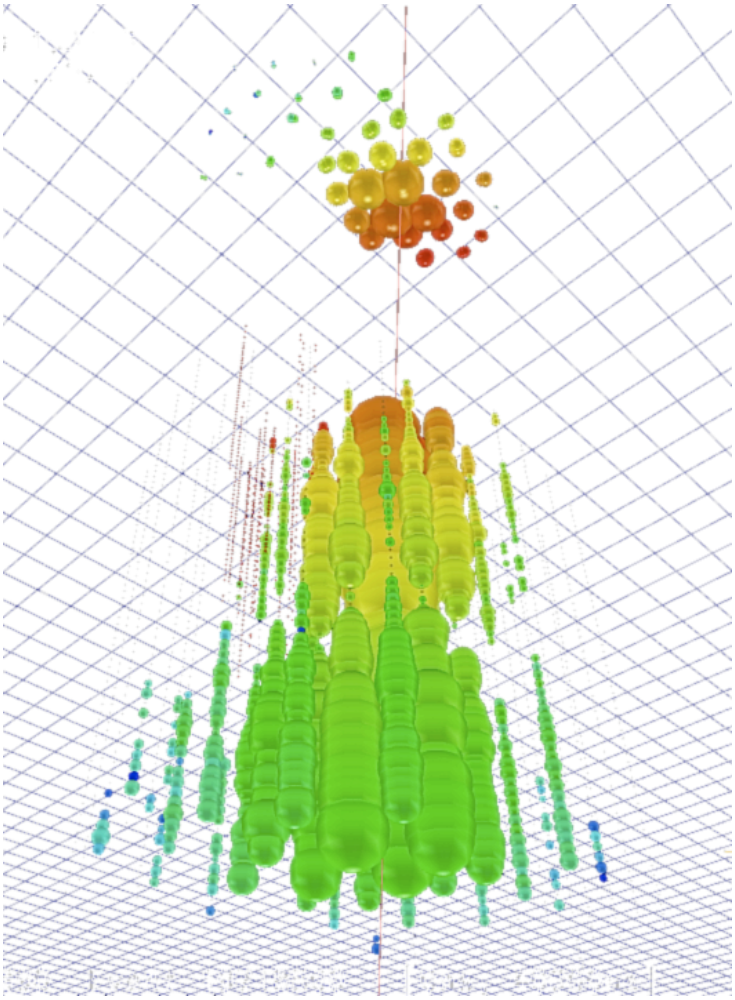

Figure 10.3: Event display of a large coincident air shower and atmospheric muon event in the 40-string IceCube detector, as seen from below. The colors indicate the relative timing of the photon hits, with red being early and blue being late. Note the effect of the large dust layer (the "pinched" region above the green-colored hits).

Bibliography

[1] W. Pauli, in "Rapports du Septième Conseil de Physique Solvay," Brussels, 1933 (Gauthier-Villars, Paris, 1934).

[2] F. Reines and C. L. Cowan, Phys. Rev. **92**, 830 (1953); C. L. Cowan *et al.*, Science **124**, 103 (1956).

[3] S. N. Ahmed *et al.*, Phys. Rev. Lett. **92**, 181301 (2004).

[4] T. Araki *et al.*, Nature **436**, 499 (2005).

[5] K. Hirata *et al.*, Phys. Rev. Lett. **58**, 1490 (1987).

[6] R. M. Bionta *et al.*, Phys. Rev. Lett. **58**, 1494 (1987).

[7] E. Roulet, astro-ph/0011570.

[8] V. F. Hess, Phys. Z. **13**, 1804 (1912).

[9] A. Hillas, astro-ph/0607109.

[10] J. S. Warren *et al.*, Astrophys. Jour. **634**, 376 (2005).

[11] T. K. Gaisser, *Cosmic Rays and Particle Physics* (Cambridge University, U.K., 1991).

[12] J. Hörandel, Astropart. Phys. **21**, 241 (2004).

[13] T. K. Gaisser, Jour. of Phys. Conf. Ser. **47**, 15 (2004).

[14] B. Louis *et al.*, Los Alamos Sci. **25**, 16 (1997).

[15] T. K. Gaisser and M. Honda, Annu. Rev. Nucl. Part. Sci. **52**, 153 (2002).

[16] G. Barr *et al.*, Phys. Rev. D **70**, 023006 (2004).

[17] M. Honda *et al.*, Phys. Rev. D **75**, 043006 (2007).

[18] D. Chirkin, hep-ph/0407078.

[19] P. Lipari, Astropart. Phys. **14**, 153 (2000).

[20] Particle Data Group, *Particle Physics Booklet*. AIP, New York, USA (2004).

[21] Y. Ashie *et al.*, Phys. Rev. Lett. **93**, 101801 (2004).

[22] M. Sanchez *et al.*, Phys. Rev. D **68**, 113004 (2003).

[23] M. Ambrisio *et al.*, Phys. Lett. B **566**, 35 (2003).

[24] Q. R. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 (2002).

[25] M. H. Ahn *et al.*, Phys. Rev. Lett. **90**, 041801 (2003).

[26] M. Maltoni, T. Schwetz, M. Tórtola, and J. W. F. Valle, New Jour. of Phys. **6**, 122 (2004).

[27] R. Gambini and J. Pullin, Phys. Rev. D **59**, 124021 (1999).

[28] J. Madore, S. Schraml, P. Schupp, and J. Wess, Eur. Phys. Jour. C **16**, 161 (2000).

[29] V. A. Kostelecký and S. Samuel, Phys. Rev. D **39**, 683 (1989).

[30] G. Amelino-Camelia, C. Lämmerzahl, A. Macias, and H. Müller, Gravitation and Cosmology: AIP Conf. Proc. **758**, 30 (2005).

[31] D. Mattingly, Living Rev. Relativity **8**, 5 (2005).

[32] S. W. Hawking, Commun. Math. Phys. **87**, 395 (1982).

[33] V. A. Balkanov *et al.*, Astropart. Phys. **12**, 75 (1999).

[34] E. Andrés *et al.*, Nature **410**, 441 (2001).

[35] J. A. Aguilar *et al.*, Astropart. Phys. **26**, 314 (2006).

[36] J. Ahrens *et al.*, Astropart. Phys. **20**, 507 (2004).

[37] D. Colladay and V. A. Kostelecký, Phys. Rev. D **58**, 116002 (1998).

[38] V. A. Kostelecký and M. Mewes, Phys. Rev. D **69**, 016005 (2004).

[39] V. A. Kostelecký and M. Mewes, Phys. Rev. D **70**, 031902 (2004).

[40] T. Katori, V. A. Kostelecký, and R. Tayloe, Phys. Rev. D **74**, 105009 (2006).

[41] S. Coleman and S. L. Glashow, Phys. Rev. D **59**, 116008 (1999).

[42] S. L. Glashow, hep-ph/0407087.

[43] M. C. González-García and M. Maltoni, Phys. Rev. D **70**, 033010 (2004).

[44] G. Battistoni *et al.*, Phys. Lett. B **615**, 14 (2005).

[45] G. L. Fogli, E. Lisi, A. Marrone, and G. Scioscia, Phys. Rev. D **60**, 053006 (1999).

[46] J. Ahrens and J. L. Kelley *et al.*, Proc. of 30th ICRC (Mérida), 2007, arXiv:0711.0353.

[47] D. Morgan, E. Winstanley, J. Brunner, and L. F. Thompson, Astropart. Phys. **29**, 345 (2008).

[48] J. Alfaro, H. Morales-Técotl, and F. Urrutia, Phys. Rev. Lett. **84**, 2318 (2000).

[49] R. Brustein, D. Eichler, and S. Foffa, Phys. Rev. D **65**, 105006 (2002).

[50] M. Gasperini, Phys. Rev. D **38**, 2635 (1988).

[51] A. Halprin and C. N. Leung, Phys. Rev. Lett. **67**, 14 (1991).

[52] G. Z. Adunas, E. Rodriquez-Milla, and D. V. Ahluwalia, Phys. Lett. B **485**, 215 (2000).

[53] J. Ellis, J. S. Hagelin, D. V. Nanopoulos, and M. Srednicki, Nucl. Phys. B **241**, 381 (1984).

[54] G. Lindblad, Commun. Math. Phys. **48**, 119 (1976).

[55] A. M. Gago, E. M. Santos, W. J. C. Teves, and R. Zukanovich Funchal, hep-ph/0208166.

[56] N. E. Mavromatos, gr-qc/0407005.

[57] F. Benatti and R. Floreanini, JHEP **02**, 032 (2000).

[58] D. Morgan, E. Winstanley, J. Brunner, and L. Thompson, Astropart. Phys. **25**, 311 (2006).

[59] G. Barenboim, N. E. Mavromatos, S. Sarkar, and A. Waldron-Lauda, Nucl. Phys. B **758**, 90 (2006).

[60] J. Ellis, N. E. Mavromatos, and D. V. Nanopoulos, Mod. Phys. Lett. A **12**, 1759 (1997); J. Ellis, N. E. Mavromatos, D. V. Nanopoulos, and E. Winstanley, Mod. Phys. Lett. A **12**, 243 (1997).

[61] E. Lisi, A. Marrone, and D. Montanino, Phys. Rev. Lett. **85**, 1166 (2000).

[62] G. L. Fogli, E. Lisi, A. Marrone, and D. Montanino, Phys. Rev. D **67**, 093006 (2003).

[63] G. L. Fogli *et al.*, Phys. Rev. D **76**, 33006 (2007).

[64] R. Gandhi *et al.*, Phys. Rev. D. **58**, 093009 (1998).

[65] K. Greisen, Annu. Rev. Nucl. Sci. **10**, 63 (1960).

[66] F. Reines, Annu. Rev. Nucl. Sci. **10**, 1 (1960).

[67] J. Lundberg, Ph.D. thesis, Uppsala Universitet, Uppsala, Sweden (2008).

[68] S. Eidelman *et al.*, Phys. Lett. B **592**, 252 (2004).

[69] E. Andrés *et* al., Astropart. Phys. **13**, 1 (2000).

[70] M. Kowalski, Ph.D. thesis, Humboldt-Universität zu Berlin, Germany (2004).

[71] M. S. Longair, *High Energy Astrophysics Vol. 1* (Cambridge, U.K., 2004), p. 49.

[72] D. Chirkin and W. Rhode, hep-ph/0407075.

[73] T. R. DeYoung, Ph.D. thesis, Univ. of Wisconsin, Madison, U.S.A. (2001).

[74] J. Ahrens *et al.*, Nucl. Inst. Meth. A **524**, 169 (2004).

[75] M. Ackermann *et al.*, J. Geophys. Res. **111**, D13203 (2006).

[76] G. C. Hill, Astropart. Phys. **6**, 215 (1997).

[77] J. Lundberg *et al.*, Nucl. Inst. and Meth. A **581**, 619 (2007).

[78] S. Hundertmark, Proc. of the Workshop on Simulation and Analysis Methods for Large Neutrino Telescopes, DESY-PROC-1999-01, 276 (1999); S. Hundertmark, Ph.D. thesis, Humboldt-Universität zu Berlin, Germany (1999).

[79] D. Heck *et al.*, Tech. Rep. FZKA 6019, Forschungszentrum Karlsruhe (1998).

[80] M. Ackermann, Ph.D. thesis, Humboldt-Universität zu Berlin, Germany (2006).

[81] D. Pandel, Diploma thesis, Humboldt-Universität zu Berlin, Germany (1996).

[82] A. Achterberg *et al.*, Phys. Rev. D **75**, 102001 (2007).

[83] T. Neunhöffer, Astropart. Phys. **25**, 220 (2006).

[84] A. Pohl, Ph.D. thesis, Uppsala Universitet, Uppsala, Sweden (2004).

[85] G. J. Feldman and R. D. Cousins, Phys. Rev. **D57**, 873 (1998).

[86] A. Stuart, K. Ord, and S. Arnold, *Kendall's Advanced Theory of Statistics* (Arnold, London, 1999), vol. 2A, 6th and earlier editions.

[87] F. James and M. Roos, Comp. Phys. Comm. **10**, 343 (1975).

[88] W. A. Rolke and A. M. Lopez, Nucl. Instrum. Meth. A **458**, 745 (2001).

[89] W. A. Rolke, A. M. Lopez, and J. Conrad, physics/0403059.

[90] G. J. Feldman, "Multiple measurements and parameters in the unified approach," Workshop on Confidence Limits, Fermilab (2000), `http://www.hepl.harvard.edu/~feldman/Journeys.pdf`.

[91] R. Cousins, in *Statistical Problems in Particle Physics, Astrophysics and Cosmology: Proceedings of PHYSTAT05*, edited by L. Lyons and M. Ünel (Univ. of Oxford, U.K., 2005).

[92] M. Sanchez, Ph.D. thesis, Tufts University, Massachusetts, U.S.A. (2003).

[93] G. Punzi, Proceedings of PHYSTAT 05, Oxford, U.K. (2005), arXiv:physics/0511202.

[94] K. Cranmer, Proceedings of PHYSTAT 03, Palo Alto, U.S.A. (2003), arXiv:physics/0310108.

[95] K. Cranmer, physics/0511028.

[96] M. Kowalski and A. Gazizov, astro-ph/0312202.

[97] H. L. Lai *et al.*, Eur. Phys. Jour. C **12**, 375 (2000).

[98] A. D. Martin, R. G. Roberts, and W. J. Stirling, Phys. Lett. B **354**, 155 (1995).

[99] T. K. Gaisser, M. Honda, P. Lipari, and T. Stanev, Proc. of 27th ICRC (Hamburg) **1**, 1643 (2001).

[100] C. G. S. Costa, Astropart. Phys. **16**, 193 (2001).

[101] G. Fiorentini, A. Naumov, and F. L. Villante, Phys. Lett. B **510**, 173 (2001); E. V. Bugaev *et al.*, Nuovo Cimento C **12**, 41 (1989).

[102] R. Enberg, M. H. Reno, I. Sarcevic, Phys. Rev. D **78**, 043005 (2008).

[103] A. D. Martin, M. G. Ryskin, and A. M. Stasto, Acta Phys. Polon. B **34**, 3273 (2003).

[104] T. K. Gaisser, hep-ph/0209195.

[105] V. Agrawal, T. K. Gaisser, P. Lipari, and T. Stanev, Phys. Rev. D **53** 3, 1314 (1996).

[106] R. S. Fletcher, T. K. Gaisser, P. Lipari, and T. Stanev, Phys. Rev. D **50**, 5710 (1994); R. Engel, T. K. Gaisser, P. Lipari, and T. Stanev, Proc. of 26th ICRC (Salt Lake City) **1**, 415 (1999).

[107] K. Werner *et al.*, Phys. Rev. C **74**, 014026 (2006).

[108] S. Ostapchenko, Phys. Lett. B **636**, 40 (2006); S. Ostapchenko, Phys. Rev. D **74**, 014026 (2006).

[109] S. Mizobuchi *et al.*, Proc. of 29th ICRC (Pune) **5**, 323 (2005); A. Djannati-Ataï *et al.*, A&A **350**, 17 (1999).

[110] M. C. González-García, F. Halzen, and M. Maltoni, Phys. Rev. D **71**, 093010 (2005).

[111] M. C. González-García, M. Maltoni, and J. Rojo, JHEP **0610**, 075 (2006).

[112] K. Daum, W. Rhode, P. Bareyre, and R. Barloutaud, Z. Phys. C **66**, 417 (1995).

[113] A. Wiedemann, private communication. This is a corrected version of results that appeared in K. Münich and J. Lünemann *et al.*, Proc. of 30th ICRC (Mérida), 2007, arXiv:0711.0353.

[114] L. A. Anchordoqui *et al.*, Phys. Rev. D **72**, 065019 (2005).

[115] Images from `http://gallery.icecube.wisc.edu`. Material is based upon work supported by the National Science Foundation under grant nos. OPP-9980474 (AMANDA) and OPP-0236449 (IceCube), University of Wisconsin–Madison.

[116] D. Hooper, D. Morgan, and E. Winstanley, Phys. Rev. D **72**, 14 (2005).

[117] G. Amelino-Camelia, Intl. Jour. of Mod. Phys. D **12**, 1633 (2003).

[118] U. Jacob and T. Piran, Nature Phys. **3**, 87 (2007).

[119] M. C. González-García and F. Halzen, JCAP **2**, 008 (2007).

[120] L. Lyons, *Statistics for nuclear and particle physicists*, Cambridge University Press (1986), 12-13.

Appendix A

Non-technical Summary

The 20th century saw the creation of two great theories of physics — one governing the very big, the other the very small. Albert Einstein's General Theory of Relativity, which explains how massive objects produce gravity, revolutionized our view of space and time. The theory predicts all sorts of strange effects, like clocks running slower on Earth than in outer space. Such predictions might seem outlandish, but Global Positioning Satellites (GPS) could not function correctly without taking into account Einstein's theory.

The other great theory is that of quantum mechanics, which governs the behavior of the very small. It operates in the realm of the atom, explaining how tiny elementary particles come together to make up everything around us. If the predictions of general relativity are strange, the world of quantum mechanics is truly bizarre: particles traveling through solid walls, being in two places at once, and other mind-boggling ideas. Despite its strangeness, quantum mechanics has made possible the technology of today, including the computer and the Internet.

Despite many attempts, no one has been able to devise any experimental test that general relativity or quantum mechanics cannot pass. Yet these two great theories cannot apparently be reconciled. If one imagines a scenario to which both theories apply — something very heavy and very small, say — one will quickly run into major problems. What this tells us is that while each theory works amazingly well in its own domain, each by itself is incomplete. What we need is a new theory that combines the two, a theory of *quantum gravity*.

Physicists have been trying for decades to build such a theory. Current efforts go by names such as string theory, loop quantum gravity, non-commutative geometry, and causal set theory, to

name just a few. But no one has succeeded yet, so we continue to look for hints in nature. Perhaps we will be able to find something that doesn't quite agree with relativity or quantum mechanics, and that will lead us toward a theory of quantum gravity.

Studying very high-energy subatomic particles in nature is one way to look for such hints. Conveniently, our Galaxy is already filled with them: extremely high-energy particles called *cosmic rays* bombard Earth all the time. Fortunately, the Earth's atmosphere protects us on the surface, but the cosmic rays still slam into the upper atmosphere and create huge showers of other particles, a train wreck high in the sky with pieces flying to the ground. Within this wreckage is a zoo of other particles: electrons, positrons, muons, pions, kaons, and tiny neutrinos.

The neutrino is related to the particles that make up matter around us, like protons, neutrons, and electrons. For example, if you could pluck a neutron out of the chair you are sitting in, after about 15 minutes, the neutron would decay into a proton, an electron, and a neutrino. The Sun also produces neutrinos all the time as it burns hydrogen into helium via nuclear fusion. But we never notice them; they zip right through us as if we weren't there. Detecting neutrinos requires huge experiments located deep underground, underwater, or under ice, in order to shield the sensitive detectors from the other particles produced in the cosmic ray air showers.

AMANDA-II is one such neutrino detector. It is built deep into the ice at the geographic South Pole. The ice not only acts as a shield, but has another advantage: if you drill deep enough, it is extremely clear. AMANDA-II can use the huge ice sheet at the South Pole as a target for neutrinos. While most of them pass right through without stopping, about every hour, one will crash into the ice and produce another particle called a muon, which emits light as it continues through the detector. By putting very sensitive light detectors into holes drilled into the ice, we can see this light. And by only looking for muons coming from below the detector, we use the entire Earth as a filter to block out other particles. An "up-going" muon could only have been produced by a neutrino that made it most of the way through the Earth.

The neutrinos produced when cosmic rays hit the Earth are known as *atmospheric neutrinos*. During seven years of taking data, AMANDA-II has detected over 5000 of these neutrinos, from various directions and of different energies. By looking for certain unexpected features — for example,

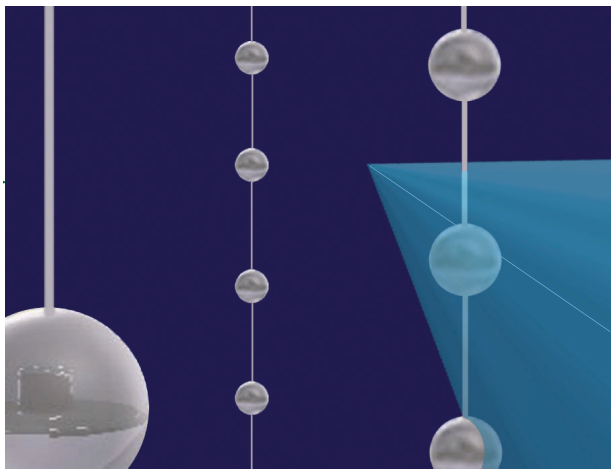

Figure A.1: A high-energy muon, produced from a neutrino collision with the ice or rock, emits light as it travels through the ice. Sensitive light detectors deployed on cables detect this light to track the muon and "see" which direction the neutrino came from.

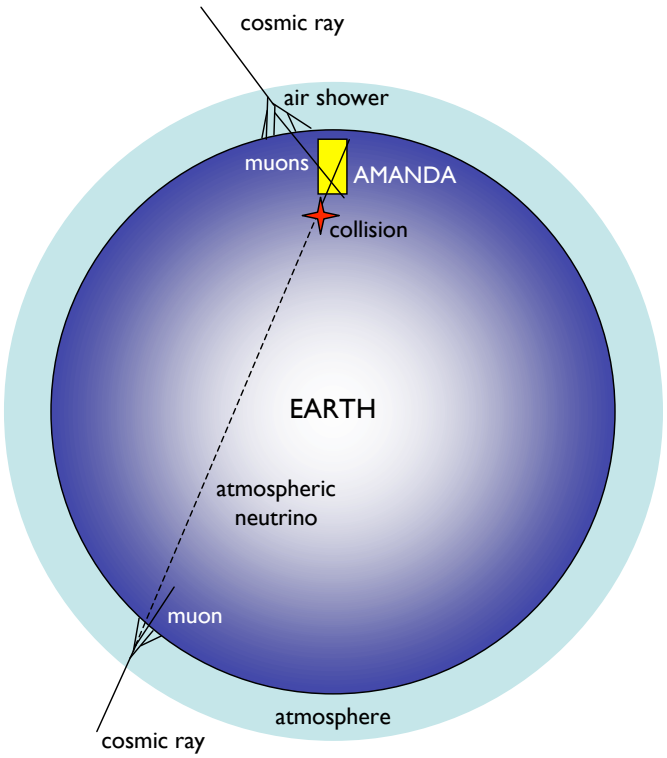

Figure A.2: Diagram (not to scale!) demonstrating the use of Earth as a filter to screen out everything but neutrinos.

missing neutrinos at the highest energies from a certain direction — we might detect a hint of quantum gravity.

How would a neutrino "go missing"? To understand this, we must delve a bit more into the types of particles that make up the Universe. In actuality, there is not just one type of neutrino, but three: electron neutrinos, muon neutrinos, and tau neutrinos. Muon neutrinos produce the light-emitting muon that give them away inside of AMANDA-II. Electron and tau neutrinos can, not surprisingly, produce electrons and tau particles, but these do not create the same nice track of light through the detector that the muon does.

And now we meet again some of the strangeness of quantum mechanics: these three types of neutrinos can actually transform into one another as they travel. When scientists first measured the number of electron neutrinos coming from the Sun, they were shocked to find only one-third of the number they expected. It was only in the 1980s that we understood that the electron neutrinos were transforming into other types on their way to Earth. Exactly why this happens is an interesting but complicated story, but suffice it to say that the fact that the neutrinos have a tiny bit of mass allows these "neutrino oscillations" to occur.

We do not expect this type of neutrino oscillation in the atmospheric neutrinos detected by AMANDA-II, but transformations of neutrinos from one type to another can also be caused by quantum gravity. Neutrinos could travel at slightly different speeds than what we expect, or they could run into a "frothiness" of space itself caused by quantum gravity. Either of these possibilities[1] would cause our atmospheric muon neutrinos to change into another type, so if we just counted the muon neutrinos, we'd find some were missing. And that would be a tell-tale sign of quantum gravity.

This analysis has done just that — used AMANDA-II to count the muon neutrinos coming from different parts of the sky, at different energies, and looked to see if any are missing. As it turns out, they are all there. Predictions of how many to expect, without invoking quantum gravity, are right. So, the search for quantum gravity continues — but we can now put a limit on how big these effects from quantum gravity are.

---

[1]The technical terms for each of these quantum gravity effects are "violation of Lorentz invariance" or VLI, and "quantum decoherence" or QD, both of which you will see throughout this thesis.

Appendix B

Effective Area

The neutrino effective area is defined so that the number of events detected is

$$N_{\text{events}} = \int dE_\nu \; d\Omega \; dt \; \Phi(E_\nu, \theta, \phi) \; A^\nu_{\text{eff}}(E_\nu, \theta, \phi) \tag{B.1}$$

for a differential neutrino flux $\Phi(E_\nu, \theta, \phi)$. Given the lack of a suitable neutrino calibration beam, we generally calculate effective area using simulation. If we have a set of neutrino MC events that we can weight to an arbitrary flux, we can use the same events to find the effective area.

Given a set of N_{gen} unweighted MC events, the number of detected events given a flux Φ is

$$N_{\text{events}} = \sum_{i=1}^{N_{\text{gen}}} \Phi(E_i, \theta_i, \phi_i) \; w_i \; , \tag{B.2}$$

where w_i is the per-event weight needed to reweight from the generation spectrum back to an E^0 flux. Combining the previous two equations, we have

$$\int dE \; d\Omega \; dt \; \Phi(E_\nu, \theta, \phi) \; A^\nu_{\text{eff}}(E_\nu, \theta, \phi) = \sum_{i=1}^{N_{\text{gen}}} \Phi(E_i, \theta_i, \phi_i) \; w_i \tag{B.3}$$

for any flux Φ.

If we want to calculate the effective area for a given energy E_0 and angles (θ_0, ϕ_0), we can calculate the average over a small energy range $E_0 \pm \Delta E$ and solid angle range $\Omega_0 \pm \Delta\Omega$:

$$T \int_{E_0-\Delta E/2}^{E_0+\Delta E/2} dE_\nu \int_{\Omega_0-\Delta\Omega/2}^{\Omega_0+\Delta\Omega/2} d\Omega\ \Phi(E_\nu,\theta,\phi)\ A^\nu_{\rm eff}(E_\nu,\theta,\phi) = \sum_i^{E,\Omega\ \text{range}} \Phi(E_i,\theta_i,\phi_i)\ w_i \tag{B.4}$$

where we have also assumed the flux and effective area are independent of time, so we can integrate out the livetime T. Now, we can set Φ to whatever makes the calculation easiest, so we choose $\Phi = E^0$:

$$T \int_{E_0-\Delta E/2}^{E_0+\Delta E/2} dE_\nu \int_{\Omega_0-\Delta\Omega/2}^{\Omega_0+\Delta\Omega/2} d\Omega\ A^\nu_{\rm eff}(E_\nu,\theta,\phi) = \sum_i^{E,\Omega\ \text{range}} w_i\ . \tag{B.5}$$

So, approximately, we have

$$T\ A^\nu_{\rm eff}(E_0,\Omega_0)\ \Delta E\ \Delta\Omega = \sum_i^{E,\Omega\ \text{range}} w_i \tag{B.6}$$

and thus

$$A^\nu_{\rm eff}(E_0,\theta_0,\phi_0) = \sum_i^{E,\Omega\ \text{range}} \frac{w_i}{\Delta E\ \Delta\Omega\ T}\ . \tag{B.7}$$

We note that in practice, an easy way to compute the above is to form a histogram of the MC events versus true neutrino energy. For example, we can calculate the azimuth-averaged effective area in a given zenith angle range $\Delta\cos\theta$ by forming a histogram of the events versus true neutrino energy, weighted with the quantity $w_i/(2\pi\Delta\cos\theta\ \Delta E\ T)$, where ΔE is the histogram bin width. If one wishes to visualize the effective area versus $\log_{10} E$, we can easily do this by substituting the flux $\Phi = E^{-1}$ into eq. B.4:

$$T \int_{\Delta E} dE \int_{\Delta\Omega} d\Omega\ \frac{A^\nu_{\rm eff}(E,\theta,\phi)}{E} = \sum_i^{E,\Omega\ \text{range}} \frac{w_i}{E_i}\ , \tag{B.8}$$

so

$$T \int_{\Delta \log E} d\log E \int_{\Delta\Omega} d\Omega\ \ln 10\ A^\nu_{\rm eff}(E,\theta,\phi) = \sum_i^{E,\Omega\ \text{range}} \frac{w_i}{E_i} \tag{B.9}$$

and thus

$$A^{\nu}_{\text{eff}}(E_0, \theta_0, \phi_0) = \sum_i^{E,\Omega \text{ range}} \frac{w_i}{\ln 10 \cdot E_i \cdot \Delta \log E \cdot \Delta\Omega \cdot T} \ . \tag{B.10}$$

The event weight for a histogram over $\log_{10} E$ thus changes to the summand of the previous equation.

The above calculation is valid either for ν or $\bar{\nu}$ — the effective area for neutrinos and antineutrinos is in general not the same because of the different cross sections. If we wish to calculate the average effective area for ν and $\bar{\nu}$, we can sum the weights of both but add a factor of $\frac{1}{2}$:

$$A^{\nu,\bar{\nu}}_{\text{eff}}(E_0, \theta_0) = \sum_i^{E,\Omega \text{ range}} \frac{w_i^{\nu,\bar{\nu}}}{2 \ln 10 \cdot E_i \cdot \Delta \log E \cdot \Delta\Omega \cdot T} \ . \tag{B.11}$$

Finally, we put in explicit quantities for the weights w_i to demonstrate the equivalence of this approach with other working definitions of the effective area. Suppose we have a MC sample generated with a power-law spectrum $E^{-\gamma}$ from E_L to E_H, with N_{gen} events *each* of neutrinos and antineutrinos. Then the weight w_i is

$$w_i = \frac{P_i}{C \ E_i^{-\gamma}} \ , \tag{B.12}$$

where P_i is the interaction probability. The flux normalization constant C must be chosen such that

$$N_{\text{gen}} = \int_{E_L}^{E_H} dE \int_{A_{\text{gen}}} dA \int_{\Omega_{\text{gen}}} d\Omega \int_T dt \ C \ E^{-\gamma} \tag{B.13}$$

for a generation area of A_{gen} (which could depend on, for example, the zenith angle), solid angle generation of Ω_{gen}, and livetime T. Solving for C, we find

$$C = \frac{N_{\text{gen}}}{C_E \cdot A_{\text{gen}} \cdot \Omega_{\text{gen}} \cdot T} \ , \tag{B.14}$$

where the constant C_E is given by the definite integral[1]

$$C_E = \int_{E_L}^{E_H} E^{-\gamma} \ dE \ . \tag{B.15}$$

[1] As an example, for $\gamma = 1, E_L = 10$ GeV, and $E_H = 10^8$ GeV, $C_E = 16.12$.

While expressions B.7 and B.10 are useful for calculation, one may come across other working definitions of the effective area, based on the number of detected events vs. the number generated. To see the equivalence, note that the fraction of generated events in a region $(\Delta \log E, \Delta\Omega)$ is

$$\begin{aligned} \Delta N_{\rm gen} &= N_{\rm gen} \, \frac{\int_{\Delta \log E} \int_{\Delta\Omega} d\log E \; d\Omega \; \ln 10 \; E \; E^{-\gamma}}{C_E} \\ &\approx N_{\rm gen} \, \frac{\ln 10 \cdot E^{-\gamma+1} \cdot \Delta \log E \cdot \Delta\Omega}{C_E \; \Omega_{\rm gen}} \; . \end{aligned} \tag{B.16}$$

Using this with the event weight given in eq. B.12 and the expression for the effective area in eq. B.10, we find

$$\begin{aligned} A^{\nu}_{\rm eff}(E_0, \theta_0, \phi_0) &= \sum_{i}^{E,\Omega \text{ range}} \frac{P_i}{C \; E_i^{-\gamma}} \cdot \frac{1}{\ln 10 \cdot E_i \cdot \Delta \log E \cdot \Delta\Omega \cdot T} \\ &= \sum_{i}^{E,\Omega \text{ range}} P_i \; A_{\rm gen} \, \frac{C_E \; \Omega_{\rm gen}}{N_{\rm gen} \cdot \ln 10 \cdot E_i^{-\gamma+1} \cdot \Delta \log E \cdot \Delta\Omega} \\ &= \frac{\sum_{i}^{E,\Omega \text{ range}} P_i \; A_{\rm gen}}{\Delta N_{\rm gen}} \; . \end{aligned} \tag{B.17}$$

The latter expression is often used as a definition of effective area; we note here its equivalence with the arguably more fundamental definition given in eq. B.1.

Appendix C

Reweighting of Cosmic Ray Simulation

In order to optimize simulation time, the 2005 AMANDA dCORSIKA cosmic ray air shower Monte Carlo is generated with a harder spectrum than is present in nature. Using the DSLOPE steering file option, the primary spectrum is altered roughly from $E^{-2.7}$ to $E^{-1.7}$. More accurately, the DSLOPE applies a slope difference to each component of the slope parametrization used, which in our case is Hörandel. The generated events must then be reweighted to the original spectrum and an appropriate normalization factor applied.

C.1 Event Weighting (Single Power Law)

In this section, we derive a simpler reweighting result for a single-component power-law spectrum $E^{-\gamma}$ generated with minimum energy E_L and maximum energy E_H. We first derive the normalization on the flux generating N events:

$$N = \int_{E_L}^{E_H} dE\ A\ E^{-\gamma}\ , \tag{C.1}$$

and in our case, since $E_H \gg E_L$, we approximate $E_H \approx \infty$. We then have for the normalization factor

$$A \approx \frac{N\ (\gamma - 1)}{E_L^{-\gamma+1}}\ . \tag{C.2}$$

So, suppose we are generating a sample of N events with a modified slope of $\tilde{\gamma} = \gamma + \Delta$, where $\Delta =$ DSLOPE (note the sign convention here). This sample corresponds to a flux of

$$\frac{N\ (\tilde{\gamma}-1)}{E_L^{-\tilde{\gamma}+1}}\ E^{-\tilde{\gamma}}\ . \tag{C.3}$$

We will apply two weighting factors to weight the sample as if it were N events of the original slope: $w_E(E)$ to correct the slope, and w_N to correct the normalization. That is,

$$w_N\ w_E(E)\ \frac{N\ (\tilde{\gamma}-1)}{E_L^{-\tilde{\gamma}+1}}\ E^{-\tilde{\gamma}}\ = \frac{N\ (\gamma-1)}{E_L^{-\gamma+1}}\ E^{-\gamma}\ . \tag{C.4}$$

By inspection, we have:

$$w_N = \frac{\gamma-1}{\tilde{\gamma}-1}\ \frac{E_L^{-\tilde{\gamma}+1}}{E_L^{-\gamma+1}} \tag{C.5}$$

and

$$w_E(E) = \frac{E^{-\gamma}}{E^{-\tilde{\gamma}}}\ . \tag{C.6}$$

Therefore, the event weight for the ith event w_i is:

$$\begin{aligned} w_i &= w_N\ w_E(E_i) \\ &= \frac{\gamma-1}{\tilde{\gamma}-1}\ \frac{E_L^{-\tilde{\gamma}+1}}{E_L^{-\gamma+1}} \frac{E_i^{-\gamma}}{E_i^{-\tilde{\gamma}}} \\ &= \frac{\gamma-1}{\gamma-1+\Delta}\ E_L^{-\Delta}\ E_i^{\Delta}\ . \end{aligned} \tag{C.7}$$

In the specific case of the 2005 dCORSIKA simulation, $\Delta = -1$, $\gamma = 2.7$, and $E_L = 800$ GeV, so the approximate weight w_i would be

$$w_i = \frac{1.7}{0.7}\ \frac{800\ \text{GeV}}{E_i\ (\text{GeV})}\ . \tag{C.8}$$

However, see section C.3 for a more accurate result.

C.2 Livetime

The advantage of reweighting the normalization back to N events of the unmodified spectrum is that it makes the livetime reweighting simple, since one can now use the livetime of an unmodified dCORSIKA run. We simply use an additional factor w_L,

$$w_L = \frac{T}{N_{file}\ t_{file}}\ , \tag{C.9}$$

where T is the data livetime (including any prescaling factors), N_{file} is the number of MC files, and t_{file} is the livetime of one file with an *unmodified* spectrum (with the same number of events, of course). This technique avoids any confusion about dCORSIKA / ucr calculation of livetimes on runs with modified DSLOPE.

For the 2005 CORSIKA MC, $t_{file} = 0.0787$ s (obtained with one run with an unmodified spectrum), and the livetime of the filtered data set is 199.25 d. This results in a final livetime-adjusted weight of:

$$w_L\ w_i = \frac{1.7}{0.7}\ \frac{800\ \text{GeV}}{E_i\ (\text{GeV})}\ \frac{199.25 \cdot 86400\ \text{s}}{N_{file} \cdot 0.0787\ \text{s}}\ . \tag{C.10}$$

C.3 Event Weighting (Hörandel)

One finds in practice that the expression in eq. C.10 is only accurate to about 30% when applied to dCORSIKA generated with Hörandel. This is for two reasons: first, there are multiple components in the flux, each with different spectral slope γ_k; second, with the SPRIC steering card enabled, the minimum energy for a primary with mass A_k amu is $A_k \cdot E_L$, not E_L.

In theory, one could construct a composite expression using the above equations for each component, knowing the parameters of the Hörandel flux; however, dCORSIKA makes our life easier by providing a composite integral FLUXSUM in the log file. The only trick is that the energy integral is calculated internally in units of TeV, so when using it in our reweighting expression, we need to correct for this.

The analogue to expression C.7 using the FLUXSUM approach is

$$w_i = \frac{\tilde{F}_\gamma}{F_\gamma}\ 1000^{-\Delta}\ E_i^{\Delta}\ , \tag{C.11}$$

where, as before, $\tilde{\gamma} = \gamma + \Delta$, and $\tilde{F}_\gamma$ is the FLUXSUM for one file with slope $\tilde{\gamma}$, and F_γ is the FLUXSUM for one file with slope γ.

The FLUXSUM ratio allows us to use the original unweighted livetime in the full weight, so adding the same livetime weight w_L, we have

$$w = w_i\ w_L = \frac{\tilde{F}_\gamma}{F_\gamma}\ 1000^{-\Delta}\ E_i^{\Delta}\ \frac{T}{N_{file}\ t_{file}}\ , \tag{C.12}$$

where T is the data livetime (including any prescaling factors), N_{file} is the number of MC files, and t_{file} is the livetime of one file with an unmodified spectrum (with the same number of events). The $1000^{-\Delta}$ term corrects the units to GeV.

However, knowing the unit conversion of FLUXSUM allows us to use the modified livetime as calculated by ucr — in which case we can calculate w without using ratios between two different dCORSIKA runs. Equivalently,

$$w = 1000^{-\Delta}\ E_i^{\Delta}\ \frac{T}{N_{file}\ \tilde{t}_{file}}\ , \tag{C.13}$$

where $\tilde{t}_{file}$ is the livetime reported by ucr for one file with *modified* spectrum.

For the original 2005 dCORSIKA generation, with Hörandel spectrum, $E_L = 800$ GeV, DSLOPE $= -1$, and 10K events/file, this results in a final reweighting expression of

$$w(\Delta = -1) = \frac{1000}{E_i(\mathrm{GeV})}\ \frac{199.25 \cdot 86400\ \mathrm{s}}{N_{file}\ 0.0306\ \mathrm{s}}\ . \tag{C.14}$$

More recently, we have generated another sample of 2005 MC with DSLOPE $= -0.4$ and 1M events/file. For this sample, $\tilde{t}_{file} = 6.20$ s, so the weight is

$$w(\Delta = -0.4, \mathrm{1M\ events}) = \frac{15.8}{\mathrm{E_i^{0.4}}}\ \frac{199.25 \cdot 86400\ \mathrm{s}}{\mathrm{N_{file}}\ 6.20\ \mathrm{s}}\ . \tag{C.15}$$

Appendix D

Effective Livetimes and their Applications

We consider the problem of determining the effective livetime of a sample of weighted events (such as in Monte Carlo simulations). We derive the expression for an effective livetime, then provide a few illustrative applications: optimization of cosmic ray simulation, and (more speculatively) estimating the statistical error on zero Monte Carlo events.

D.1 Formalism

We first present the idea of an effective number of events n_{eff}, as in [120]. For a set of n weighted events with observable x_i, each with weight w_i, the total number of weighted events T, given n simulated events, is

$$T = \sum_{i=1}^{n} w_i \ , \tag{D.1}$$

and the variance σ^2 is

$$\sigma^2 = \sum_{i=1}^{n} w_i^2 \ . \tag{D.2}$$

This leads naturally to the idea of an effective number of events n_{eff}, defined so that the fractional Poisson error on n_{eff} is the same as the weighted sample[1]:

[1] In ROOT, n_{eff} can be computed with TH1::GetEffectiveEntries().

$$\frac{n_{\text{eff}}^2}{(\sqrt{n_{\text{eff}}})^2} = \frac{T^2}{\sigma^2} \tag{D.3}$$

$$n_{\text{eff}} = \frac{T^2}{\sigma^2} = \frac{(\sum_{i=1}^{n} w_i)^2}{\sum_{i=1}^{n} w_i^2}. \tag{D.4}$$

D.1.1 Constant Event Weight

For a constant weight $w_i = w \; \forall \; i$, the effective number of events is just the unweighted number of Monte Carlo events:

$$n_{\text{eff}} = \frac{(\sum_{i=1}^{n} w_i)^2}{\sum_{i=1}^{n} w_i^2} = \frac{n^2 w^2}{n \; w^2} = n \; . \tag{D.5}$$

Equivalently, one can view the weight w as the ratio between the weighted number of events T and n_{eff}:

$$\frac{T}{n_{\text{eff}}} = \frac{n \; w}{n} = w \; . \tag{D.6}$$

Also note that in the case of constant weight w, the error σ on the weighted number of events T is just $w\sqrt{n_{\text{eff}}}$:

$$\sigma = \sqrt{\sum_{i=1}^{n} w_i^2} = w\sqrt{n} = w\sqrt{n_{\text{eff}}} \; . \tag{D.7}$$

One can view the weight w in terms of an effective livetime for the Monte Carlo sample, which provides a more intuitive feeling of how the errors are scaling. Specifically, if we are simulating a data sample (or integer-valued distribution) with livetime L, using Monte Carlo events with weight w, our effective livetime L_{eff} for the Monte Carlo sample is simply

$$L_{\text{eff}} = \frac{L}{w} \; . \tag{D.8}$$

Viewed this way, w is the fraction L/L_{eff} by which we must scale the Monte Carlo distribution to result in one that has a Poisson variance. This also is equivalent to how we normally calculate

simulation livetimes in the case of constant event weights, using the ratio of the data events (or weighted MC events normalized to data) to the number of simulated MC events:

$$L_{\text{conv}} = L\ \frac{n}{N_{\text{data}}} = L\ \frac{n}{T} = L\ \frac{n}{\sum_{i=1}^{n} w} = \frac{L}{w} = L_{\text{eff}}\ . \tag{D.9}$$

D.1.2 Variable Event Weights

In many cases, Monte Carlo events have variable weights (such as in the case of spectral reweighting). We want to find the effective "average" weight $\tilde{w}$ that we can use to calculate an effective livetime. To do this, we generalize equation D.6:

$$\begin{aligned} \tilde{w} &= \frac{T}{n_{\text{eff}}} \\ &= \frac{\sum w_i\ \sum w_i^2}{(\sum w_i)^2} \\ &= \frac{\sum w_i^2}{\sum w_i}\ . \end{aligned} \tag{D.10}$$

The weight $\tilde{w}$ is the *contraharmonic mean* of the w_i, and for $w_i = w\ \forall\ i$, one can check that the above reduces to $\tilde{w} = w$. We also note that this definition of $\tilde{w}$ is equivalent to generalizing equation D.7, so that

$$\sigma = \tilde{w}\sqrt{n_{\text{eff}}}\ . \tag{D.11}$$

In the language of livetimes, we are now in the position to define a effective livetime of a Monte Carlo subsample with variable event weights. Specifically, for a sample of events with weights w_i representing a data sample with livetime L, the effective livetime is

$$L_{\text{eff}} = \frac{L}{\tilde{w}} = \frac{L\ \sum w_i}{\sum w_i^2}\ . \tag{D.12}$$

Because $\tilde{w}$ is a function of the event subsample, one can define concepts like "the effective livetime in bin 10" or "the effective livetime above 100 GeV".

Note that while we derived this expression based on our definition of $\tilde{w}$, equation D.12 is equivalent to another intuitive definition of L_{eff} based on n_{eff}:

$$L\frac{n_{\text{eff}}}{T} = \frac{L\ \sum w_i}{\sum w_i^2} = L_{\text{eff}}\ . \tag{D.13}$$

This is the variable-event-weight analogue to expression D.9.

D.2 Application 1: Cosmic Ray Simulation

As a first application and sanity check of this definition, we apply the above formalism to the problem of cosmic ray simulation, in which one frequently simulates a harder spectrum than desired and then reweights to the original.

Specifically, we consider simulation of a power law spectrum $E^{-\gamma}$, using different spectral slopes $E^{-(\gamma+\Delta)}$. The event weights w_i for this case are

$$w_i = \frac{\gamma - 1}{\gamma - 1 + \Delta}\ E_L^{-\Delta}\ E_i^{\Delta}\ , \tag{D.14}$$

where E_L is the low-energy bound for the simulation, and where the high-energy bound $E_H \gg E_L$ (see appendix C).

One can then generate a small sample of events with different Δ and compare the effective livetimes of events that trigger our detector (AMANDA-II, in this case), as shown in table D.1. First, we note that the effective livetime is behaving as desired, and the effective livetime of high-energy events keeps rising as the spectrum gets harder. The effective livetime of low-energy events, however, starts to get worse as we oversample high energies and then reweight to a steep power law.

Because of the energy-dependent effective area of our detector, this leads to an optimal Δ to maximize the effective livetime of events at trigger level (in this case, $\Delta_{\text{best,L0}} = -0.6$). We can also find the energy range of events that survive to higher filter levels (say, level 3 of the 2005 filtering) and use this to estimate the best Δ for maximizing livetime at L3 (in this case, $\Delta_{\text{best,L3}} = -0.8$, because the energy peak at L3 is slightly higher than at L0).

Furthermore, one can take into account the variable (in some cases, nonlinear) simulation times for the different spectra (see the runtime column in table D.1). Then one can choose the spectrum

Δ	$\gamma + \Delta$	Runtime (s)	Trig. events	L_{eff} (s)	L_{eff} (s) $E < 5$ TeV	L_{eff} (s) $E > 5$ TeV	L_{eff} (s) est. L3	L_{eff}/ runtime est. L3
0	-2.7	154	33	0.39	0.39	0.39	0.39	0.0025
-0.2	-2.5	176	44	0.59	0.46	0.67	0.59	0.0034
-0.4	-2.3	299	99	1.0	0.55	1.1	1.0	0.0034
-0.6	-2.1	508	188	1.3	0.57	1.9	1.3	0.0025
-0.8	-1.9	1454	361	1.2	0.54	2.4	1.5	0.0011
-1.0	-1.7	3745	875	1.2	0.50	3.5	1.5	0.0004

Table D.1: Effective livetimes for cosmic ray MC samples with varying spectral slope. 50K events were simulated with dCORSIKA + SIBYLL, triggering AMANDA-II using AMASIM.

with the highest livetime to runtime ratio. For optimizing effective livetime to runtime at level 3, $\Delta_{\text{opt,L3}} = -0.4$. Note that simulation with $\Delta = -1$ is a factor of 6 times less efficient that using no slope change at all!

Of course, because the effective livetime depends on the event sample, the optimal Δ will depend on the specific filtering scenario for which one is optimizing. For high-energy filters, the harder slopes may be better, but keep in mind that this is only true if one has removed most of the low-energy events — otherwise their large weights will lower the livetime.

D.3 Application 2: The Error on Zero

Consider a Monte Carlo simulation of some binned distribution $f_i(x)$ of an event observable x (f is integer-valued in bins i), which falls off to zero at high x. A simulation of this distribution will fall to zero at some $x > x_0$. We argue that the statistical error on this bin must depend on the number of simulated events n (unweighted) with $x < x_0$.

D.3.1 A Worst-case Scenario

Consider a worst-case scenario in which we have a single Monte Carlo event in bin j representing $f_i(x)$, that is, $n_j = 1$. Then the weight for this event in bin j is roughly the number of events f_j, if (as is most likely) the distribution peaks in bin j. The number of simulated events in bin $j + 1$ is zero by construction, but the number of expected events f_{j+1} can be arbitrarily large depending on the distribution. Intuitively, we expect that the error on the simulated value $n_{j+1} = 0$ should be

quite large, and ideally should cover the expected value f_{j+1}.

Specifically, let's suppose the expected number of events in bin j is 100, and the expected number of events in bin $j+1$ is 75, and that our single event is in bin j. So we could expect the weight $w \approx 100$ if the distribution peaks around this value, and thus

$$\sigma \approx w\sqrt{n} = 100 \tag{D.15}$$

and $T_j = 100 \pm 100$.

With an idea toward extending this to the $j+1$ bin, instead of using the error $\sqrt{n}$ above, we might also consider the Feldman-Cousins confidence interval [85] for $n_{obs} = 1$, which gives $\mu_{1\sigma} \in [0.37, 2.75]$, where μ is the "true" number of expected events (with infinite Monte Carlo). Then the weighted confidence interval is $w \cdot \mu \in [37, 275]$, or $T_j = 100^{+175}_{-63}$.

Now, in the $j+1$ bin, we have $n_{obs} = 0$, but now the event weight w_i is undefined. However, we're still considering the case of a constant event weight, so we set $w = 100$ again. Now for $n_{obs} = 0$, the Feldman-Cousins confidence interval for the mean μ is $\mu_{1\sigma} \in [0, 1.29]$. Then the weighted confidence interval for this bin is $[0, 129]$, that is

$$T_{j+1} \approx 0^{+129}_{-0} \ . \tag{D.16}$$

Our hypothetical expected value for T_{j+1}, 75, lies within this interval, but we note this is because a) the weight w is a decent approximation for the expected value in bin j, and b) the expected value of bin $j+1$ is close to that of bin j. These are heuristic conditions for this approximation to remain meaningful.

Despite all the hand waving, we are better off than before in that we have a handle on the statistical error on the simulated zero events in bin $j+1$, and we have an idea of how this depends on the event weighting.

Specifically, for constant event weight w and $n_{j+1} = 0$, we have

$$T_{j+1} \approx 0^{+w\cdot\mu_{CL}}_{-0} \ , \tag{D.17}$$

where $\mu_{1\sigma} = 1.29$ and $\mu_{90} = 2.44$.

D.3.2 Variable Event Weights

For variable event weights, we return to our "average" weight $\tilde{w}$ as defined in equation D.10. We still have the problem, however, of the event weights being undefined in the zero bin. To approximate the weighting in this region, we construct a sequence $\tilde{w}_1, \tilde{w}_2, \tilde{w}_3, ...$ where

$$\tilde{w}_k = \frac{\sum_{bin=j-k-1}^{j} w_i^2}{\sum_{bin=j-k-1}^{j} w_i} \tag{D.18}$$

or, alternatively,

$$\tilde{w}_k = \frac{\sum_{bin=j-k-1} w_i^2}{\sum_{bin=j-k-1} w_i} \tag{D.19}$$

and bin $j+1$ is the first bin with zero simulated events. Then we construct an approximate limit (which is really just an extrapolation)

$$\tilde{w}_0 = \lim_{k \to 0} \tilde{w}_k \ . \tag{D.20}$$

Then we use the estimated $\tilde{w}_0$ to construct the error on the zero bin $j+1$:

$$T_{j+1} \approx 0^{+\tilde{w}_0 \cdot \mu_{CL}}_{-0} \ . \tag{D.21}$$

From the viewpoint of effective livetimes, the sequence of $\tilde{w}_k$ extrapolated to $\tilde{w}_0$ can be seen as a sequence of effective livetimes L_k extrapolated to some estimated livetime representing the bin with zero events, L_0. The contents and error on that bin can equivalently be written as $0^{+(L/L_0)\cdot\mu_{CL}}_{-0}$.

Currently, we make no statement about the coverage of this modified confidence interval, as the accuracy of this approximation is dependent specifically on the weighting scheme and the shape of the observable distribution.

D.3.3 An Example

As an illustration of this error procedure, we consider the simulation of the number of optical modules hit (N_{ch}) in AMANDA-II by cosmic-ray muons, an energy-correlated observable. A plot of this distribution, simulated with a harder spectrum ($\Delta = -1.0$, so $\gamma = -1.7$), is shown in Figure D.1. One notes that the high-energy bins have rather small errors (sub-Poissonian).

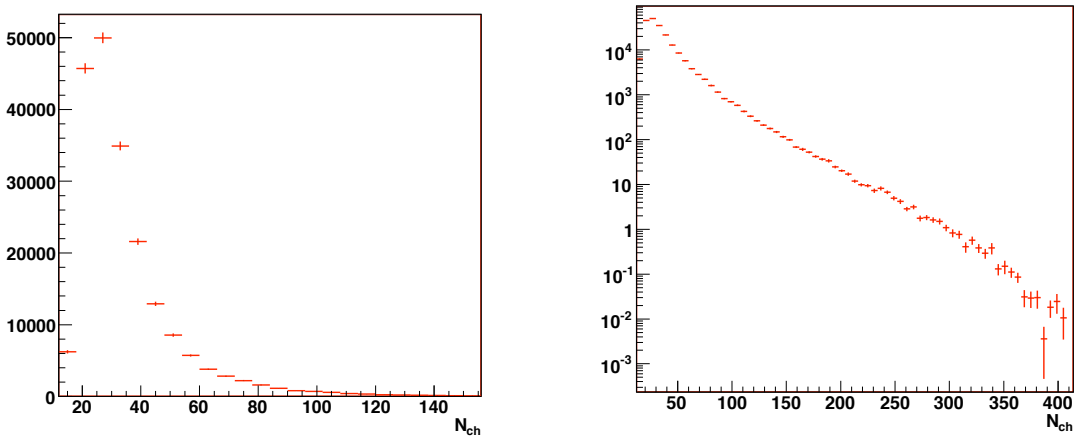

Figure D.1: Number of optical modules hit, from simulation of atmospheric muons with $\Delta = -1.0$.

In Figure D.2 one can see the effective weight $\tilde{w}$ calculated for each bin (as in eq. D.19), and also summing backward from the high-N_{ch} bin (as in eq. D.18). At low N_{ch}, the weight is significantly larger than 1, indicating that the statistics are worse than Poissonian, while at high N_{ch}, the situation is reversed. We note that because the energy, and thus the weights w_i, are correlated with N_{ch}, $\tilde{w}$ varies smoothly across the distribution. Thus we can fairly easily extrapolate to $\tilde{w}_0$ for the bin $(414 < N_{ch} < 420)$ — by eye, $\tilde{w}_0 \approx 0.006$, so

$$T_{414<N_{ch}<420} \approx 0^{+0.02}_{-0} \tag{D.22}$$

at the 90% confidence level. We note this error is quite reasonable given the values and errors of the final nonzero bins in figure D.1.

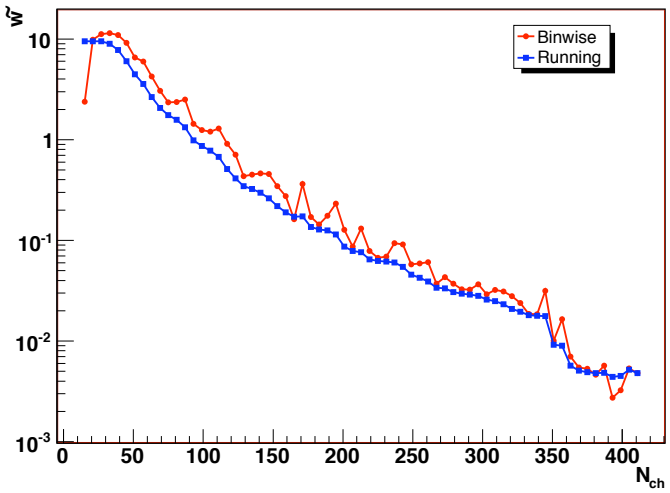

Figure D.2: The effective weight $\tilde{w}$ calculated both for each bin of the N_{ch} distribution as well as the sample running back from the final bin.

D.3.4 A Caveat

The procedure to define the error on the zero bin with constant event weight w is always well-defined (by D.17). It may be the case, however, that the sequence defined in eq. D.18 is not well-behaved. This can happen if the event weight w_i is not correlated with the observable chosen in the binning. In this case, it may not be possible to determine a limit or extrapolation of the $\tilde{w}_k$. One may at least, however, be able to estimate the order of magnitude of $\tilde{w}_0$.