

**A STUDY OF ATMOSPHERIC NEUTRINO OSCILLATION
USING THE ICECUBE DEEPCORE DETECTOR**

by

Laura Gladstone

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2014

Date of final oral examination: (3 October 2014)

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

Albrecht Karle (advisor), Department Chairperson and Professor, Physics

Francis Halzen, Gregory Breit Professor and Hilledale Professor, Physics

Vernon Barger, Vilas Professor and Van Vleck Professor, Physics

Justin Vandenbroucke, Assistant Professor, Physics

Dmitry Chirkin, Associate Scientist, Physics

© Copyright by Laura Gladstone 2014

All Rights Reserved

For my family

PRINCE HAL: *And Crispin Crispian shall ne'er go by,
From this day to the ending of the world,
But we in it shall be remembered—
We few, we happy few, we band of brothers;
For he to-day that sheds his blood with me
Shall be my brother; be he ne'er so vile,
This day shall gentle his condition;
And gentlemen in England now-a-bed
Shall think themselves accurs'd they were not here,
And hold their manhoods cheap whiles any speaks
That fought with us upon Saint Crispin's day.*

— WILLIAM SHAKESPEARE (1599)

ACKNOWLEDGMENTS

Sincere thanks to the community that got me through the long and arduous journey of this degree; I could never have done this alone, nor would I have wanted to try. Thanks especially to my advisor Albrecht Karle for the chance to study a subtle phenomenon at the edge of our detectable range, and for always encouraging me to pursue details.

Thanks to the local oscillations team! I was lucky to get extra feedback and advice from officemate-postdoc Melanie Day, upcoming student Moriah Tobin, Chris Wendt, Dima Chirkin, and Justin Vandenbroucke. Thanks to the whole group for countless suggestions, reality checks, practice talk comments, and making sure I didn't flounder or get sucked too far down any rabbit holes. Moriah, I hope I'm leaving this analysis in good enough shape that all your puzzles have to do with real physics.

Thanks to the professors who wrote letters of recommendation on my behalf: Albrecht, Reina Maruyama, Francis Halzen, and Darren Grant; and to the wonderfully supportive WIPAC professors Stefan Westerhoff and Mike Duvernois. Their regular advice, inspiration and leadership has guided me through the intricacies of working in a large collaboration, and helped me find the emotional stamina to persevere through this process.

Thanks to fellow students and partners-in-crime Christine Lewis, Valerie Plaus, Eli Parke, Karen Andeen, Nathan Whitehorn, Jakob van Santen, Marcos Santander, Antonia Hubbard, and Jake Feintzeig. You have made my time at Wisconsin the crazy fun exploration it was.

Thanks for several chapters of editing and proofreading from the DM-Ice Dissertators Peer Editing Club: Antonia Hubbard, Benedikt Riedel, Walter Pettus, and Zachary Pierpoint. Thanks to Jason Loch for final editing, plus consistent writing advice and encouragement over the last year.

A postdoc with an open door and an interesting story to answer to any question, no matter how silly sounding, is a lifeline for a graduate student. Thanks to David Boersma and Jason Koskinen for teaching me how to program, and how to follow up on the "huh, that's weird..." moments that so often herald discovery.

The "postdoc" lunch crowd (even though few of us were technically postdocs) was a wonderful professional mentorship group in a casual setting: thanks to Paolo Desiati, Juan Carlos Díaz Vélez, John Kelley, Segev BenZvi, Naoko Kurahashi Neilson, Markus Ahlers, Claudio Kopper, the Osteria Papavero restaurant, and the Ancora and Barrique's coffee shops.

In a very literal sense, I'm in physics today because of Janet Conrad's mentorship and encouragement through college and grad school. Thank you to her and her network, especially Lindley Winslow, Georgia Karagiorgi, Teppei Katori, and Ben Jones. And remember, practice talks are good for everyone, for every talk.

The moment when I decided to pursue physics research as a career was at the end of a night shift monitoring the MiniBooNE detector during my first of three summers at Fermilab. I was wrestling with the question of whether I had a place in the research world (and staring out at the sunrise) when Janet walked in and said "we should talk about thesis topics for you." Not only was this postulating that I belonged in the neutrino group, it gave me hope for a specific direction and guidance. I thought "I could do this forever," and I've never looked back.

Most of all, this would not have been possible without my loving and hardworking parents and sister—Igor, Becky, and Anna—for persistent support and inspiration and for teaching me to love puzzles.

DISCARD THIS PAGE

TABLE OF CONTENTS

	Page
LIST OF TABLES	vii
LIST OF FIGURES	viii
ABSTRACT	xiii
1 The History and Theory of Neutrino Oscillations	1
1.1 Discovery of Neutrinos	1
1.2 The Solar Neutrino Problem	4
1.3 Neutrino Oscillations	5
1.4 Approximations from Three to Two Flavor Oscillations	9
1.5 Two Flavor Oscillations	9
1.6 Oscillations in Matter	11
1.7 Summary of This Work	13
2 Neutrino Detectors	14
2.1 Interactions and Cross Sections	14
2.2 A Good Oscillation Experiment	17
2.3 Atmospheric Neutrinos	19
2.4 Muon Energy Losses in Matter	23
2.5 Cherenkov Light	25
2.6 Current Neutrino Detectors for Measuring Atmospheric Mixing	28
2.6.1 The Super-Kamiokande Experiment	28
2.6.2 The Tokai to Kamioka Experiment	29
2.6.3 The MINOS Experiment	30
3 The IceCube DeepCore Detector	32
3.1 IceCube Subdetectors	32
3.2 IceCube and DeepCore Geometry	33
3.3 Modeling the Ice Properties	39
3.3.1 Gathering Ice Data	39
3.3.2 Creating an Ice Model	40
3.4 Digital Optical Modules	42
3.4.1 Digitizing the Signal: ATWD and fADC	42
3.5 Triggers and Filters	44

	Page
4 DOM Testing	46
4.1 Production Tests	46
4.2 South Pole Testing	49
4.3 Specialty Testing	49
4.3.1 Dependence of Noise Rates on Temperature and Glass Presence	49
4.3.2 The Vuvuzela Noise Model	54
4.3.3 Internal DOM Temperature Tests	54
5 Reconstruction Terminology and Methods	58
5.1 Basic Terminology	58
5.2 “I3Particle” Definition	59
5.3 Hit Cleaning: SRT and TW	60
5.4 Topological Trigger	61
5.5 Noise Engine	61
5.6 Linefit and Improved Linefit	63
5.7 CLast	63
5.8 Describing Ice with Photonics Tables	64
5.9 The Gulliver Software Suite	66
5.10 SPE and MPE	67
5.11 The Millipede Software Library	68
5.11.1 Monopod	70
5.11.2 MuMillipede Fit	70
5.11.3 Igelfit	70
5.11.4 Other Millipede Configurations	72
6 The Shadow of the Moon in Cosmic Ray Muons	73
6.1 The Moon Filter	73
6.2 The IC22 Binned Analysis	75
6.3 The IC40 Binned Analysis	77
6.4 The IC40 and IC59 Unbinned Analyses	78
6.5 Future Studies	79
7 Event Selection	84
7.1 Limiting Bias Using Blind Analysis	85
7.2 Monte Carlo Simulation	85
7.3 The Online Filter	86
7.4 Preliminary Cut Levels	89
7.4.1 Cut Level 3: First Fits	89
7.4.2 Cut Level 4: NCh NVeto	91
7.4.3 Cut Level 5: TTrigger Splitter	91
7.4.4 Cut Level 6: Improved Linefit and SPE32	92
7.4.5 Cut Level 7: Causality-based veto	92
7.5 Newly Optimized Cuts	92
7.5.1 Cut Levels 8: Monopod	92

Appendix

	Page
7.5.2 Cut Level 9: Igelfit	94
7.5.3 Cut Level 10:	94
7.6 Event Selection Summary	95
8 Analysis Method	103
8.1 Maximum Likelihood Technique	103
8.2 Nuisance Parameters: Physics and Systematics	106
8.3 Frequentist Confidence Region	107
8.4 Feldman-Cousins Error Contours	109
8.5 Software Specifics	111
9 Round Trip Simulation Trials	115
9.1 Input Oscillation Parameters	115
9.2 Statistical Increase	116
9.3 Placement of the L10 Cuts	116
9.3.1 Zenith	116
9.3.2 Energy	121
9.4 Variations in Cosmic Ray μ Rates	121
9.5 Variations in ν_τ Rates	128
9.6 Comparing Wilks Theorem and Feldman-Cousins Statistics	128
10 Results and Conclusions	133
10.1 Single Dimension Results	134
10.2 Likelihood Scan Results with Variations in the Cosmic Ray Muon Rates	134
10.3 Specific Ideas for Future Improvements	140
10.3.1 Additional Cosmic Ray Muon Simulation	141
10.3.2 Spline Tables	142
10.3.3 Noise Modeling	142
10.3.4 Updating the Genie Simulation Generator	143
10.3.5 Particle Reconstructions	143
10.4 Conclusion	144
LIST OF REFERENCES	145
APPENDICES	
Appendix A: Supplemental Plots for Level 8 Cuts	153
Appendix B: Supplemental Plots for Level 9 Cuts	166
Appendix C: Supplemental Plots for Level 10 Cuts	184
Appendix D: The IceCube Moon Shadow Paper	201

DISCARD THIS PAGE

LIST OF TABLES

Table	Page
1.1 Current values of oscillation parameters	7
3.1 Comparison of the various configurations of IceCube detector	39
3.2 Digitizer readout description	42
4.1 Temperatures at which temperature probes were tested	55
7.1 Event rates passing each cut level for each type of data	102
7.2 Number of events passing each cut level for each type of data	102
8.1 Conversions from probability to significance	109

DISCARD THIS PAGE

LIST OF FIGURES

Figure	Page
1.1 A beta decay energy spectrum cartoon	2
1.2 A detailed beta decay energy spectrum prediction endpoint plot	3
1.3 The Kamland oscillation result	6
1.4 Oscillation probabilities as a function of energy, with length fixed at the diameter of the Earth	9
2.1 Feynman diagrams of the interactions through which we observe neutrinos	15
2.2 Neutrino $\bar{\nu}_e e^- \rightarrow \bar{\nu}_e e^-$ charged current scattering cross section as a function of energy, with representative pictures of sources at several energies	16
2.3 Neutrino and anti-neutrino cross sections as a function of energy	18
2.4 Oscillation probabilities for atmospheric neutrinos visible to IceCube.	20
2.5 Fractional contributions to cosmic ray composition as a function of energy	21
2.6 Observations of cosmic rays over many orders of magnitude in energy	22
2.7 The energy-dependent fraction of ν_μ and μ produced in the atmosphere from pions and kaons	23
2.8 Predicted and measured atmospheric neutrino fluxes as a function of energy	24
2.9 Muon energy losses as a function of energy	26
2.10 The Cherenkov angle and wavefront relative to the particle track	27
2.11 DOM light acceptance as a function of light wavelength	28
2.12 Atmospheric oscillation limits from Super-K	29
2.13 Atmospheric oscillation limits from T2K	30
2.14 MINOS ν_μ disappearance oscillation measurement using beam and atmospheric data	31
2.15 First MINOS+ results: ν_μ survival probability curve	31
3.1 The detector layout from the side as if through the ice sheet	34
3.2 The detector layout from above with distances between strings	35

Figure	Page
3.3 The detector layout from above with distances between strings	36
3.4 A diagram of the layout of the DeepCore detector	37
3.5 The absorption and scattering parameters in the SPICE-MIE ice model	41
3.6 DOM diagram and photo	43
4.1 Photos of DOMs set up for production tests	47
4.2 Photo of the South Pole DOM testing site	50
4.3 Average noise rates as a function of temperature for high quantum efficiency PMTs within DOMs	51
4.4 Dark noise rates as a function of temperature comparing standard versus high quantum efficiency PMTs and DOMs versus bare PMTs.	52
4.5 Dark noise rate tests on non-IceCube PMTs	53
4.6 Photograph of the temperature probe locations within the DOM.	56
4.7 Temperature dependence of the temperature measurement deviations between various points on a DOM compared to the ambient freezer temperature	57
5.1 The rates of reconstructed pulses, with various cleanings applied, versus the “time residual”	62
5.2 A multi-panel plot showing how we combine information from many PMT waveforms to reconstruct the direction of cascade-type events	71
5.3 A hedgehog, or “igel”	72
6.1 Moon filter rates as a function of time, over two one-month cycles	74
6.2 Moon filter rates versus reconstructed azimuth, showing a peaked structure	75
6.3 IC22 Moon Filter sample after cuts, plotted in right ascension and declination relative to the Moon	76
6.4 The IC22 declination bands relative to the Moon	80
6.5 The IC22 Moon shadow final histogram	80
6.6 IC40 Moon shadow histogram of event counts	81
6.7 IC40 Moon shadow deviations in the region around the Moon	81
6.8 Gaussian fit to the significances of deviations in the region around the IC40 Moon	82
6.9 An alternate Moon shadow search using annular bins around the Moon position	83
7.1 The Genie-NuGen cross-weighting method	87

Appendix Figure	Page
7.2 NCh versus rate at cut level 10, showing all simulation components	88
7.3 Particle speeds reconstructed by the DeepCore filter algorithm	90
7.4 Monopod reconstructed depth versus energy resolution	93
7.5 Event rate versus zenith angle reconstruction error for Igelfit and Monopod.	94
7.6 Event rate versus energy reconstruction error for Igelfit and Monopod.	96
7.7 NCh at all cut levels	97
7.8 Monte Carlo true energy at all cut levels	98
7.9 Monte Carlo true zenith angle at all cut levels	99
7.10 Reconstructed energy at all available cut levels	100
7.11 Reconstructed zenith angle at all available cut levels	101
8.1 Component histograms of the simulated rates prediction	105
8.2 How to optimize the step size of a likelihood minimizer	108
8.3 Normal distribution with significances labeled	110
8.4 Example histogram of pseudodata trials outcomes for Feldman-Cousins contour construction	112
8.5 Example pseudocode explaining the Feldman Cousins statistical method	113
9.1 Round trip trials varying the input oscillation parameters	117
9.2 Results of two round-trip trials comparing long (7500.71 hours) and short (1566.39 hours) livetime.	118
9.3 Results of round-trip trials comparing three different binning configurations for zenith, part 1 of 2	119
9.4 Results of round-trip trials comparing three different binning configurations for zenith, part 2 of 2	120
9.5 Results of round-trip trials comparing three different binning configuration for energy, part 1 of 2.	122
9.6 Results of round-trip trials comparing three different binning configuration for energy, part 2 of 2.	123
9.7 Round trip trials varying the rate of cosmic ray muon contamination.	124
9.8 Round trip trials varying the rate of cosmic ray muon contamination.	125
9.9 Round trip trials comparing smoothed and unsmoothed CORSIKA distributions	127
9.10 Round trip trials varying the rate of ν_τ contamination.	129
9.11 Round trip trials varying the rate of ν_τ contamination.	130

Appendix	Page
Figure	
9.12 Results of round-trip trials using the same setup to compare Wilks' theorem with Feldman-Cousins as methods for placing significance contours	132
10.1 Rate versus reconstructed energy, comparing the full data to simulation predictions in both the expected oscillation and no oscillations cases.	135
10.2 Rate versus reconstructed zenith angle, comparing the full data to simulation predictions in both the expected oscillation and no oscillations cases.	136
10.3 Results assuming that the cosmic ray muon background is distributed as in the simulation.	137
10.4 Results assuming that the cosmic ray muon background is distributed as in a hand-smoothed shape.	138
10.5 Results assuming that all events not accounted for by the simulated neutrino rates are cosmic ray muon background.	139
A.1 Cut Level 7: NCh, after hit cleaning	154
A.2 Cut Level 7: NString, after hit cleaning	155
A.3 Cut Level 7: Total charge, after hit cleaning	156
A.4 Cut Level 7: CLast reconstructed vertex time	157
A.5 Cut Level 7: Distance between the central string and center-of-gravity reconstructed vertex, squared	158
A.6 Cut Level 7: Monopod reconstructed depth	159
A.7 Cut Level 7: Distance between the central string and the Monopod reconstructed vertex, squared	160
A.8 Cut Level 7: Number of direct hits relative to Monopod	161
A.9 Cut Level 7: Number of early pulses relative to the Monopod reconstructed interaction vertex	162
A.10 Cut Level 7: Number of late pulses relative to the Monopod reconstructed interaction vertex	163
A.11 Cut Level 7: Monopod reconstructed energy	164
A.12 Cut Level 7: Monopod Reconstructed Zenith	165
B.1 Cut Level 8: NCh, after hit cleaning	167
B.2 Cut Level 8: NString, after hit cleaning	168
B.3 Cut Level 8: Total charge, after hit cleaning	169
B.4 Cut Level 8: CLast reconstructed vertex time	170
B.5 Cut Level 8: Distance between the central string and center-of-gravity reconstructed vertex, squared	171
B.6 Cut Level 8: Monopod reconstructed depth	172

Appendix Figure	Page
B.7 Cut Level 8: Distance between the central string and the Monopod reconstructed vertex, squared	173
B.8 Cut Level 8: Number of direct hits relative to Monopod	174
B.9 Cut Level 8: Number of early pulses relative to the Monopod reconstructed interaction vertex	175
B.10 Cut Level 8: Number of late pulses relative to the Monopod reconstructed interaction vertex	176
B.11 Cut Level 8: Monopod reconstructed energy	177
B.12 Cut Level 8: Monopod reconstructed zenith	178
B.13 Cut Level 9: Number of direct hits relative to Igelfit	180
B.14 Cut Level 9: Number of early pulses relative to the Igelfit reconstructed track	181
B.15 Cut Level 9: Number of late pulses relative to the Igelfit reconstructed track	182
B.16 Cut Level 9: Igelfit reconstructed zenith	183
C.1 Cut Level 10: NCh, after hit cleaning	185
C.2 Cut Level 10: NString, after hit cleaning	186
C.3 Cut Level 10: Total charge, after hit cleaning	187
C.4 Cut Level 10: CLast reconstructed vertex time	188
C.5 Cut Level 10: Distance between the central string and center-of-gravity reconstructed vertex, squared	189
C.6 Cut Level 10: Monopod reconstructed depth	190
C.7 Cut Level 10: Distance between the central string and the Monopod reconstructed vertex, squared	191
C.8 Cut Level 10: Number of direct hits relative to Monopod	192
C.9 Cut Level 10: Number of early pulses relative to the Monopod reconstructed interaction vertex	193
C.10 Cut Level 10: Number of late pulses relative to the Monopod reconstructed interaction vertex	194
C.11 Cut Level 10: Monopod reconstructed energy	195
C.12 Cut Level 10: Monopod reconstructed zenith	196
C.13 Cut Level 10: Number of direct hits relative to Igelfit	197
C.14 Cut Level 10: Number of early pulses relative to the Igelfit reconstructed track	198
C.15 Cut Level 10: Number of late pulses relative to the Igelfit reconstructed track	199
C.16 Cut Level 10: Igelfit reconstructed zenith	200

A STUDY OF ATMOSPHERIC NEUTRINO OSCILLATION USING THE ICECUBE DEEPCORE DETECTOR

Laura Gladstone

Under the supervision of Professor Albrecht Karle

At the University of Wisconsin-Madison

The IceCube detector at the South Pole is a cubic-kilometer-scale neutrino detector designed to observe TeV-range charged particle secondaries from neutrino interactions, and thus do neutrino astronomy. As a main background to the search for astrophysical point sources of neutrinos, IceCube also observes muons and neutrinos from the atmospheric interactions of cosmic rays. By observing a spectrum of atmospheric neutrinos and comparing it to independent predictions of atmospheric fluxes with and without oscillations, IceCube can test various values of oscillation parameters.

Neutrino oscillations have been observed experimentally for several decades; IceCube is the first experiment to extend this measurement to the 10-20 GeV energy range. An initial analysis has established that IceCube can see oscillations using the 79-string detector configuration (“IC79”), which was the first year of data that included the DeepCore detector. As a follow-up, this analysis uses a less restrictive event selection and thus a higher total event count, around 3,000 for one year of data. The fit is a Poisson likelihood fit of a two-dimensional rate histogram, using both oscillated length and observed energy. The arrival zenith angle of the muon is used as a proxy for oscillation length. The error contours are dominated by systematic effects more than by statistical limitations of the data. Major systematics include uncertainties in the atmospheric neutrino flux at high energies and uncertainties in the distribution of the cosmic ray muon background.

This analysis was designed to produce limits on the mixing angle θ_{23} that are competitive with other current experiments, although this is still uncertain as error analysis work is ongoing. Future work will further refine the event selection and systematic error analysis; the statistical methods and software used here are expected to become the IceCube oscillations standard.

This thesis also includes background information about the history of neutrino physics, basics of oscillation formalism, and the IceCube DeepCore detector. A brief chapter is included on the search for a deficit in cosmic ray muons from the direction of the Moon, and the implications of that deficit for the pointing resolution of IceCube.

Albrecht Karle

ABSTRACT

The IceCube detector at the South Pole is a cubic-kilometer-scale neutrino detector designed to observe TeV-range charged particle secondaries from neutrino interactions, and thus do neutrino astronomy. As a main background to the search for astrophysical point sources of neutrinos, IceCube also observes muons and neutrinos from the atmospheric interactions of cosmic rays. By observing a spectrum of atmospheric neutrinos and comparing it to independent predictions of atmospheric fluxes with and without oscillations, IceCube can test various values of oscillation parameters.

Neutrino oscillations have been observed experimentally for several decades; IceCube is the first experiment to extend this measurement to the 10-20 GeV energy range. An initial analysis has established that IceCube can see oscillations using the 79-string detector configuration (“IC79”), which was the first year of data that included the DeepCore detector. As a follow-up, this analysis uses a less restrictive event selection and thus a higher total event count, around 3,000 for one year of data. The fit is a Poisson likelihood fit of a two-dimensional rate histogram, using both oscillated length and observed energy. The arrival zenith angle of the muon is used as a proxy for oscillation length. The error contours are dominated by systematic effects more than by statistical limitations of the data. Major systematics include uncertainties in the atmospheric neutrino flux at high energies and uncertainties in the distribution of the cosmic ray muon background.

This analysis was designed to produce limits on the mixing angle θ_{23} that are competitive with other current experiments, although this is still uncertain as error analysis work is ongoing. Future work will further refine the event selection and systematic error analysis; the statistical methods and software used here are expected to become the IceCube oscillations standard.

This thesis also includes background information about the history of neutrino physics, basics of oscillation formalism, and the IceCube DeepCore detector. A brief chapter is included on the search for a deficit in cosmic ray muons from the direction of the Moon, and the implications of that deficit for the pointing resolution of IceCube.

Chapter 1

The History and Theory of Neutrino Oscillations

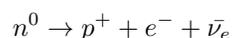
*When you have eliminated the impossible, whatever remains,
however improbable, must be the truth.*

— SHERLOCK HOLMES, 1890 [1]

1.1 Discovery of Neutrinos

In 1930, the conservation of energy, one of the best-tested precepts of modern science, was facing the possibility of defeat by counterexample. Observations of radioactive beta decay energy spectra did not seem to conserve energy the way they should, and it was here that neutrinos first appeared on the scene: to save energy conservation. The only remaining possibility was a new invisible particle, which seemed unlikely; thus began the broad and active research field of neutrino physics.

Nuclear beta decay occurs when a nucleus emits an electron (originally called a beta particle) while changing a neutron to a proton:



Charge is conserved in this reaction whether or not the neutrino is included: the original neutron is neutral, and the final sum of proton-plus-electron is neutral. The question is whether energy conservation forces us to include the third particle, the neutrino. The conservation of the sum of mass and energy can be similarly investigated by looking at the sum of the final products: the rest mass plus kinetic energy of the proton and electron. Assume for a moment that these are the only two product particles: this will lead to a problem that the neutrino can solve. The total mass is a relativistically invariant quantity, so consider (for the sake of simplicity) the frame in which the neutron is at rest: its total energy is just from its rest mass, a definite value, not a spectrum. The sum of the energy of the resulting proton and electron should add up to this definite mass. But the observed energy of the outgoing proton and electron was a spectrum between 0 and the neutron mass. The experiments of Lise Meitner and Otto Hahn in 1911 found this spectrum instead of a line, which is illustrated in Fig. 1.1.

This contradiction can be explained by positing an extra particle in the interaction that can carry away a random fraction of the energy but remain undetected. The upper limit of the curve is very close to the neutron mass, implying that the new particle has at most a minuscule mass of its own. Such a particle explains the spectrum but seems

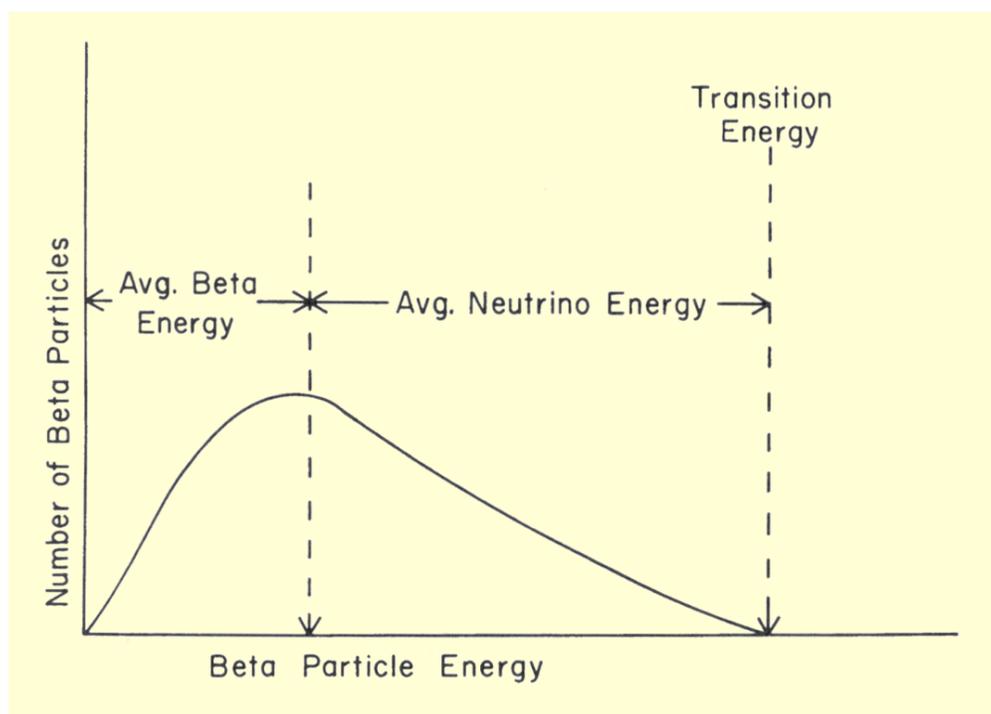


Figure 1.1: A beta decay energy spectrum cartoon. The x-axis shows the measured energy of the emitted beta particle, and the y-axis shows the frequency of that detected energy. The fact that this is a spectrum instead of a single line at the transition energy implies a 3, rather than 2, body decay.

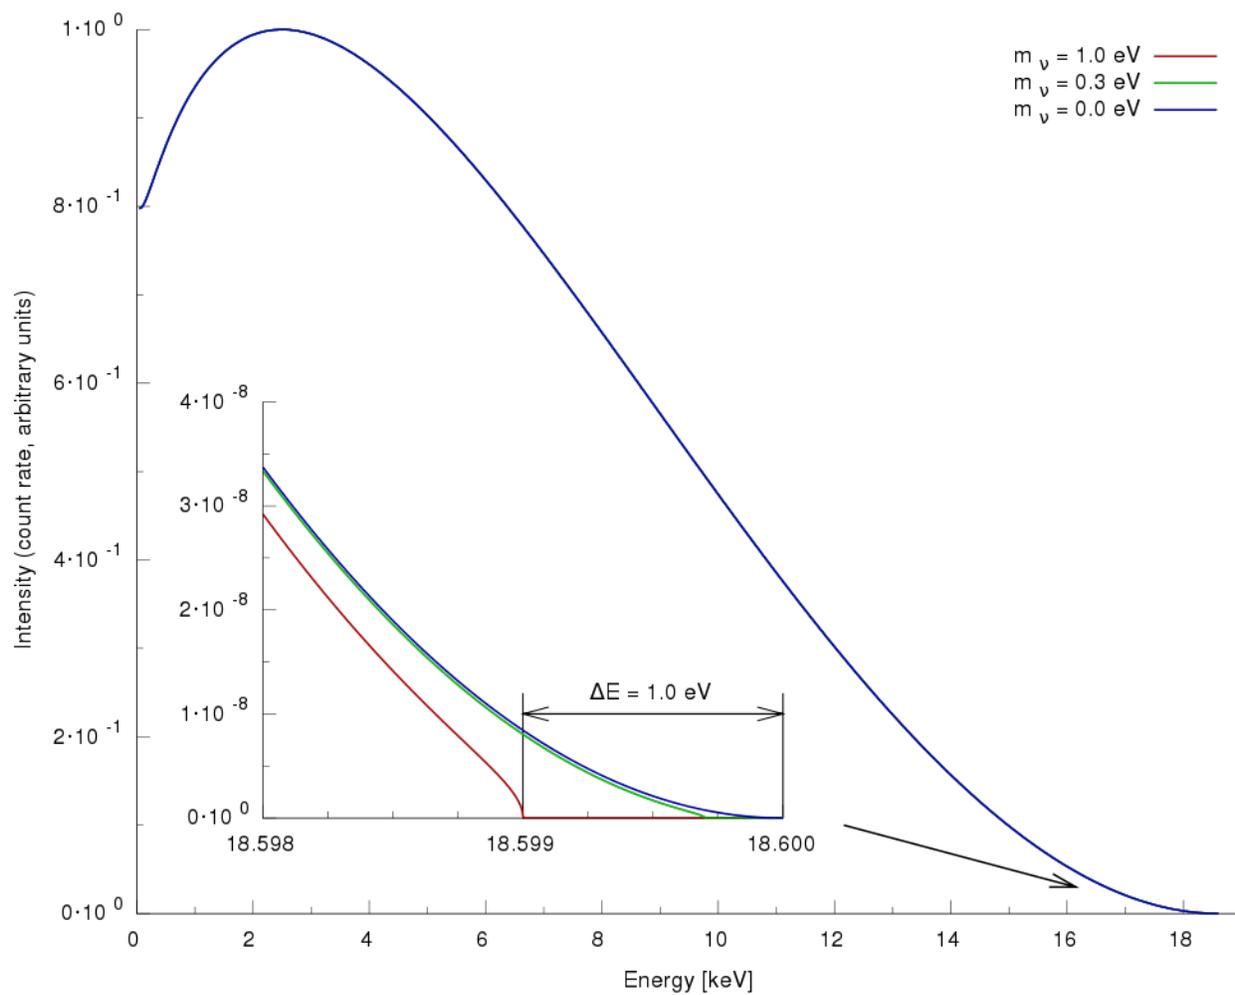
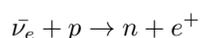


Figure 1.2: A detailed beta decay energy spectrum prediction endpoint plot. The inset shows only the endpoint, focusing on the difference between the zero-mass neutrino case (far right) and the the massive neutrino case, using known limits on neutrino mass. [2]

unattractive because of the requirement that it remain undetected. How could a particle change the energy of a reaction so dramatically and yet avoid detection?

The first written record of someone proposing such a particle to explain the beta decay spectrum was a letter from Wolfgang Pauli [3] to a conference, in lieu of speaking at the conference. He famously addressed the letter “Dear Radioactive Ladies and Gentlemen,” and apologized for choosing a social engagement over the conference: “since I am indispensable here in Zurich because of a ball on the night from December 6 to 7”. He recognizes the audacity of proposing an invisible particle, but proposes it anyway.¹ This idea was later formalized by Enrico Fermi, [5][6] but it was considered too speculative for English-language journals for several years. The particle name that Pauli proposed was “neutron”, which means “neutral one” in Italian. However, only a few years later a different nuclear particle was discovered and named the “neutron” by James Chadwick [7] (who won the Nobel Prize in 1935 for the discovery), so the name of the smaller particle was changed from “neutron” to “neutrino”, which means “little neutral one”.

Neutrinos were first observed in 1956 by Clyde L. Cowan and Frederick Reines [8]. They used a scintillator detector next to a nuclear reactor to look for neutrinos inducing inverse beta decay in chlorine:



They identified a neutrino interaction by looking for a two-part signal. First, a flash from the positron annihilating with an electron to create two 511 keV gamma rays. Second, a gamma ray from the neutron being absorbed by cadmium. The timing of the neutrino absorption could be confirmation of a neutrino event: 5 μ sec is the neutron lifetime in cadmium.

Their experiment was originally located at Hanford, Washington, but later moved to Savannah River near Aiken, South Carolina, because there was better shielding available from cosmic rays: this experimental hall was located 12 m underground. The experiment included two tanks, with ~ 200 L of water, and ~ 40 kg of dissolved CdCl_2 . Each tank was lined with 5 inch photomultiplier tubes.

They identified ~ 3 neutrino events per hour. This was within their expectations: the predicted cross section was $6 \times 10^{-44} \text{cm}^2$, and their measured rate translated to $6.3 \times 10^{-44} \text{cm}^2$.

This experiment was accepted as conclusively observing, and thus discovering, the neutrino. Clyde Cowan died in 1974, but Frederick Reines received the Nobel Prize in 1995 “for the detection of the neutrino” [9].

1.2 The Solar Neutrino Problem

After the existence of neutrinos had been confirmed by experiment, measurements continued in increasing detail. A discrepancy arose [10][11][12] between the number of neutrinos predicted to arrive from the Sun and the number observed; this is known as the Solar Neutrino Problem.

¹As an interesting historical note, the letter was addressed specifically to Lise Meitner and Hans Geiger [4] because of their beta decay expertise. This beta decay work makes Lise Meitner (not only the first woman in Germany to be a full professor of physics, but also) the woman most closely tied to the early history of neutrinos.

A resolution was proposed: including neutrino oscillations in the predicted rates [13], which could then explain the data. Another proposed solution (which has since been ruled out by more precise measurements of several solar processes in favor of oscillations) was to change the model of solar interactions. However, any changes to this model decreased the overall goodness of the fit to the data [14].

An early piece of evidence [15] for neutrino oscillations came from the Super-Kamiokande experiment looking at atmospheric neutrinos. This showed that the number of neutrinos depended on the arrival direction, that is, the number of neutrinos of a given energy changed depending on how far they had traveled. This was consistent with neutrino oscillations, but did not directly prove that neutrinos from the Sun were oscillating.

The definitive evidence [16] came from the SNO experiment, which had fine enough resolution to distinguish between neutral current events and charged current electron flavor events. SNO used heavy water: water with deuterium in the place of hydrogen. In neutral current interactions, deuterium atoms were disassociated. The resulting free neutron captured onto a different nucleus, emitting a characteristic ~ 6 MeV gamma ray. By contrast, a charged current electron neutrino event produced an electron of $\sim 5 - 15$ MeV. The charged current contributions from other flavors were negligible because the energy of solar neutrinos is below the mass of the muon or tau charged leptons. The neutral current even rate fit with the expected rate before oscillations, and the electron flavor charged current event rate fit with the expected rate after oscillations. This evidence shifted the scientific consensus to favor oscillations.

Further evidence came from the KamLAND reactor experiment, first with a disappearance consistent with oscillations [17]. This was followed by an analysis [18] precise enough that the oscillation effect is visible in a plot of the reconstructed ratio of length the neutrino traveled to neutrino energy: L/E ; see Fig. 1.3.

1.3 Neutrino Oscillations

Neutrino oscillations have been observed in several contexts, one of which was the Solar Neutrino Problem from the last section. Several more are discussed in the next chapter. This section introduces the mathematical formalism behind these observations, including the relevant parameters for measuring oscillations. It ends with setting up the oscillation probability amplitude. Full oscillation probability formulas are discussed in the following sections of this chapter.

Mathematically, neutrino oscillations can be understood as a result of the difference between the weak-interaction flavor eigenstates and mass eigenstates. It follows that a flavor eigenstate (indexed by letter) can be written as a linear superposition of mass eigenstates (indexed by number):

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle.$$

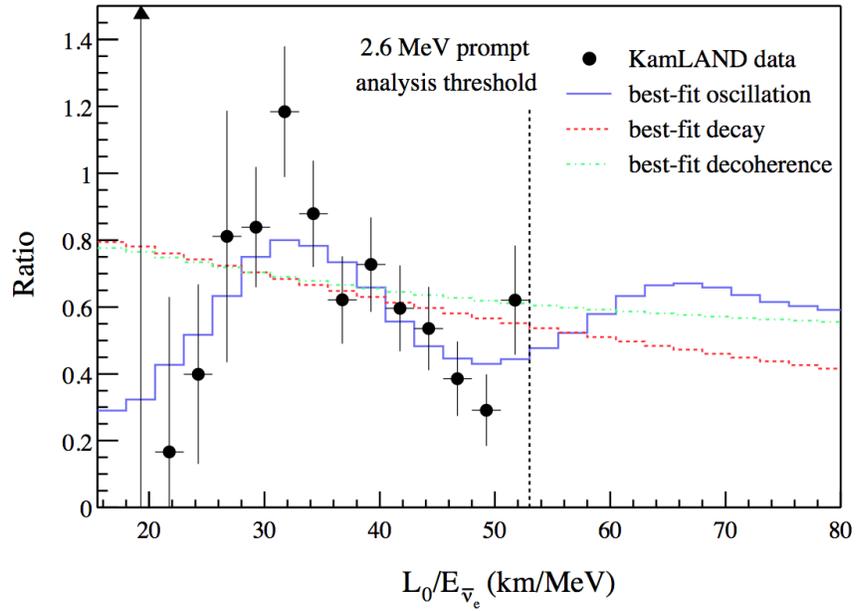


Figure 1.3: The Kamland oscillation result, from [18]: Ratio of the observed ν_e spectrum to the expectation for no-oscillation versus L_0/E . The curves show the expectation for the best-fit oscillation, best-fit decay and best-fit decoherence models taking into account the individual time-dependent flux variations of all reactors and detector effects. The data points and models are plotted with $L_0 = 180\text{km}$, as if all anti-neutrinos detected in KamLAND were due to a single reactor at this distance.

Expanding this notation, we can write the full matrix that transforms between the flavor basis and the mass basis:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}.$$

This was first proposed theoretically in the 1960s by Maki, Nakagawa, Sakata [19], and Pontecorvo [20], and thus the key matrix (the $U_{\alpha i}$ matrix) in this theory is often called MNSP (or PMNS) matrix. This three-dimensional rotation matrix can be factorized into three two-dimensional rotations. To fit this within the margins, we adopt the notational condensations that $\sin(\theta) \equiv s\theta$ and $\cos(\theta) \equiv c\theta$.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c\theta_{23} & s\theta_{23} \\ 0 & -s\theta_{23} & c\theta_{23} \end{pmatrix} \begin{pmatrix} c\theta_{13} & 0 & s\theta_{13}e^{i\delta} \\ 0 & 1 & 0 \\ -s\theta_{13} & 0 & c\theta_{13} \end{pmatrix} \begin{pmatrix} c\theta_{12} & s\theta_{12} & 0 \\ -s\theta_{12} & c\theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}.$$

The three mixing angles θ_{23} , θ_{13} , and θ_{12} , together with the CP-violating phase δ are the parameters of the theory that can be experimentally constrained. This factorization corresponds conveniently to the various types of experiments that have observed oscillations, and we name the angle of each of these rotations accordingly. The ‘‘atmospheric’’ mixing angle θ_{23} is observed in atmospheric experiments like IceCube and SuperK, and long-baseline accelerator experiments like MINOS; ‘‘reactor mixing’’ θ_{13} is observed in reactor experiments and short baseline accelerator experiments; and ‘‘solar mixing’’ θ_{12} in solar neutrino experiments like SNO. The phase δ has not yet been constrained but is a major upcoming question for the field. A summary of the current values of these parameters is shown in Table 1.1.

Parameter	Best Fit $\pm 1\sigma$
$\sin^2(\theta_{23})$	$0.386^{+0.024}_{-0.021}$
$\sin^2(\theta_{12})$	$0.307^{+0.018}_{-0.016}$
$\sin^2(\theta_{13})$	0.0241 ± 0.0025
Δm_{23}^2	$2.43^{+0.06}_{-0.10} \times 10^{-3} \text{eV}^2$
Δm_{12}^2	$7.54^{+0.26}_{-0.22} \times 10^{-5} \text{eV}^2$

Table 1.1: Current world best fit values and constraints [21] on neutrino oscillation parameters: mixing angles and mass differences. CP-violating phase, majorana mass, and all sterile neutrino parameters are neglected.

Oscillation experiments try to observe neutrinos that start in one flavor, ν_α , and finish in flavor ν_β , after traveling some baseline distance L . We talk about this as the ‘‘oscillation probability’’ $P(\nu_\alpha \rightarrow \nu_\beta)$; if the start and end flavors are the same, this is a ‘‘survival probability’’.

The notation convention used here is that numbers refer to mass states (1,2,3, or generally i or j) and letters refer to flavor states (e, ν, τ , or generally α or β).

The initial state of a neutrino is taken as the neutrino at time $t = 0$ with flavor transformation:

$$\langle \nu_\alpha | = \langle \nu_i | U_{\alpha i}^* \quad (1.1)$$

$$| \nu_\beta \rangle = U_{\beta j} | \nu_j \rangle \quad (1.2)$$

This state evolves in time with the standard energy-dependent propagator:

$$| \nu_\beta(t) \rangle = U_{\beta j} | \nu_j(t) \rangle \quad (1.3)$$

$$= e^{-iE_j t} U_{\beta j} | \nu_j(t=0) \rangle. \quad (1.4)$$

To find the general time-dependent oscillation probability amplitude, we take the product of these two states:

$$A(\nu_\alpha \rightarrow \nu_\beta | t) = \langle \nu_\alpha | \nu_\beta(t) \rangle \quad (1.5)$$

$$= \langle \nu_i | U_{\alpha i}^* e^{-iE_j t} U_{\beta j} | \nu_j \rangle \quad (1.6)$$

$$= U_{\alpha i}^* e^{-iE_j t} U_{\beta j} \langle \nu_i | \nu_j \rangle \quad (1.7)$$

$$= U_{\alpha i}^* e^{-iE_j t} U_{\beta j} \delta_{ij} \quad (1.8)$$

$$= U_{\alpha j}^* e^{-iE_j t} U_{\beta j}. \quad (1.9)$$

The δ_{ij} here is the Kronecker delta, which arises from the fact that the $| \nu_j \rangle$ s and $\langle \nu_i |$ s are both parts of an orthonormal basis. Without loss of generality, we chose to keep the j index instead of the i index, to avoid confusion with the imaginary integer i . We sum over the repeated j index.

Having found the probability amplitude, we get the probability by taking the square: amplitude times amplitude-conjugate.

$$P(\nu_\alpha \rightarrow \nu_\beta | t) = |A(\nu_\alpha \rightarrow \nu_\beta | t)|^2 \quad (1.10)$$

$$= |U_{\alpha j}^* e^{-iE_j t} U_{\beta j}|^2 \quad (1.11)$$

Keeping track of all the entries of PMNS matrix, this expression quickly becomes cumbersome to calculate.

The probability of three neutrino mixing in a vacuum is given by the following expression, in which flavors are indexed by Greek letters (α, β) and masses are indexed by Latin letters (i, j).

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2(\Delta m_{ij}^2 \frac{L}{4E}) + 2 \sum_{i>j} \Im(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2(\Delta m_{ij}^2 \frac{L}{2E})$$

In this expression, the second term includes a sum over the real parts of the matrix product; selecting the real part is denoted by \Re . Similarly, the third term includes a sum over the imaginary parts of the matrix product, denoted by \Im . For the case of muon and tau neutrinos traveling the diameter of the Earth, this expression is plotted in Fig. 1.4.

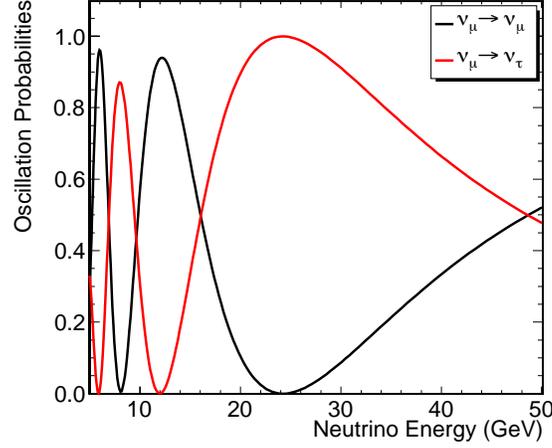


Figure 1.4: Oscillation Probabilities as a function of energy, with length fixed at the diameter of the Earth (12,742 km). This probability includes contributions from all 3 neutrino flavors.

1.4 Approximations from Three to Two Flavor Oscillations

There are two approximations needed to simplify the math from the three neutrino case to the more tractable two neutrino case. The first is that the oscillations are driven by a single mass splitting. This happens when the other mass splitting is much smaller than the energies involved:

$$\frac{\Delta m_{12}^2 L}{E} \ll 1$$

(It should be clear at this point that L/E has dimensions of $1/\text{mass}^2$). We can then approximate $\Delta m_{12}^2 \approx 0$.

The second approximation is that the other mixing angles are negligible: specifically, that

$$\sin \theta_{13} \approx 0,$$

or equivalently that the matrix element $|U_{e3}| \approx 0$. In the major neutrino result of 2012, reactor experiments [22][23] showed that this approximation is not completely precisely true (as can be seen in Table 1.1), but θ_{13} is still small enough that this approximation is within the precision of our experiment.

1.5 Two Flavor Oscillations

In many experimental cases, oscillation probabilities can be approximated with two flavors, neglecting the third. This assumption simplifies calculation. This assumption is used for the analysis presented in this thesis, after confirming that the difference between two and three flavor oscillations was indeed negligible in this case. Within the simplified 2 neutrino case, the calculation from probability amplitude to probability is concise enough to include here.

This section culminates with the standard two-neutrino oscillation probability formula in the form most useful for long-baseline experiments like IceCube.

If we consider only two mass states and two flavor states, we only have one mixing angle (θ_{23}) and one mass difference (Δm_{32}^2) to deal with:

$$\begin{pmatrix} \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} \cos \theta_{23} & \sin \theta_{23} \\ -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix} \begin{pmatrix} \nu_2 \\ \nu_3 \end{pmatrix}.$$

Through the remainder of this section, the subscript is dropped from the mixing angle.

Plugging this into the probability equation 1.10, we get a nice exercise in trigonometric identities:

$$P(\nu_\mu \rightarrow \nu_\tau | t) = \left| U_{\mu j}^* e^{-iE_j t} U_{\tau j} \right|^2 \quad (1.12)$$

$$= \left| -\cos \theta (e^{-iE_2 t}) \sin \theta + \sin \theta (e^{-iE_3 t}) \cos \theta \right|^2 \quad (1.13)$$

$$= \cos^2 \theta \sin^2 \theta \left| -e^{-iE_2 t} + e^{-iE_3 t} \right|^2 \quad (1.14)$$

Consider the factors outside the square: these can be simplified using the half-angle formula ($2 \sin(\theta) \cos(\theta) = \sin(2\theta)$):

$$\cos^2 \theta \sin^2 \theta = \frac{1}{4} \sin^2(2\theta) \quad (1.15)$$

Then, consider the bit inside the square: we are free to shift the phase by factoring out any imaginary exponent because it cancels itself out with its complex conjugate when we square. If we choose carefully, we can find one that leaves a trigonometric substitution remaining:

$$\left| e^{-iE_3 t} - e^{-iE_2 t} \right|^2 = \left| \left(e^{\frac{-it}{2}(E_2 - E_3)} \right) \left(e^{\frac{+it}{2}(E_2 - E_3)} \right) (e^{-iE_3 t} - e^{-iE_2 t}) \right|^2 \quad (1.16)$$

$$= \left| \left(e^{\frac{-it}{2}(E_2 - E_3)} \right) \left(e^{\frac{it}{2}(E_2 - E_3)} - e^{\frac{it}{2}(E_2 - E_3)} \right) \right|^2 \quad (1.17)$$

$$= \left| \left(e^{\frac{1}{2}it(E_2 - E_3)} - e^{-\frac{1}{2}it(E_2 - E_3)} \right) \right|^2 \quad (1.18)$$

$$= \left| 2i \sin \left(\frac{E_3 - E_2}{2} t \right) \right|^2 \quad (1.19)$$

$$= 4 \sin^2 \left(\frac{E_3 - E_2}{2} t \right) \quad (1.20)$$

The last ingredient of this calculation is to plug in an energy for each mass state. We will put this first in a convenient format, starting from the familiar formula: $E_i^2 = p_i^2 + m_i^2$.

$$E_i = \sqrt{p^2 + m_i^2} \quad (1.21)$$

$$= p \left(1 + \frac{m_i^2}{p^2} \right)^{1/2} \quad (1.22)$$

$$\approx p \left(1 + \frac{1}{2} \frac{m_i^2}{p^2} \right) \quad (1.23)$$

$$= \left(p + \frac{m_i^2}{2p} \right) \quad (1.24)$$

$$(1.25)$$

The approximation here is the binomial expansion $(1+x)^n \approx 1 + nx + \frac{n(n-1)}{2!}x^2 + \dots$, from which we keep the first two terms. This is valid when x is small, or in our case when momentum is much larger than mass, i.e. when the particle is relativistic. This assumption holds for any neutrino we know how to experiment with.

Trivially, then, the difference between two energies is:

$$E_3 - E_2 = \left(p + \frac{m_3^2}{2p} \right) - \left(p + \frac{m_2^2}{2p} \right) \quad (1.26)$$

$$= \frac{\Delta m_{32}^2}{2p} \quad (1.27)$$

Putting those three expressions (1.15, 1.20, and 1.27) back together into equation 1.14, we get a full probability:

$$P(\nu_\mu \rightarrow \nu_\tau) = \frac{1}{4} \sin^2(2\theta) 4 \sin^2 \left(\frac{E_3 - E_2}{2} t \right) \quad (1.28)$$

$$= \sin^2(2\theta) \sin^2 \left(\frac{\Delta m_{23}^2 t}{4p} \right) \quad (1.29)$$

$$= \sin^2(2\theta) \sin^2 \left(\frac{\Delta m_{23}^2 t}{4E} \right) \quad (1.30)$$

The final substitution of E for p works because, again, the neutrino is relativistic. Using the units appropriate for atmospheric oscillations, namely kilometers for length and GeV for energy, equation 1.30 simplifies to the following:

$$\boxed{P(\nu_\mu \rightarrow \nu_\tau | L(\text{km}), E(\text{GeV})) = \sin^2 2\theta_{23} \sin^2 \left(1.27 \frac{\Delta m_{23}^2 L}{E} \right)} \quad (1.31)$$

This formula is central to all the work that follows, and should be memorized by any student working on oscillations. More details about this calculation, or how to put in the neglected factors of c and \hbar , can be found in [24]. Further analysis of this formula from an experimental point of view can be found in Ch. 2.

1.6 Oscillations in Matter

The neutrino oscillations we have described up to this point are valid in a vacuum. When neutrinos travel through matter, the oscillations are modified.

Electron flavor neutrinos can undergo charged current interactions with electrons in the matter, described by an effective potential:

$$V_{\text{CC}} = \pm\sqrt{2}G_F N_e$$

where N_e is the number density of electrons in the matter, G_F is the Fermi constant, N_n is the number density of neutrons in the matter, and the sign is for two cases: positive for neutrinos, negative for antineutrinos. The potential does not apply to muon or tau flavor neutrinos because of the lack of muon and tau leptons in normal matter, but these flavors are affected through three-flavor oscillations with electron flavor neutrinos.²

The familiar two-flavor oscillation formula, when this additional potential is included, is modified by introducing two “effective oscillation parameters”. The modified oscillation probability is:

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2 2\theta_m \sin^2\left(\frac{\pi L}{l_m}\right).$$

The two modified oscillation parameters in matter are an effective oscillation length, l_m , and an effective mixing angle, θ_m , which are defined in terms of the vacuum oscillation mixing angle θ_o :

$$l_m = \frac{2\pi}{\left(\frac{\Delta m^2}{2E} \cos 2\theta_o - \sqrt{2}G_F N_e\right)^2 + \left(\frac{\Delta m^2}{2E}\right)^2 \sin^2 2\theta_o}$$

$$\sin^2 2\theta_m = \frac{\left(\frac{\Delta m^2}{2E}\right)^2 \sin^2 2\theta_o}{\left(\frac{\Delta m^2}{2E} \cos 2\theta_o - \sqrt{2}G_F N_e\right)^2 + \left(\frac{\Delta m^2}{2E}\right)^2 \sin^2 2\theta_o}$$

These matter-modified oscillation parameters both exhibit resonance behavior when the following condition is met:

$$\frac{\Delta m^2}{2E} \cos 2\theta_o = \sqrt{2}G_F N_e.$$

This resonance is called the “MSW effect” after Mikheev, Smirnov, and Wolfenstein, who proposed it [25][26]. When this condition is met, mixing can be large or even maximal, even if the vacuum mixing angle θ_o is small. A method for extending this correction to 3 flavor oscillations has been developed [27].

This resonance has been observed in ν_e states produced in the Sun [16]. Additionally, matter effects can be used to determine the signs of the mass splittings because electrons always have the same charge (and we have no positrons in normal matter). This approach has been used to measure the sign of Δm_{12}^2 [28]. Currently, work is in progress to see if IceCube or an IceCube extension such as PINGU [29] could observe matter effects in atmospheric oscillations, which could constrain the sign of Δm_{23}^2 .

²It should be clear that two-flavor oscillations between ν_μ and ν_τ are unaffected by normal matter.

1.7 Summary of This Work

This dissertation uses the IceCube DeepCore detector to study atmospheric neutrino oscillations. The following chapter discusses neutrino oscillation experiments, first in general terms and then discussing a few experiments other than IceCube. Chapter 3 is a dedicated description of the IceCube detector. Chapter 4 discusses the testing process of the detector modules. Chapter 5 is an overview of the many reconstruction methods used in the event selection, and may be used as a reference for the discussions in later chapters. Chapter 6 provides documentation of the early phases of the IceCube Moon shadow analysis in which I played a major role (the final appendix is the paper that resulted from Chapter 6). Chapter 7 returns to the central topic of this thesis with an overview of the event selection, including a discussion of how it was optimized. The analysis method is described in Chapter 8 and tested in Chapter 9. Finally, Chapter 10 includes results of the analysis and discussion about its implications for further work.

Chapter 2

Neutrino Detectors

*Neutrinos, they are very small.
 They have no charge and have no mass
 And do not interact at all.
 The earth is just a silly ball
 To them, through which they simply pass,
 Like dustmaids down a drafty hall
 Or photons through a sheet of glass.*

— JOHN UPDIKE, 1960 [30]

Neutrino oscillations have been observed in many different contexts, including stellar core processes driving our Sun [16], hadronic decays in the Earth’s atmosphere [15], nuclear reactors [8], and particle accelerator laboratories [31].

This chapter covers the neutrino oscillations material most relevant to an experimentalist. This includes a discussion of interactions and cross sections (§2.1) and the implications of the two-neutrino approximation of the oscillation probability formula (§2.2). The next three sections are more specifically tied to this IceCube analysis. Atmospheric neutrinos (§2.3) produce the signal of this thesis; they can produce muons that interact in the glacial ice (§2.4), particularly through Cherenkov light (§2.5). Finally, brief sketches are included of two other experiments that probe the atmospheric oscillation parameters: Super-Kamiokande (§2.6.1) and MINOS (§2.6.3).

2.1 Interactions and Cross Sections

Neutrinos interact with the ice by the weak force: by exchanging charged W^\pm bosons or neutral Z^0 bosons with the quarks in the nuclei in the atoms in the molecules of ice. The weak mediating bosons were discovered in 1983 and have relatively large masses: $M_W = 80.6\text{GeV}/c^2$, and $M_Z = 91.2\text{GeV}/c^2$. Because of these large masses, the weak force has an extremely short range of $\sim 10^{-18}$ m. If a charged boson mediates the exchange, it is called “charged current,” and in our case (when the progenitor lepton is a neutrino) a charged current interaction necessarily produces a charged lepton in the final state. Feynman diagrams for these interactions are shown in Fig. 2.1. This is the only way we can observe neutrinos; since they do not interact electromagnetically, we cannot see light from them directly, but we can see the products of their interactions.

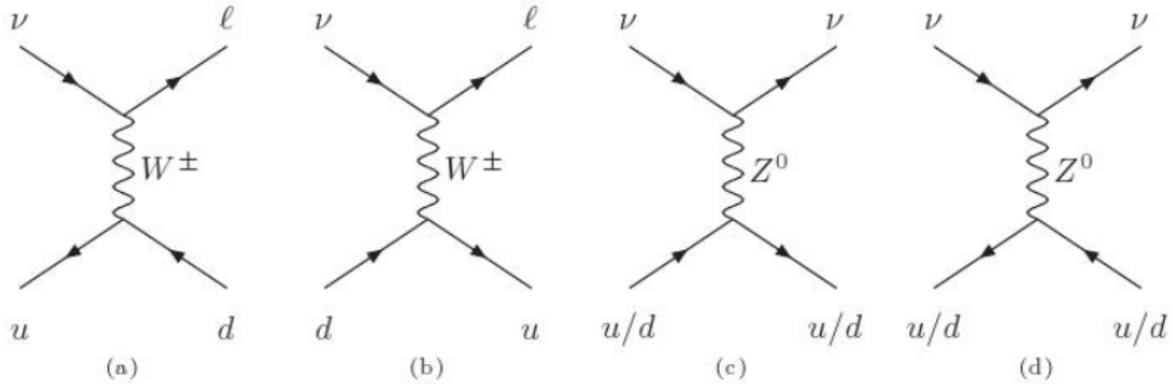


Figure 2.1: Feynman diagrams of the interactions through which we observe neutrinos. Interactions (a) and (b) are called *charged current* because the mediating boson is charged; similarly, (c) and (d) are called *neutral current*. Note that the charged current interactions each have a charged lepton output particle. Observing these charged leptons is how IceCube tells a ν_μ event apart from a ν_e or neutral current event. However, the IceCube detector resolution is not fine enough to distinguish between ν_e and neutral current events.

Neutrino cross sections are extremely energy dependent [32] with a rough direct proportionality to their energy. This gives rise to the main shape visible in Fig. 2.2. The other main feature in Fig. 2.2 is the resonance from the W^- mass above 1 PeV, called the Glashow Resonance because in the 1960s, Glashow proposed [33] using this resonance to directly observe the W^- .

Cross sections have units of area, most often reported in cm^2 or barns; one barn is 10^{-24}cm^2 . The name “barn” as a unit was first used to describe neutrons scattering on uranium, calling a uranium nucleus “as big as a barn.” [35].

A popular neutrino trivium is that a neutrino can pass through a light year of lead without interacting, or more precisely, a sample of many neutrinos can travel through 22 light years of lead with only half of them interacting. This is only true at energies around 1 MeV, the energy of neutrinos created by fusion within the Sun. For a typical DeepCore 50 GeV neutrino sample, it takes only 2 millilightyears of lead (140 astronomical units) to stop half the neutrinos. For the most extreme IceCube events on the order of 1 PeV, the Earth itself is opaque [36].

Through different energy ranges, several different interaction mechanisms become dominant as the neutrino can probe different size objects coherently (see Fig. 2.3), reaching the idealized simple quark picture presented in Fig. 2.1 only when the neutrino reaches the Deep Inelastic Scattering (DIS) regime above about 500 GeV. This is where neutrinos interact with nuclear partons as point particles, and the interaction dissociates the parton from the nucleus.

This is the energy range most relevant for IceCube. However, DeepCore operates at the edge of this assumption, in the crossover between DIS and the more complex quasi-elastic scattering. At typical DeepCore energies of 20 GeV,

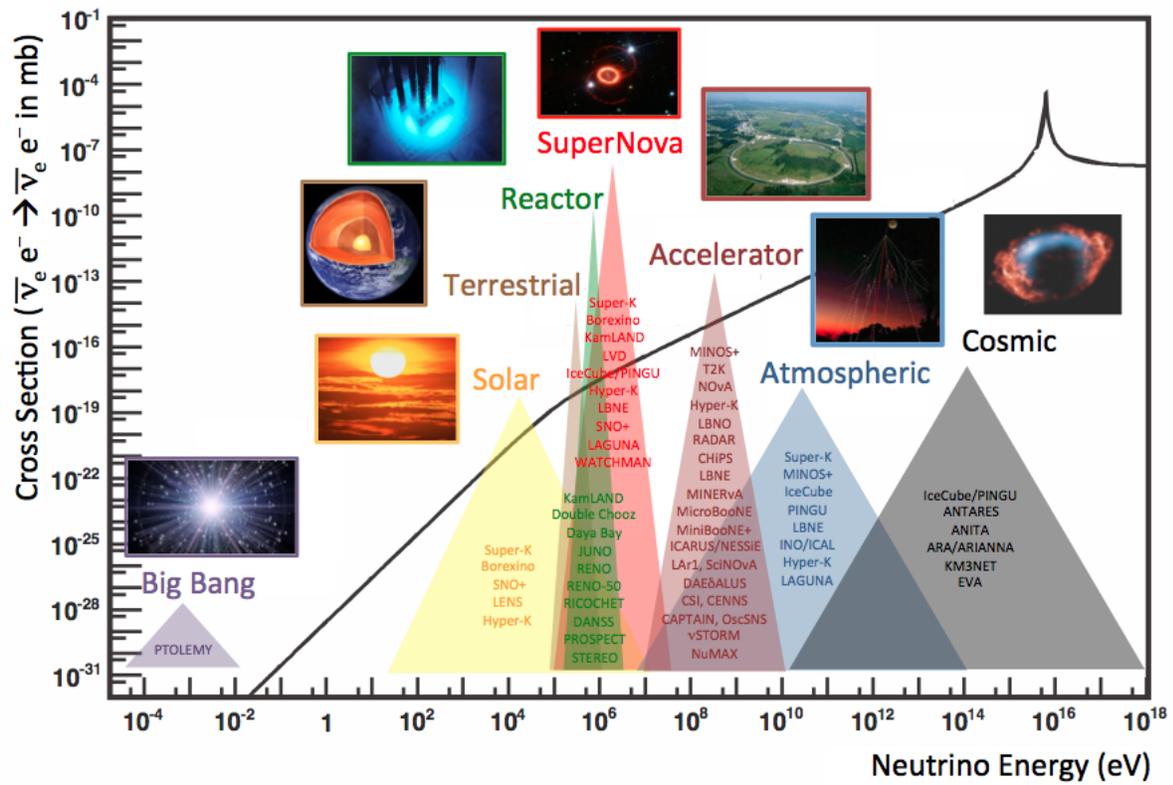


Figure 2.2: Neutrino $\bar{\nu}_e e^- \rightarrow \bar{\nu}_e e^-$ charged current scattering cross section as a function of energy, with representative pictures of sources at several energies; this is the keystone figure from [32], modified for inclusion in [34]. The peak at 10^{16} eV is due to the W^- resonance.

the charged current cross section is $\sim 1.4 \times 10^{-37} \text{cm}^2$, i.e. one neutrino out of every 10^{10} interacts in a cubic kilometer of ice. At much higher energies around 1 PeV, this cross section increases to $\sim 1 \times 10^{-31}$, or 1 neutrino in 10^4 . The neutrino cross section is about a factor of 3 larger than the antineutrino cross section because the antineutrino cross section is helicity suppressed [37]¹.

This small probability of interacting is the major reason that neutrino experiments need detectors that are, as a rule, larger than detectors for other types of physics: you need many nuclei to give the neutrinos many chances to interact, if you want any hope of statistical significance for your experiment. As an example, consider the difference in physical size between an optical telescope, around a meter or perhaps as much as tens of meters, and our neutrino telescope, at a cubic kilometer.

Considering effective areas instead of physical areas makes it easier to compare neutrino telescopes to other types of telescopes. A more useful way of looking at interaction cross sections and probability is to combine it with detector efficiencies; that way, instead of answering the question “given N neutrinos passing through, how many interact?” we can instead answer “how many should we see, assuming our estimate of neutrino flux is correct and we understand our detector?” This is the motivation for reporting energy dependent effective areas and effective volumes. These measures depend on the data selection, so they are typically reported with each analysis. As an example, a recent IceCube search for neutrino sources [38] reported energy dependent effective areas as a function of arriving zenith angle. For neutrinos at 10 TeV and within 30° of the zenith (i.e., downgoing), the effective area was 10m^2 ; this is close to the physical size of an optical telescope.

2.2 A Good Oscillation Experiment

Recall the 2 neutrino approximation of the oscillation probability formula from Ch. 1:

$$P(\nu_\mu \rightarrow \nu_\tau | L(\text{km}), E(\text{GeV})) = \sin^2 2\theta_{23} \sin^2 \left(1.27 \frac{\Delta m_{23}^2 L}{E} \right)$$

This function splits up neatly into two sine functions: one that depends only on the mixing angle, and one that depends on the mass splitting and experimental parameters L and E .

The first sine is often called the “amplitude” of the oscillation, because it multiplies the rest. In experiments, when we search for a deficit due to oscillations, the depth of this deficit is determined by the amplitude, and thus by the mixing angle. The errors on these measurements are dominated either by statistics (typically, neutrino experiments have some of the smallest number of observations in modern particle physics) or by the errors on flux predictions. An experiment that can start from an accurate flux prediction, with small errors, and that can produce a large sample of neutrinos could measure mixing angle particularly well.

¹Specifically, §12.8 of [37].

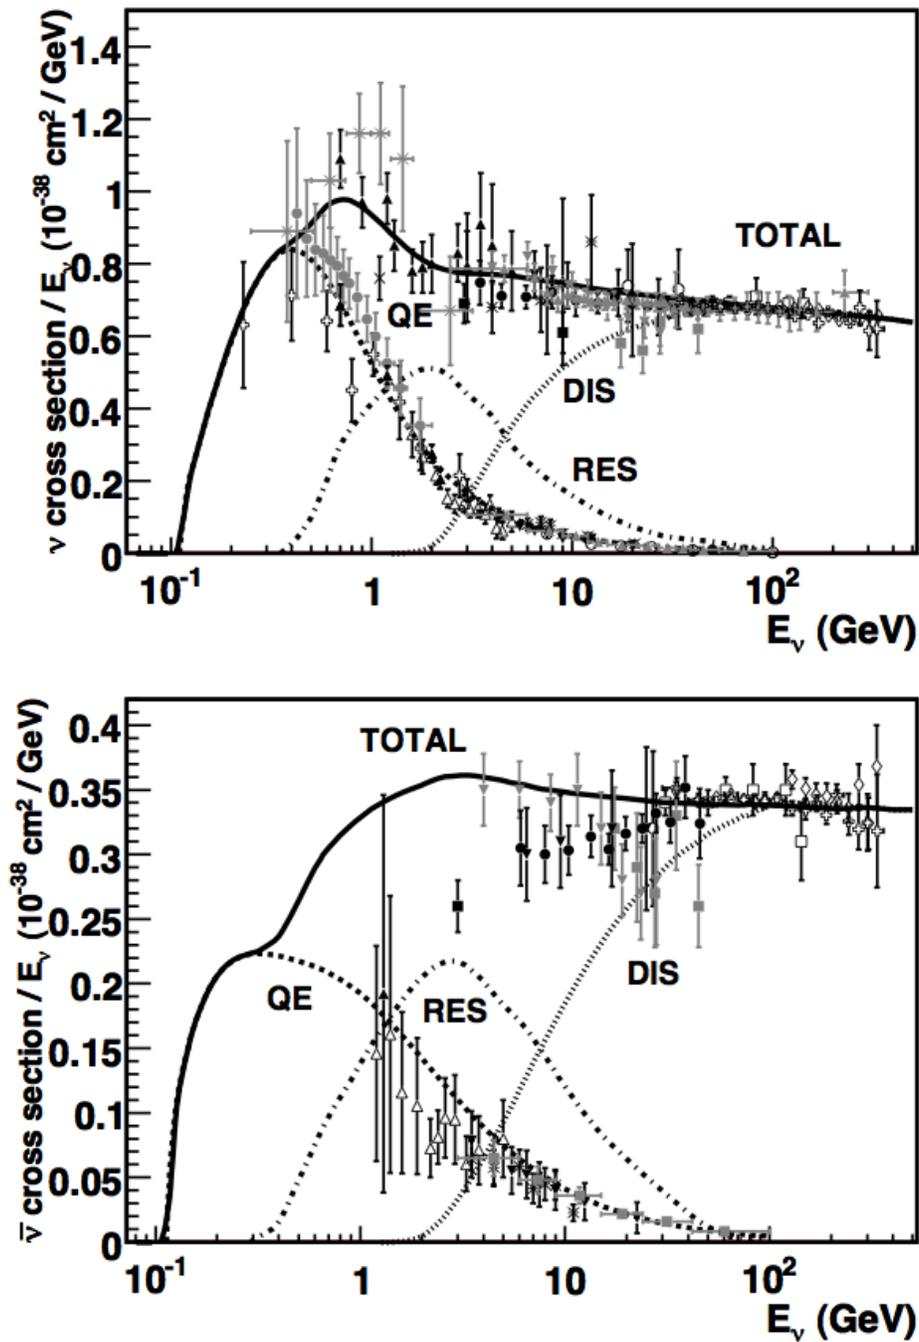


Figure 2.3: Total neutrino and antineutrino per nucleon CC cross sections (for an isoscalar target) divided by neutrino energy and plotted as a function of energy, from [32]. Note that the y-axis is divided by neutrino energy, to focus more on smaller structures; also note the different y-axis scales between the two plots. The predicted cross section (solid black line) is the sum of three contributing processes: quasi-elastic scattering (QE), resonance production (RES), and deep inelastic scattering (DIS).

The second sine function depends on three parameters: the mass difference Δm^2 (set by nature) and the experimental parameters L and E . To measure the mass difference, we look at the position of the flux deficit in our energy spectrum; systematic shifts in reconstructed energy produce errors in our Δm^2 measurement. Thus, a good experiment for measuring Δm^2 needs good energy resolution.

The energy distribution of atmospheric neutrinos (approximately describable as a power law) is very different from the energy distribution of neutrinos created in an accelerator complex (approximately describable as a normal distribution around a single energy), even though both have uncertainties in both flux normalization and shape. Additionally, the length the neutrinos travel while oscillating varies from negligible to the diameter of the Earth ($\approx 12,000$ km). This is a direct function of the zenith angle between the point of production and the detector (at the South Pole), so the oscillation probability can be plotted equivalently either as a function of zenith or as a function of $\cos(\text{zenith angle})$. Both of these are plotted in Figure 2.4.

2.3 Atmospheric Neutrinos

This work focuses on the oscillations of atmospheric neutrinos, the main source of neutrinos seen by IceCube. Atmospheric neutrinos are created when cosmic rays interact with the Earth's atmosphere.

Cosmic rays were first observed in 1912 by Victor Hess[39]. They are (see Fig.2.5) predominantly high energy protons. They strike the atmosphere nearly equally from all directions [40].

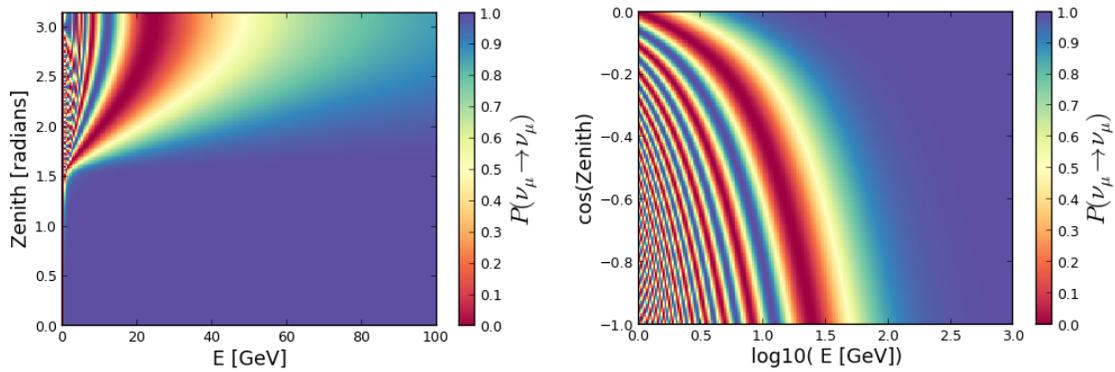
Through most of the observed energy range, cosmic rays follow (see Fig. 2.6) a power-law spectrum:

$$\frac{dN}{dE} \propto E^{-\gamma} \text{ with } \gamma \approx 2.7$$

Recent data [44] has shown that the spectral index is closer to 2.65, so that value is used as the default in this thesis. At energies around 3 PeV, the spectrum steepens; this is known as the “knee.” Around 3 EeV, it flattens out again; this is known as the “ankle.” There is considerable structure in between, which is a field of active research. The subject of this thesis is restricted to lower energies, between 1 and 100 GeV; in this region, a power-law spectrum is a good approximation.

The spectrum of atmospheric neutrinos is created from the spectrum of cosmic rays. Products of atmospheric interactions fall into three broad categories based on their energy: high energy, low energy, and a transition. The regions are defined by the relationship between interaction length and decay length for the outgoing particles. If the particles are very high energy, they are most likely to interact before decaying; if they are very low energy, they are most likely to decay before interacting.

Consider first the low-energy case: when particles have time to decay before reaching the ground or detector. This case arises when considering $\mathcal{O}(\text{MeV})$ low energy atmospheric neutrinos or distant astrophysical sources of neutrinos. Mesons are created in the resulting showers, especially pions (π) and kaons (K), which then quickly decay into



(a) Oscillation probabilities as a function of zenith angle and energy. (b) Oscillation probabilities as a function of $\cos(\text{zenith})$ and $\log(\text{energy})$. These are the variables used for the analysis in following chapters.

Figure 2.4: Oscillation probabilities for atmospheric neutrinos visible to IceCube. Neutrinos are assumed to have started in the Earth's atmosphere, so the length traveled is a function of the arrival zenith angle and the diameter of the Earth D : $L = -D \cos(\text{zenith angle})$. Note that in IceCube the visible energy cuts off around 10 GeV; lower energies are included in these plots only for illustration. These calculations assume two neutrinos and no matter effects.

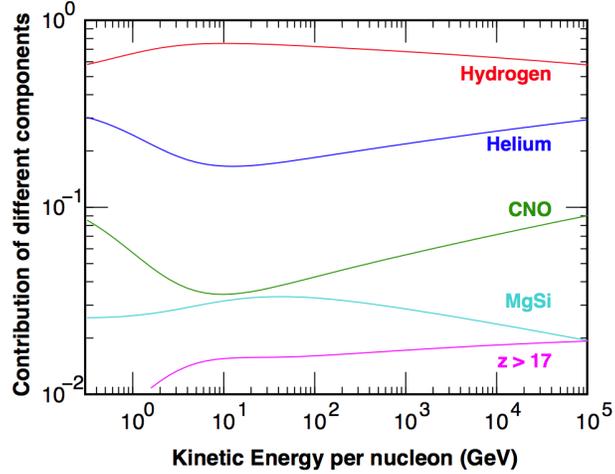
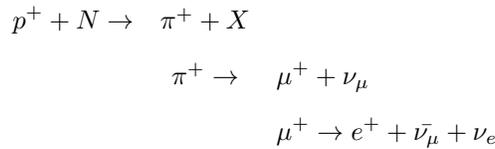


Figure 2.5: Fractional contributions to cosmic ray composition as a function of energy. Note that in the 10-50 GeV region most important for oscillations, hydrogen contributes the vast majority. Figure from [41]. Efforts to extend these measurements to higher energies are underway in IceCube; see for example [42].

neutrinos. A typical decay chain starting with a pion can take the form:



Notice that this decay chain produces twice as many electron neutrinos (taken together with antineutrinos) as muon neutrinos. The decay chain of a π^- produces the same flavor ratio, with opposite sign charge of each particle, and kaons have a similar decay chain [45]. This produces a flavor $e : \mu : \tau$ ratio of 1:2:0, which is often assumed for low-energy or far astrophysical sources, although other flavor ratios are also supported by theoretical predictions [46]. Note that this decay chain does not produce any ν_τ , and thus any ν_τ appearing must be from neutrino oscillations². The signal for this analysis is ν_μ , and the contributions from ν_e and ν_τ are confounding backgrounds.

At the other end, in the high energy limit, no particle has time to decay before reaching the detector.

The crossover between the high and low energy regions depends on the interaction and decay properties of each particle involved. The energy where the interaction length and the decay length are the same is called the “critical energy.” For pions and kaons, the transition between these two extremes takes place at $\mathcal{O}(10 - 100\text{GeV})$, exactly the target energy region for this analysis. For muons, the transition is even lower, at 1 GeV, so one may assume that all

²It is possible to create ν_τ from the “prompt” decay of hadrons containing charm quarks, e.g. $D^+ \rightarrow \tau^+ \nu_\tau$, but since these interactions only become a prominent part of the flux at much higher energies, they are neglected in this work.

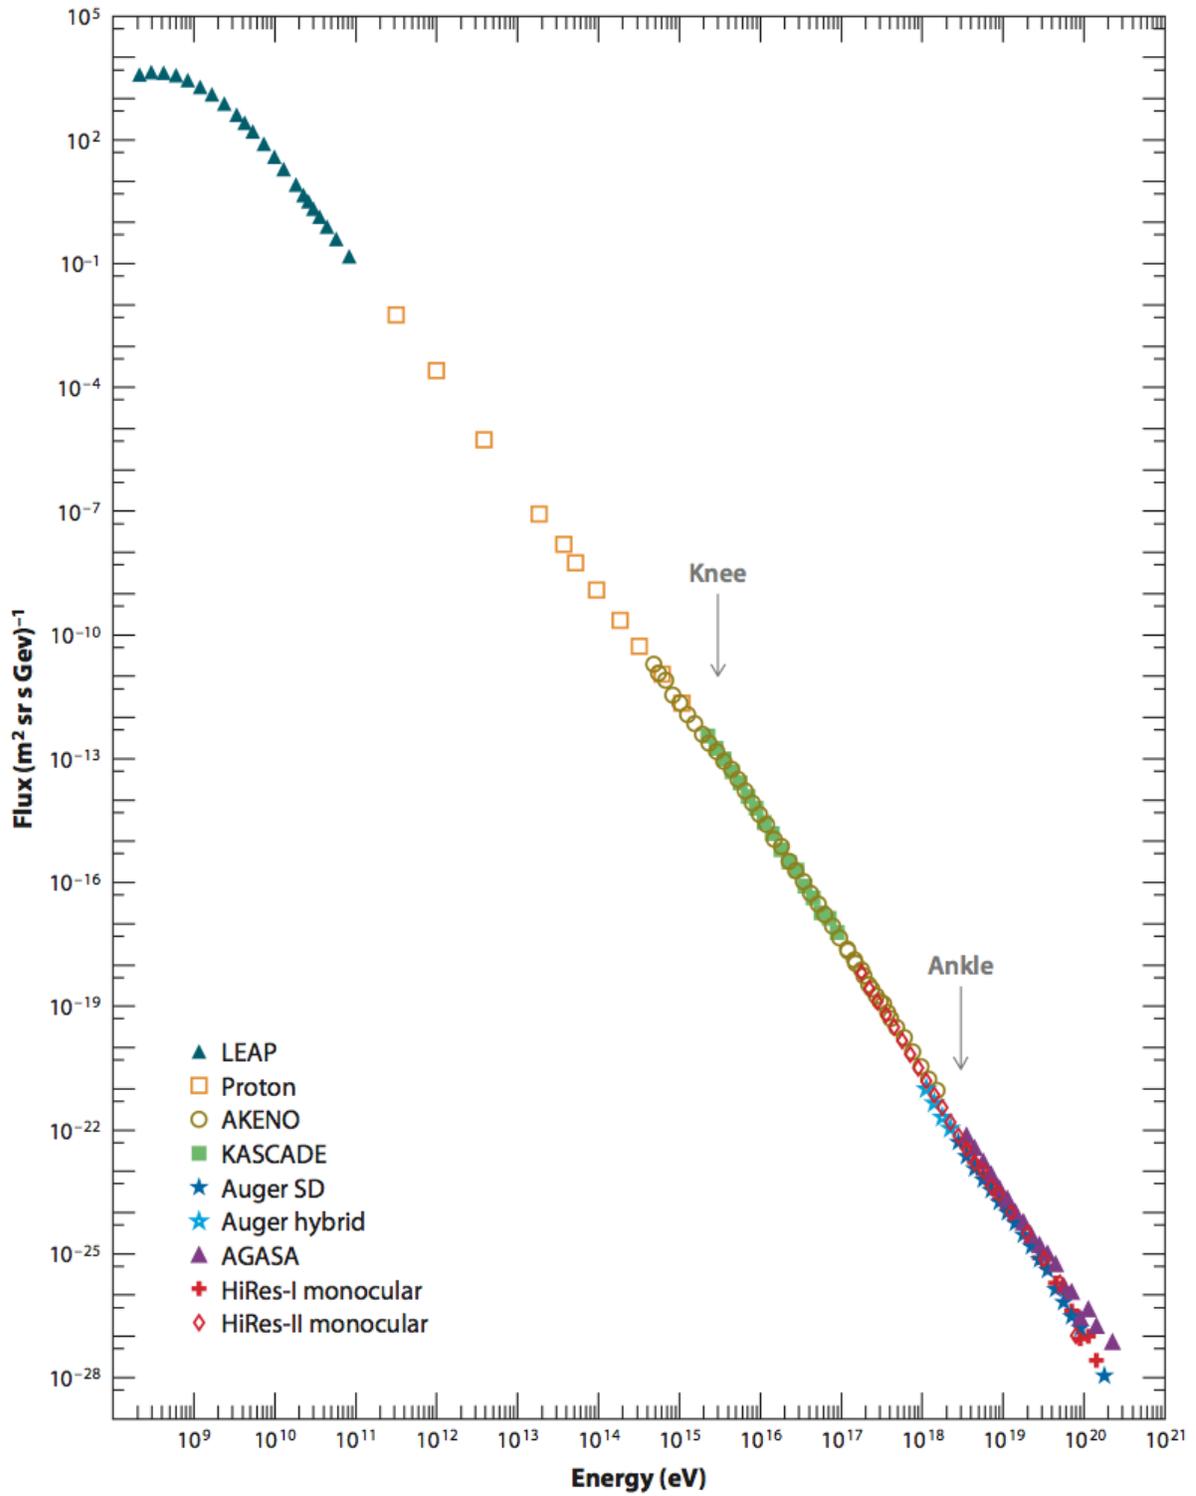


Figure 2.6: Observations of cosmic rays over many orders of magnitude in energy. From [43].

muons reach the ground from the upper atmosphere without decaying. The energy-dependent fraction of pions and kaons producing ν_μ is shown in Fig. 2.7.

Within this transition region, the spectrum of atmospheric neutrinos is generally one index steeper than the cosmic ray spectrum, that is, it follows the following power law:

$$\frac{dN}{dE} \propto E^{-3.7}$$

where E is the energy of each neutrino.

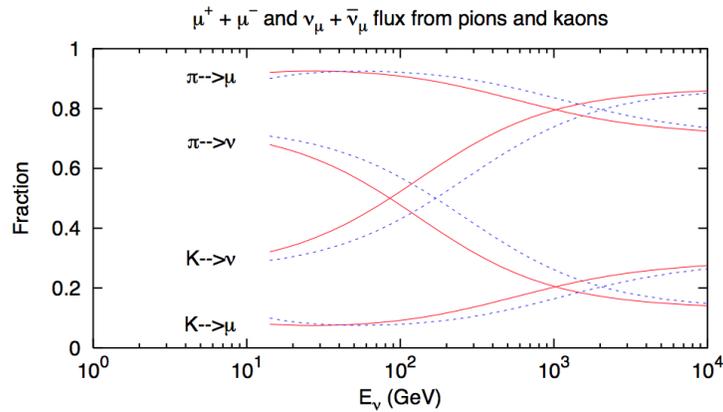


Figure 2.7: The energy-dependent fraction of ν_μ produced in the atmospheric neutrino spectrum from pions and kaons. The solid lines are from particles traveling vertically downward; the dashed lines are for a zenith angle of 60° . From [41].

These transitions and resulting spectra are discussed in detail in [48], and specific flux predictions are made in [49] and [50]. However, these calculations still leave the flux normalization somewhat uncertain; the two models disagree by 20%. Within these uncertainties, we must make assumptions to construct a simulation for analysis. Details of this simulation are discussed later in 7.2; the assumptions relevant to this section are the following. The total flux normalization is left to float (with a Gaussian penalty to the likelihood for wandering too far) within a restricted range of $\pm 50\%$ around the central value. The central value of the flux model is from Honda [49], but as a test the Bartol [50] model is also considered. The differences were found to be smaller than other systematic effects (e.g. the noise model).

2.4 Muon Energy Losses in Matter

IceCube detects neutrinos indirectly by observing the charged particles they produce. Describing the energy loss mechanisms of these charged particles is thus central to understanding the IceCube detector response to neutrinos over

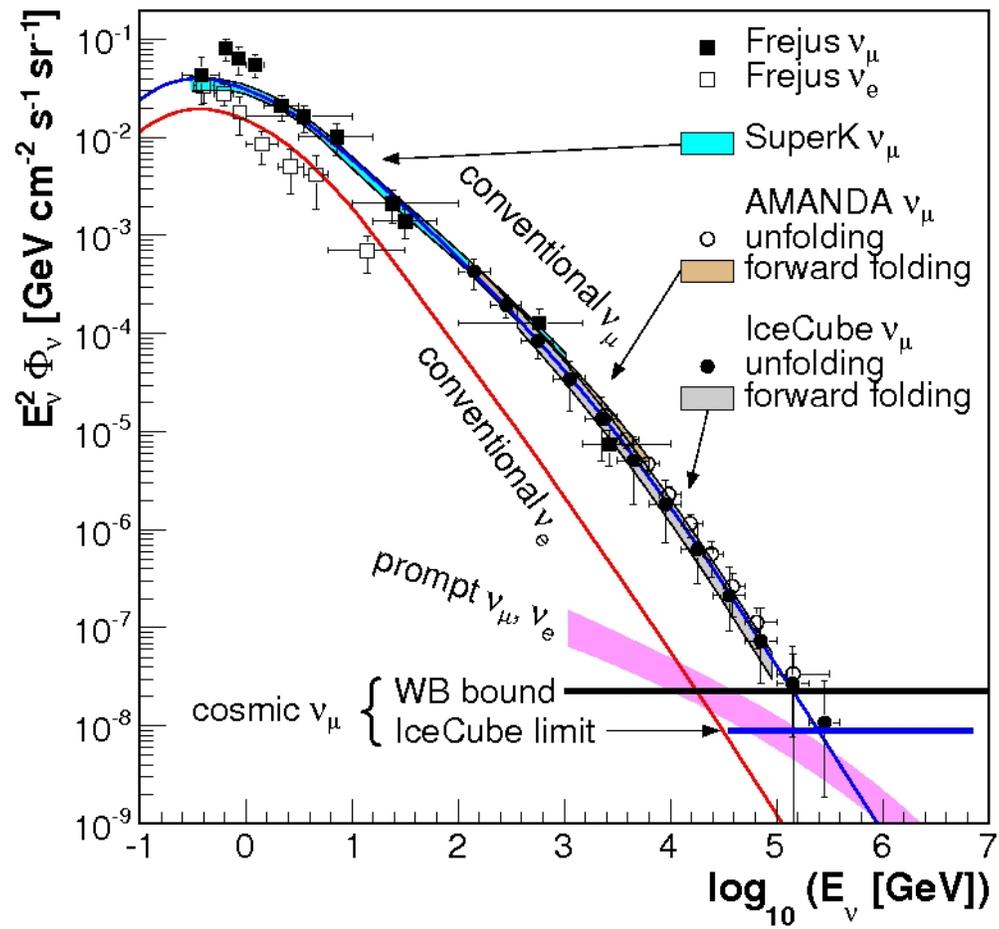


Figure 2.8: Predicted and measured atmospheric neutrino fluxes as a function of energy [47].

a wide range in energies. Muons are particularly central for understanding IceCube because their long penetration lengths match the detector's sparse inter-string spacing, making IceCube particularly good at observing muons.

As shown in Fig. 2.9, muons lose energy through different mechanisms at different energies. The analysis of this thesis is focused on the 10-50 GeV region, where muons lose energy predominantly through ionization. At higher energies (which encompass most IceCube analysis), stochastic energy loss mechanisms like bremsstrahlung dominate over ionization. This gives rise to various reconstruction methods relevant for different energy ranges (discussed in more detail in Ch. 5).

2.5 Cherenkov Light

Cherenkov light is the signal in IceCube. It is emitted by relativistic charged particles produced in neutrino interactions. Cherenkov light is a shock wave of light³. As a charged particle moves, it disturbs the electromagnetic field in its vicinity. The signal of this disturbance travels out at the speed of light in the current medium, which is ice in our case. Light travels slower in a medium than in a vacuum by a factor called the index of refraction, n .⁴ In ice, the index of refraction is about 1.33. When a charged particle travels faster than c/n , the disturbances in the field interfere constructively to create a shock wave, which is observable as light.

Because of the geometry of shocks, Cherenkov light travels out from a particle track at a specific angle given by the index of refraction n and the speed of the particle $\beta = v/c$:

$$\cos(\theta_c) = \frac{1}{n\beta}$$

This angle is illustrated in Fig. 2.10.

The number of Cherenkov photons emitted per unit track length is given by the Frank-Tamm formula [21]:

$$\begin{aligned} \frac{d^2 N}{dx d\lambda} &= \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2} \right) \\ &\approx 370 \sin^2(\theta_c) \left(\frac{E}{\text{eV}} \right) \left(\frac{d}{\text{cm}} \right) \end{aligned}$$

where α is the fine structure constant, λ is the wavelength of the light, E is the energy of the emitting particle, and d is the unit track length. The $1/\lambda^2$ dependence here means that the spectrum is mostly at shorter wavelengths, towards blue and ultraviolet. The range visible to IceCube is cut off on the ultraviolet end by the absorption by the glass in the DOM housing and PMTs at around 300 nm [51]. The DOM light acceptance as a function of light wavelength is shown in Fig. 2.11.

³ Examples of shock waves in other media include: sonic booms, when jets or whips travel faster than the speed of sound in air; and bow shocks and wakes, when boats or ducks travel faster the speed of waves across water.

⁴It is possible, then, for a particle to travel faster than light in ice but still slower than light in a vacuum, and thus still preserve the well-known speed limit of causality.

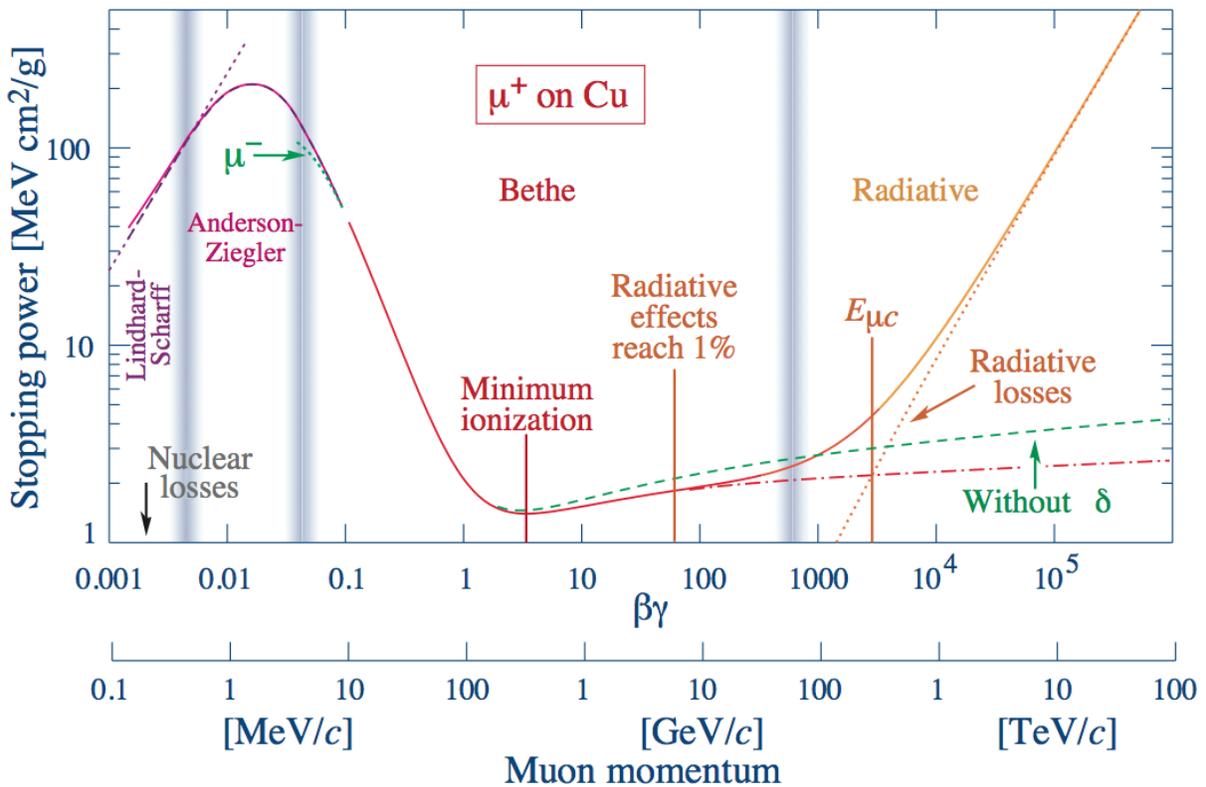


Figure 2.9: Muon energy losses as a function of energy. Although this plot illustrates muons striking copper, a similar shape arises for ice. The vertical shading represents the approximate transition energy between dominant energy loss mechanisms. From [21]

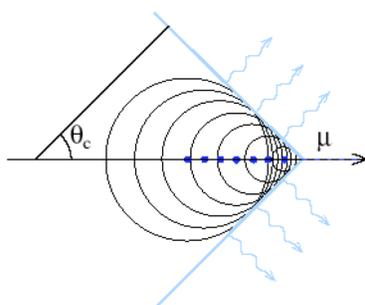


Figure 2.10: The Cherenkov angle and wavefront relative to the particle track. Note that this diagram is independent of wavelength-dependent dispersion. Each black circle represents the disturbance emitted while the particle was at a single blue point. The blue wavy arrows represent the path of the resulting shock wave as it travels forward, slightly behind the particle.

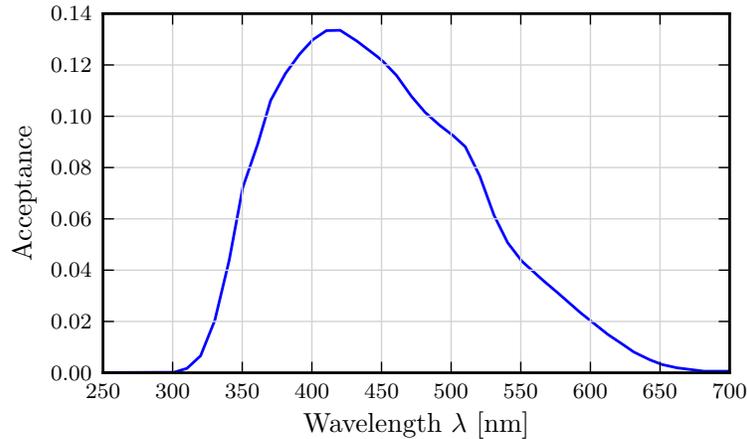


Figure 2.11: DOM light acceptance as a function of light wavelength. This shows the values used in direct photon propagation simulation. The y-axis shows the fraction of simulated photons detected by the DOM relative to total simulated photons striking the DOM.

2.6 Current Neutrino Detectors for Measuring Atmospheric Mixing

There are many experiments currently working on measuring atmospheric neutrino oscillations. Three in particular are able to measure the neutrino oscillation parameters θ_{23} and Δm_{23}^2 with precision similar to the projected IceCube DeepCore sensitivity. They are the Super-Kamiokande experiment in Japan studying atmospheric neutrinos, the Tokai to Kamioka, or “T2K,” experiment also in Japan studying accelerator neutrinos (using some of the same detectors as Super-Kamiokande), and the MINOS experiment starting from Fermilab in Illinois and studying the same L/E region but using long-baseline accelerator neutrinos. These are discussed in more detail in the following sections. Finally, the results of this thesis can be compared to a previous IceCube analysis based on the data from IC79 [52]; this detector configuration is described in Ch. 3. All of these results are compared to one another and this analysis in the concluding Ch. 10.

2.6.1 The Super-Kamiokande Experiment

The Super-Kamiokande experiment [53], or Super-K for short, was originally built to study proton decay; it is a relatively densely instrumented tank of water that was huge by the standards of neutrino experiments when it was built. Famous photos show a person or two people, in a small inflatable boat, looking tiny compared to the 39.3 m diameter and 41.4 m height of this cylindrical tank. The total fiducial volume is 22.5 kton. Super-K has been taking data since 1996, with an interruption in 2001 following a phototube implosion shock wave accident, and in 2011 following an earthquake and tsunami. They have analyzed 3903 days of atmospheric data to study oscillations [54].

The approximate energy above which the Super-K detector cannot distinguish between energies is 5 GeV; this is lower than the IceCube minimum energy threshold of 10 GeV. It has over 11,129 phototubes (comparable to IceCube's 5484, in a much smaller volume). Additionally, each Super-K phototube is larger: 50 cm to IceCube's 25 cm. This very high photocathode coverage ($\sim 40\%$) allows Super-K to perform detailed reconstructions of Cherenkov rings. Despite the excellent energy reconstruction that comes from such detailed coverage, the atmospheric analysis is less sensitive in measuring the mass-squared difference because of statistical errors. Results from 2012 [54] are shown in Fig. 2.12

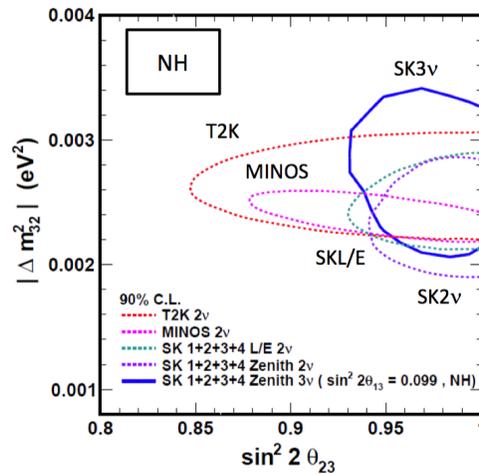


Figure 2.12: Atmospheric oscillation limits from Super-K [54].

2.6.2 The Tokai to Kamioka Experiment

The Tokai to Kamioka Experiment, typically called “T2K,” is a long-baseline accelerator neutrino oscillation experiment. Neutrinos are produced at the Japan Proton Accelerator Research Complex (J-PARC) and aimed towards the Super-K detector, 293 km away. The beam central energy is chosen at 0.6 GeV to focus on the first oscillation maximum. Detectors within J-PARC near the origin of the beam constrain systematic errors. The beamline is produced slightly off-axis (2.5°), which reduces the total flux of neutrinos but narrows the energy range of the neutrinos that reach the detectors. T2K has published a wide variety of results including cross section measurements, flux predictions, and oscillation results. In March 2014, using data collected from 6.57×10^{20} protons on target, they published [55] the most precise measurement to date of θ_{23} , including the figure included here as Fig. 2.13. The precise knowledge of the neutrino flux from the accelerator beamline allows T2K to study the oscillation mixing angle more precisely than Super-K.

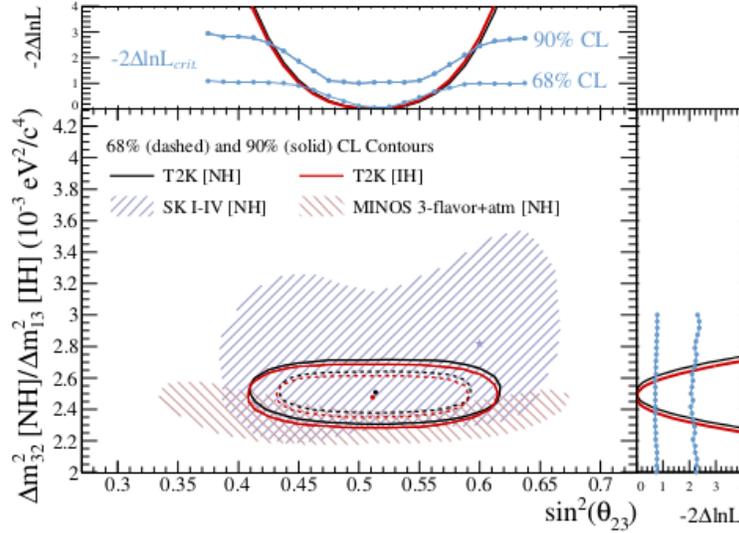


Figure 2.13: Oscillation limits from T2K [55].

2.6.3 The MINOS Experiment

The MINOS experiment was designed specifically to study neutrino oscillations. MINOS is an acronym for **M**ain **I**njector **N**eutrino **O**scillations **S**earch. It uses a 120 GeV proton beamline from the Main Injector ring at Fermilab in Batavia, Illinois to create a neutrino beam at 3 GeV. The neutrinos are nearly all muon flavor when they are created. MINOS uses two similar detectors: one near the start of the neutrino beam, at Fermilab, and one 735 km away in Soudan, Minnesota. Both detectors use alternating plates of magnetized steel and plastic scintillator. The magnetization allows MINOS to distinguish particles from antiparticles in searches for CP violation. This has also been useful in their investigations of the Moon Shadow (see Ch. 6). MINOS published constraints on θ_{23} and Δm_{32}^2 in May 2013, which are shown in Fig. 2.14.

MINOS is in the process of upgrading the detector; the new configuration is called MINOS+. It will have upgraded electronics relative to MINOS, and take advantage of the neutrino beamline passing through the MINOS detector en route to the NOvA experiment. The physics analysis will focus on precision and non-standard neutrino interaction measurements. An early oscillation result from MINOS+ is shown in Fig. 2.15.

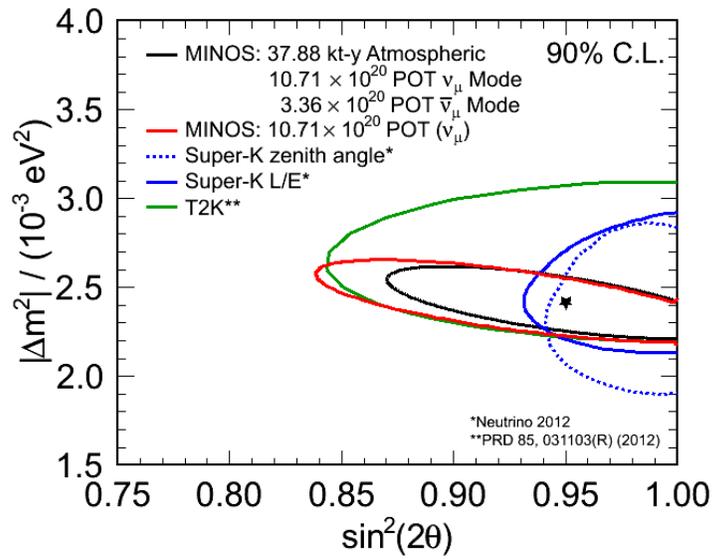


Figure 2.14: MINOS ν_μ disappearance oscillation measurement using beam and atmospheric data [56].

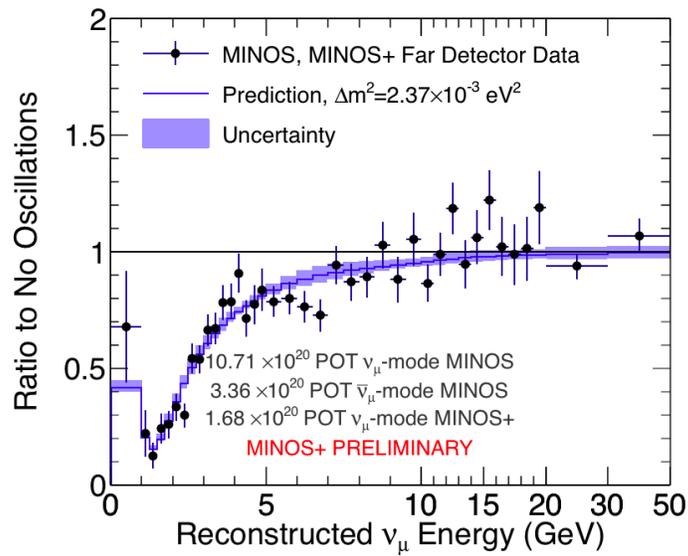


Figure 2.15: First MINOS+ results: ν_μ survival probability curve, from June 2014 [57].

Chapter 3

The IceCube DeepCore Detector

The IceCube Neutrino Observatory is designed to observe particles traveling relativistically through the Antarctic ice at the South Pole. IceCube is a cubic kilometer of ice instrumented with 5160 photomultiplier tubes. The observatory is comprised of several subdetectors. DeepCore is the low-energy infill array, which is the main sub-detector used in this analysis.

When a neutrino interacts with the ice and produces secondary particles (see §2.1) which may include charged leptons, hadronic showers, electromagnetic showers, or a combination of those with variations in relative energy. Those secondary particles can create Cherenkov light (see §2.5), which IceCube observes. The light is detected by photomultiplier tubes (PMTs); each PMT together with its electronics and housing is called a Digital Optical Module, or DOM. A set of 60 DOMs is deployed together into the ice along a single readout cable; this entire assembly is called a string. The timing and position of observed Cherenkov light is used to reconstruct the direction of the original neutrino, which can then be used as part of a sample for studying the astrophysical sources of neutrinos. The direction and energy measurements from neutrinos with the lowest energies visible to IceCube, especially those made visible by the DeepCore [58] infill array, are the topic of this thesis.

This chapter covers aspects of neutrino detection specific to IceCube but common across many analyses. §3.1 summarizes the subdetectors that together make up the observatory. §3.2 describes the geometric arrangement of the DOMs and strings making up the in-ice array and DeepCore subdetector. §3.3 discusses our knowledge and modeling of how light travels through the Antarctic glacial ice. §3.4 describes the DOMs (DOM testing is discussed in the next chapter). Finally, §3.5 outlines the first level of event selections, called “filters.”

3.1 IceCube Subdetectors

The observatory as a whole is made up of several subdetectors, each of which is designed to observe a specific energy range of particle interactions. They each complement one another.

- IceCube in-ice array [51]: the main subdetector of the IceCube Observatory. The in-ice array consists of 5160 DOMs, as described in the following section. This subdetector is designed to observe $\sim\text{TeV}$ to $\sim>\text{EeV}$ neutrinos interacting in the ice.

- DeepCore [58]: the low-energy infill array, also described in detail in the following section. DeepCore is designed to extend the energy range visible to IceCube as low as ~ 10 GeV.
- IceTop [59]: a cosmic ray detector on the surface of the ice sheet. IceTop is composed of ice tanks observed by DOMs, and can observe cosmic ray showers with ~ 100 TeV to $\sim > 1$ EeV energies. Using IceTop as an additional veto for in-ice neutrino analyses helps guarantee their astrophysical purity [60], and using the in-ice array helps IceTop measure the energy and composition of the showers it observes.
- DM-Ice [61]: a dark matter experiment deployed within the in-ice array, using NaI crystals to confirm or refute controversial claims [62]. DM-Ice shares some electronics with IceCube, including readout cables and circuit mainboard design, and uses all of IceCube as an active veto.
- PINGU [29]: the **P**recision **I**ceCube **N**ext **G**eneration **U**ppgrade. PINGU is a proposed infill array that would extend the energy range as low as 1 GeV. This would open the possibilities of further indirect dark matter and neutrino oscillation studies, including determining the mass hierarchy of neutrinos.
- AMANDA: the **A**ntarctic **M**uon **A**nd **N**eutrino **D**etector **A**rray. This was the prototype detector for IceCube, and proved that neutrinos were observable in the antarctic ice [63], and that reconstructing their directions was possible [64]. AMANDA stopped running in 2009.
- ARA [65]: the **A**skarian **R**adio **A**rray, a proposed high-energy radio extension. ARA could observe ~ 100 PeV to $\sim > 100$ EeV energy neutrino events.

3.2 IceCube and DeepCore Geometry

The main in-ice array (including DeepCore) consists of 86 strings, each containing 60 DOMs. In total there are 5160 in-ice DOMs, as shown in the drawing in Fig. 3.1. The strings are spaced in a triangular grid with ~ 125 m between strings, as shown in the reference graph Fig. 3.2.

DeepCore [58], a subarray of the IceCube detector [51], was designed to provide sensitivity to neutrinos at energies more than an order of magnitude lower than initially envisioned for the original IceCube array. This improvement is achieved through a combination of smaller DOM spacing in the clearest ice, higher quantum efficiency PMTs, and using the rest of IceCube as an active veto.

The DOMs in IceCube are deployed between a depth of 1450 m and 2820 m¹ in the Antarctic ice sheet. The ice in these regions is extremely clear to 400 nm light: in the clearest ice, around 2400 m depth, the effective scattering length is close to 50 m, and the effective absorption length is close to 190 m (see §3.3). The quantum efficiency of the PMTs

¹This 2820 m depth is comparable to the 2037 m height above sea level of Mount Mitchell, the highest peak of the Appalachian Mountains [66].

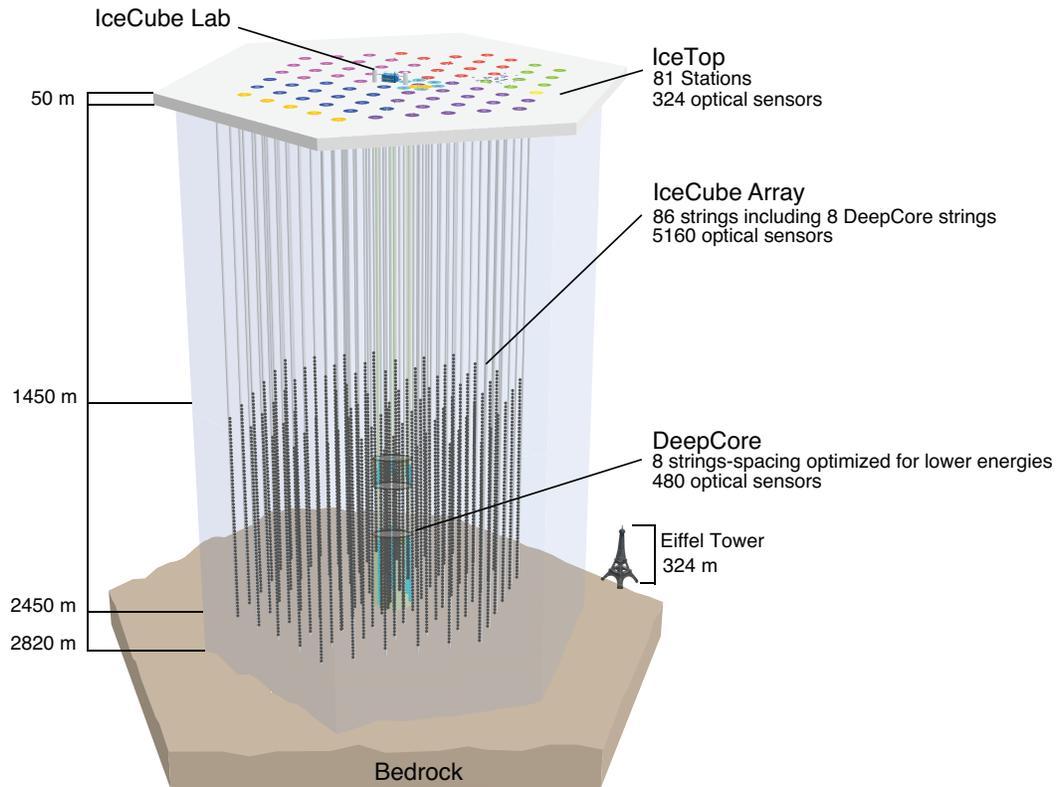


Figure 3.1: The detector layout, from the side as if through the ice sheet. The black dots each represent a DOM, which start at a deployed depth of 1.45 km and extend for an additional kilometer. The dots on the surface represent the locations of the drill holes where each string was deployed; the color of the dot corresponds to the year of deployment. The shaded blue cylinder in two parts at the center of the detector is the DeepCore infill array. The cluster of very small blue dots on the surface (in the same region as the green dots) is the deployment locations of the Amanda strings, now disconnected. The Eiffel Tower is shown on the right for a sense of scale. The IceCube lab (ICL) is a two-story building with a tower on each side to direct cables to the upper story; the ICL image has been enlarged to show some detail.

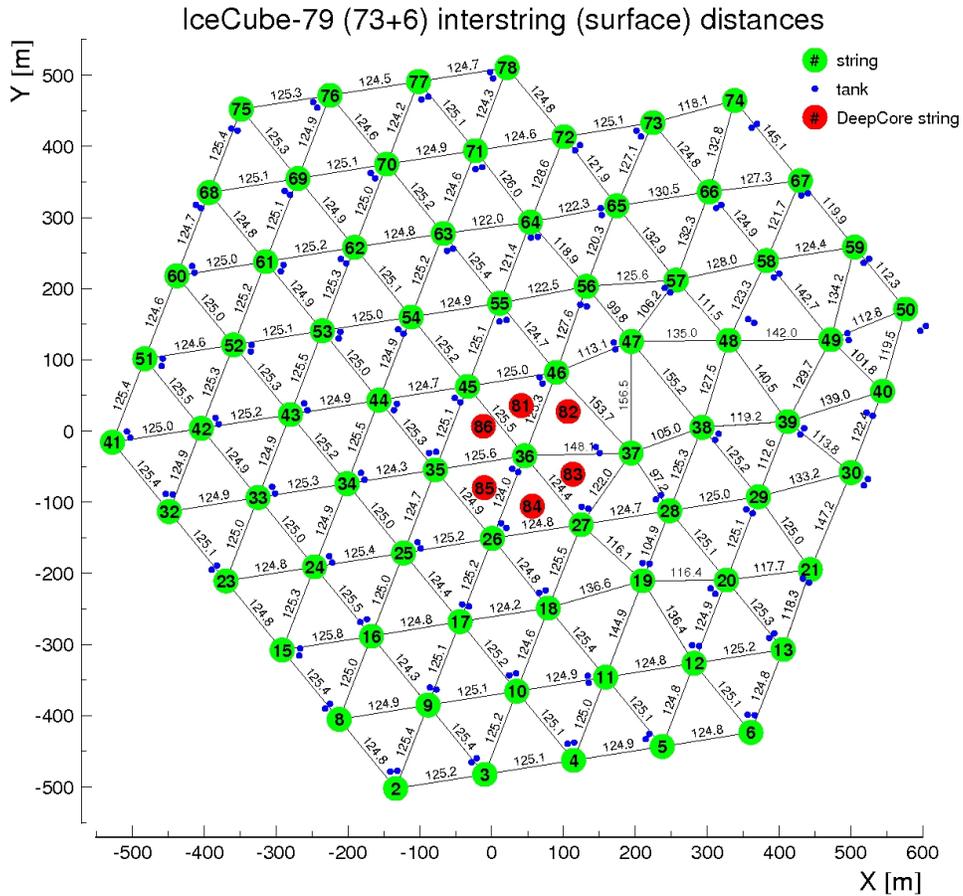


Figure 3.2: The detector layout, with distances labeled between strings, as laid out on the ice. This shows the 79 string layout: the detector configuration from which data was used for this analysis. The title “73+6” refers to the number of strings with regular quantum efficiency DOMS versus high quantum efficiency, used in DeepCore. The red dots denote these DeepCore strings.

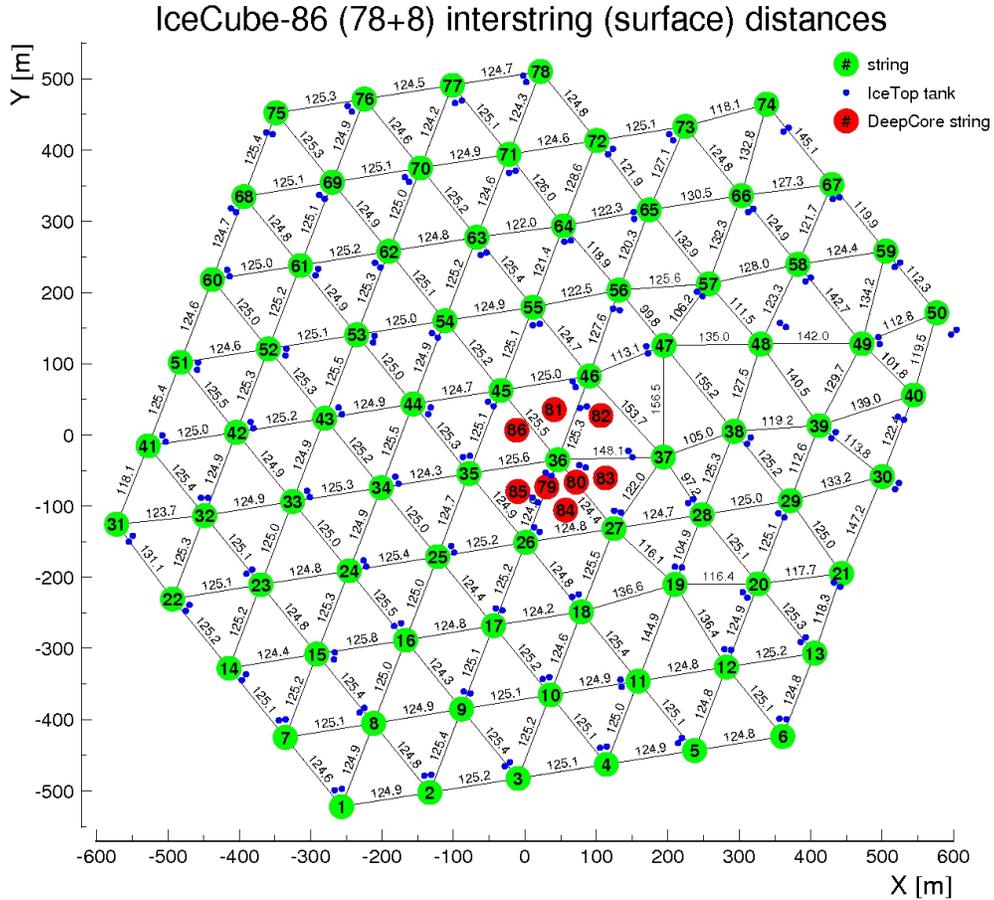


Figure 3.3: The detector layout, with distances labeled between strings, as laid out on the ice. This shows all 86 strings as contained in the full IceCube, even though the analysis in this work uses only data from the 79-string detector (before the last year of construction). The title “78+8” refers to the number of strings with regular quantum efficiency DOMS versus high quantum efficiency, used in DeepCore. The red dots denote these DeepCore strings.

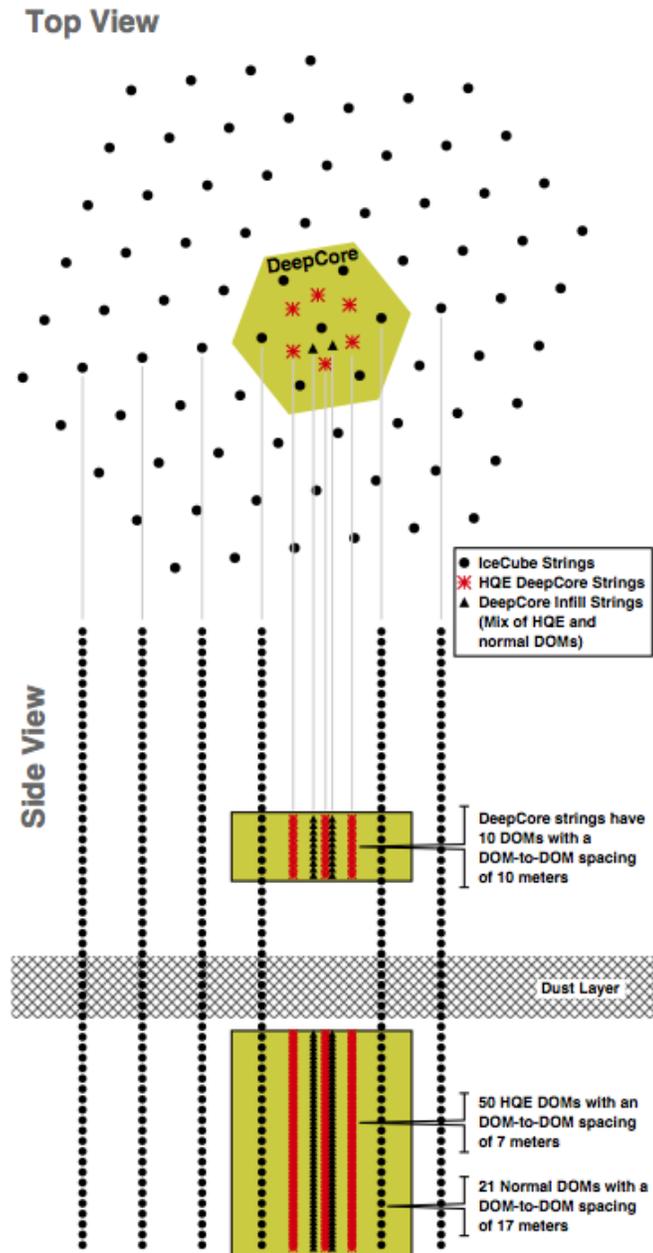


Figure 3.4: A diagram of the layout of the DeepCore detector. The upper part shows a top view, while the lower part shows a side view. The DeepCore fiducial region is highlighted in yellow. Strings consisting mostly of DOMs with high quantum efficiency PMTs are shown in red; strings with mainly regular quantum efficiency DOMs are shown in black. Some parts of the central string (string 36) are also high quantum efficiency. The shaded “dust layer” corresponds to the peak in Fig. 3.5.

was tested in labs before the DOMs were deployed, and was found to be about 20%. For the DOMs in the DeepCore subarray [67], the quantum efficiency of the PMTs was found to be about 27% (about 35% higher [67]). The PMT model used in DeepCore is the same as in the larger detector, but with a modified bialkali photocathode material. The noise rates of the PMTs vary from around 1.7 to 3.5 kHz, peaking at 2.2 kHz for regular PMTs and 2.9 kHz for high quantum efficiency PMTs. This noise rate is reduced by deadtime in the DOM electronics to 500-600 Hz, assuming normal operating temperatures of -50 to -20 °C.

The geometric parameters of the DeepCore detector that impact the reconstruction of neutrinos and vetoing of cosmic ray muons are the horizontal inter-string spacing and the vertical DOM-to-DOM spacing. At low energies (~ 1 TeV or lower), energy losses from the muon are most likely to be through ionization (see §2.4). In this energy range, a muon travels ~ 5 m in ice per 1 GeV. Considering that, the 125 m spacing in most² of the IceCube array is suitable for a minimum neutrino energy of ~ 50 -100 GeV, and an optimal response above 1 TeV. The DeepCore array decreases this spacing by filling in between strings, so the inter-string spacing is reduced from 125 m to 72 m (or, in the full detector with two additional infill strings, which is not used in this analysis, 42 m), extending the visible energies lower into the 10-100 GeV region.

The DeepCore DOMs are centered lower than the IceCube DOMs (with the same lower limit), clustering in the clearest ice at the bottom of the detector. There is a gap on each DeepCore string through the depths of 2000 to 2100 m, where the scattering and absorption lengths substantially decrease compared to deeper and shallower depths (see §3.3). The DOMs from those depths are instead used to create a “veto cap” above the central DeepCore region; this cap increases the veto efficiency against downgoing cosmic ray muons.

IceCube and DeepCore were constructed over the course of several austral summers, stopping construction each winter. This introduces a natural discreteness to datasets: the size of the detector increases in steps at the end of each year’s construction. Each year of partial-detector data is denoted by the number of strings present in the detector during that year, using the format “ICxx” to denote “xx” strings. These values are shown in Table 3.1. This analysis uses IC79 data, which was the last year of partial-detector data. IC79 was also the first year that included DeepCore strings, but it did not yet include all DeepCore strings. The first 6 strings of the DeepCore detector started acquiring physics data in April 2010; that data is the subject of this analysis. The final 2 strings³ were deployed the following season.

² String 47 was nicknamed “pork chop string” during its deployment. At the original planned location, the drill stuck on a crate of supplies that had been flown in during the annual supply air-drop. The parachute of this particular crate had not opened, so it had lodged itself many meters into the snow. The crate contained, among other things, pork chops. That explains why string 47 is off-grid compared to the rest of the detector.

³Originally, strings 79 and 80 were going to complete the corner of the highest-numbered strings. They were moved to the center of the detector to increase the string density, which extends the visible range to lower energies. This was close to the time that DeepCore was funded, so strings 79 and 80 were made with high quantum efficiency PMTs to match.

Data Start Date	Configuration Name	IceTop Stations	Number of New Strings	Number of New DeepCore Strings	Number of Total DOMs
27 Jan 2005 [68]	IC1	4	1	0	76
14 March 2006 [69]	IC9	16	8	0	604
22 May 2007 [70]	IC22	26	11	0	1424
5 April 2008 [70]	IC40	40	18	0	2560
20 May 2009 [70]	IC59	59	19	0	3776
31 May 2010 [70]	IC79	73	20	6	5032
13 May 2011 [70]	IC86	81	7	2	5484

Table 3.1: Comparison of the various configurations of IceCube detector, including the number of IceCube strings deployed each construction year. Each string contains 60 DOMs. Each IceTop station contains 4 DOMs.

3.3 Modeling the Ice Properties

Accurately modeling and reconstructing the propagation of light through the ice sheet is vital to the accuracy of IceCube [71][72]. Information about how light travels is gathered in several ways, and the process of unifying that information into a coherent model has been an ongoing process.

3.3.1 Gathering Ice Data

Laboratory investigation of the ice is obviated by the knowledge that deep South Pole ice is clearer than any laboratory ice [73][74]. Thus, all ice modeling is based on in-situ measurements.

The current ice model is based on calibration data using Light Emitting Diodes (LEDs) within each DOM. Each DOM has 12 LEDs on a dedicated flasher board. During dedicated runs (to ensure that light does not interfere with other data taking) LEDs within specific DOMs are flashed at a wavelength of 400 nm, which closely mimics the 300-600 nm wavelength of Cherenkov light. The light arrival time and amplitude are recorded from all other DOMs. These data are collected at all depths within IceCube.

Data about the ice was also gathered during the deployment process [75]. Laser “dust loggers” were used to measure the scattering of the ice with a depth resolution ~ 2 mm. These logs were taken during the deployment of 6 strings. These data make it possible to construct a map of the ice that includes tilt across the detector. For the depths within the detector, where both flasher and dust logger data are available, the two datasets agree well.

The dust logger data were additionally checked against data from an outside source: the East Dronning Maud Land ice core data [76], taken from latitude 75°S in East Antarctica. The ice core data extends deeper than IceCube. For the depths where both datasets have data, they agree with each other.

3.3.2 Creating an Ice Model

An ice model has three important parameters which vary as a function of depth: the average distance before photons absorb into the ice, the average distance a photon travels before scattering to a direction orthogonal to its original direction, and the angular probability distribution of the photon direction after scattering. These three outputs and one input (absorption length, scattering length, and scattering function, as a function of depth) form a single ice model.

The ice model used in this analysis, called “SpiceMie,” is based on data from flashers within each DOM. This model assumes that the ice properties are a function only of depth and not direction; that is, we assume that the ice is azimuthally symmetric and that the layers of variation have no tilt. This assumption simplifies the data collection process: we can use data from a single string. The SpiceMie model uses data from string 63 (see geometry layout in Fig. 3.2). Furthermore, this ice model assumes a scattering function based on two common functions for modeling scattering with impurities: Henyey-Greenstein [77] and simplified Liu [78]⁴.

To make a fit describing this flasher data, we use simulation. A full simulation of the IceCube detector (see §7.2) has LED flashes introduced into it. A specific hypothesis model is used to simulate flasher data, and predict how that data would look in each non-flashing DOM within the detector. The difference between the simulation and the data, combined across all DOMs observing all flashers, describes the quality of the hypothesis fit. The comparison of data and simulation (with these thousands of data points) is calculated with a likelihood function⁵, described in detail in [71]. Notably, this likelihood function includes a consideration of the statistical errors of the input simulation. The simulation model is then varied until ice model parameters are found that best describe the data. The output absorption and scattering lengths are shown in Fig. 3.5.

To facilitate quick changes to the ice model in this study, photons were directly propagated through the simulated detector, leaving out the lookup tables described in §5.8.

Next-generation ice models consider the possibility of anisotropic ice and tilts in the ice layers. Producing these models can be done with methods similar to producing simpler models. When using these models, however, the complexity of additional dimensions makes reconstruction computations difficult (see §5.8). This is an area of active work within the collaboration.

⁴Liu scattering is a modification of Mie scattering, from which the model “SpiceMie” gets its name

⁵This paper [71] describes the mathematical method of using a likelihood function to compare experimental data to simulation, vary the simulation, and find model values to best describe the data. Identical math can be used for two different goals: varying an ice model to describe flasher data (as described in this section) or varying oscillation parameters to describe observed neutrino rates as a function of zenith angle and energy, as described in Ch. 8. Because of this similarity, the corresponding author of [71] was on the review committee for this thesis.

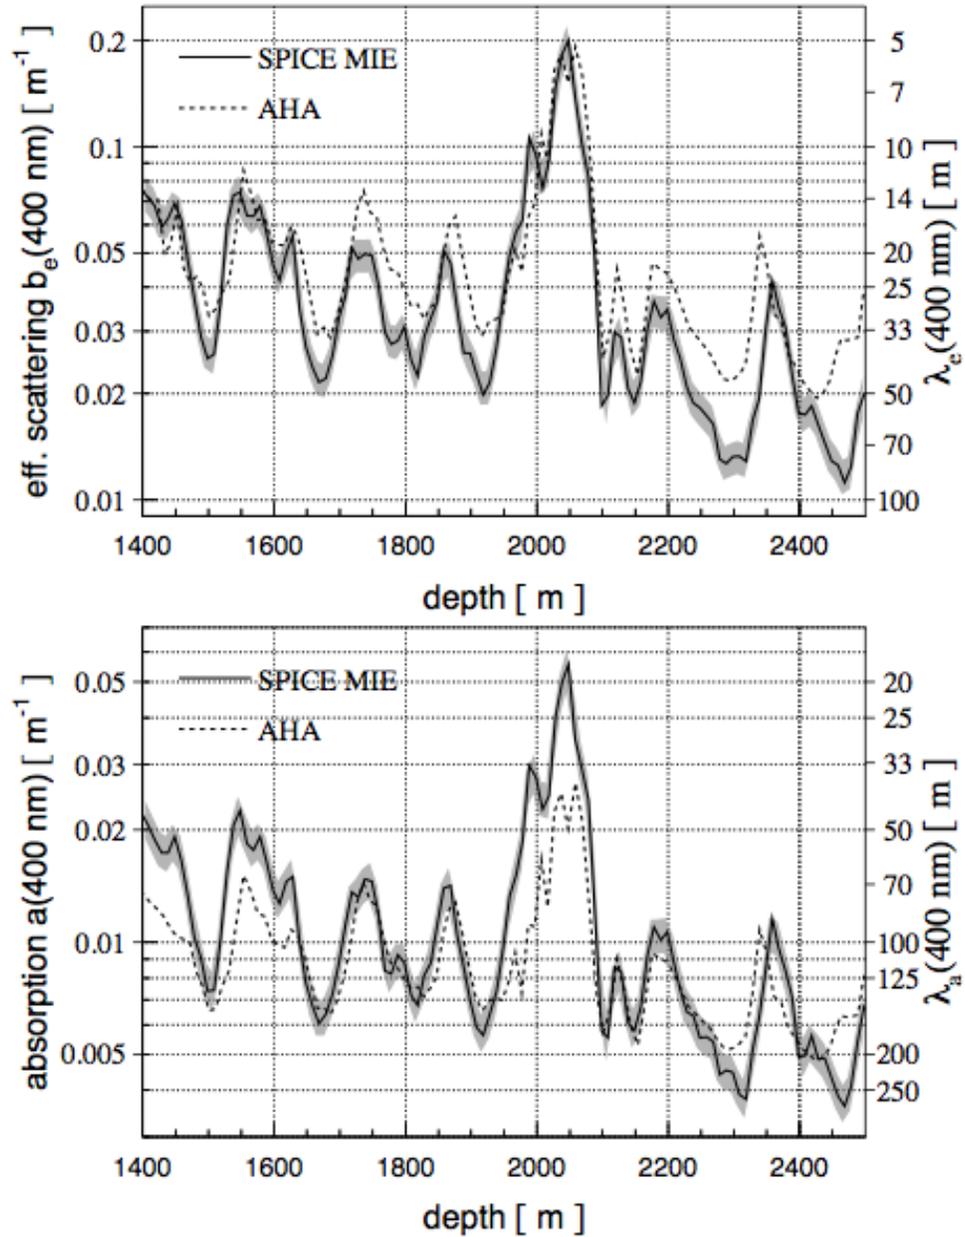


Figure 3.5: Scattering and absorption lengths in the ice as a function of depth. This figure shows the SpiceMie ice model [71], which was the most accurate azimuthally-symmetric ice model at the time of this thesis. The “dust layer” around 2000 m was caused by increased particulate matter in the atmosphere during an ice age $\sim 60\text{-}70$ ky years ago.

3.4 Digital Optical Modules

IceCube is an array of Digital Optical Modules (DOMs) [67], each of which contains a photomultiplier tube (PMT) and associated electronics within a glass pressure housing. The PMTs are 10 inches in diameter (Hamamatsu model R7081). A completed DOM is 13 inches in diameter and weighs ~ 40 pounds.

A schematic diagram and a photo of a DOM are shown in Fig. 3.6. When light strikes a DOM, the metal of the PMT photocathode undergoes the photoelectric effect and emits an electrical charge; that electrical charge signal is then digitized within the DOM before being sent to the surface of the ice.

3.4.1 Digitizing the Signal: ATWD and fADC

IceCube DOMs digitize the PMT waveforms immediately after the PMT produces the waveform.⁶ The DOMs each contain two kinds of digitizer chips [79]: Analog Transient Waveform Digitizer (ATWD) and fast Analog-to-Digital Converter (fADC). The ATWD chips record shorter waveforms with more frequent sampling; the fADC chips record longer with less frequent sampling (see Table 3.2). Including chips with both time scales increases the range of times that each DOM records: we record the first most interesting signals with great detail, with additional signal for a longer time but with less detail.

Table 3.2: Digitizer readout description

Chip Type	# per DOM	Traces per chip	# samples per trace	Sampling Rate	Bin Width	Recording Length
ATWD	2	3	128	300 MHz	3.3 ns	427 ns
fADC	1	1	256	40 MHz	25 ns	6.4 μ s

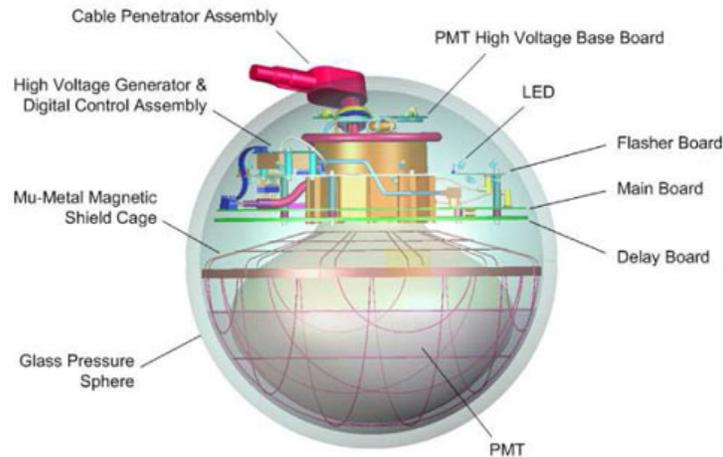
Each ATWD chips runs in 3 channels, each running at a different gain level: 16, 2, and $\frac{1}{4}$. This increases the functional range of signal amplitudes visible to the DOMs, which is limited by PMT saturation rather than digitization.

There are two kinds of signal saturation: PMT saturation and digitizer clipping. When the PMT saturates, the voltage emitted is no longer linear with the input signal. An IceCube PMT saturates at ~ 300 photoelectrons of charge within 25 ns, but this process turns on gradually with increasing observed charge, and is thus difficult to describe with a single turn-on point. The digitizer clipping is much easier to describe: the high-gain settings return their maximum values for several bins. The gain settings for the ATWD chips are chosen such that the minimum gain channel cannot be clipped before the PMT saturates. To illustrate how rare saturation is, consider the example of a ~ 1 PeV cascade

⁶ This in-situ digitization is one of the major differences between IceCube and AMANDA. Any loss in signal fidelity from the digitization step is negligible compared to the signal degradation that would have occurred from sending an analog signal 1 to 2 km over twisted cable with other signals nearby. Because of the experience trying to untangle such signals in AMANDA, IceCube was designed with in-situ digitization.

Figure 3.6: A Digital Optical Module (DOM) shown in both a schematic and photographic view. The harness assembly, including the metal waistband, is included in the bottom figure but not the top. The thickness of the glass pressure housing is more easily visible in the top than bottom, because of clear gel between the glass and PMT, designed to minimize refraction. The cable penetrator assembly (above: pink; below: black) carries the signals from the mainboard to the cable outside the DOM. There are four processing boards, shown above as green lines and below with more extent and visible attached components. The top board controls the PMT high voltage; the second board contains LED flashers and their controls for calibration; the third board contains the main digitizing electronics; the bottom board contains one long wire that preserves the analog signal while a decision is made about the Local Coincidence condition.

(a) A diagram of a DOM. The colors are purposely exaggerated to make components more easily distinguishable.



(b) A photograph of a DOM with its harness for attaching to the main cable.



reconstructed in the upper part of the detector⁷. When IceCube observed such an event, signals were observed in 500 DOMs, of which only 5 displayed PMT saturation. At the energies relevant for oscillation analysis, saturation is negligible.

Combining the signals from the fADC and ATWD chips is performed offline using a non-negative least squares fit with the known shape of single-photo-electron peaks as the basis functions. The quality of this fit relies on accurate calibrations for each DOM. The precision is limited by the PMTs rather than the readout electronics.

Deadtime on each DOM occurs while the analog signals are being digitized. The plurality of ATWD chips used in ping-pong operation reduces deadtime: one chip can record signal while the other chip is being read out. The deadtime is also reduced by requiring a coincident signal in an adjacent DOM before digitizing. The deadtime fraction for each DOM is at most $\sim 1 \times 10^{-4}$ [79].

With these precise electronics and calibrations between DOMs and across strings, IceCube has a timing resolution of less than 3 nsec. With light traveling around 1 foot/nanosecond within the ice, this means that positional reconstruction could theoretically be accurate to within about 1 m, assuming timing resolution is the limiting factor. The uncertainty due to this is small compared to, for example, uncertainties due to light scattering and absorption.

3.5 Triggers and Filters

There are limited options for getting data from the South Pole to anywhere else on the planet. We could put data onto discs and physically carry it away, but this limits the time when we can get data to the times when travel is possible to and from the Pole. We use this method for large sets of data that we do not expect to be particularly interesting, and it has a surprisingly high bandwidth (container ships are big, even if you can only take one trip per year). A better alternative is to send data over a series of satellites: these have minimal lag-time but much more limited capacity.

The first step of this process is the decision whether to write the event to disk and consider it for further processing; this decision is called “triggering”. The second step is to decide if the event is interesting enough to send North; that decision is called “filtering”, and it is based on several different types of data that IceCube participants have proposed analyses around. A specific analysis typically only uses data from a single, or a handful of, filters. The trigger, called “SMT3,” and filter, called “DeepCore Filter,” used for this analysis and many DeepCore analyses is discussed more in §7.3. Any changes to triggers and filters are considered by a collaboration board that weighs the required data transfer rate against the scientific opportunity and merit of the proposal. This is the last stage of processing common to all analysis; further processing is analysis-specific.

IceCube is allocated about 95 GB/day of data transfer to the North (that is, anywhere on Earth except the South Pole). We can send up 105 GB/day at the height of the summer season, when some of the variable rates are at their

⁷ Specifically, this refers to “Bert,” which was event number 14 in [36]

highest, but the average must be lower. Of that 95 average, 5 GB/day are dedicated to operations data, and another 5 GB/day to non-filter data, like supernova scaler reports.

Chapter 4

DOM Testing

All DOMs were tested in three stages: during assembly, at the production site after assembly (§ 4.1), and (for the DOMs that passed those tests) at the South Pole just prior to deployment (§ 4.2). A final stage of “specialty tests” was performed on only a few DOMs after construction was completed (§ 4.3). The specialty tests examined the noise rates and various changes that could effect the noise rates.

The tests during assembly were performed by the technicians who assembled the DOMs. For information about those tests, see [67]; those tests will not be discussed here.

4.1 Production Tests

DOMs deployed in the ice can never be retrieved for repairs, so pre-deployment quality assurance is essential. The purpose of DOM production testing was to determine whether each DOM was suitable for deployment into the IceCube array. When possible, failing DOMs were repaired and re-tested. These production tests were called “Final Acceptance Tests” or “FATs.”

There were three production and testing sites: one in Sweden with contributions from the Uppsala and Stockholm groups, one in Berlin at DESY-Zeuthen, and one at the Physical Sciences Laboratory (PSL) in Stoughton, Wisconsin.¹ The testing procedures were standardized across production sites. The total average production rate was more than 50 DOMs per week (stopping for each polar construction season) between June 2004 and September 2009. In total, 5801 DOMs were produced.

Tests were performed at several temperatures in special dark walk-in industrial freezers, of which the Madison group had three; each freezer holds up to 60 DOMs. DOMs were tested in batches of 60, paralleling the number of DOMs on a string, but test batches were purposely mixed onto several strings during the pre-deployment process to mitigate any systematic effects common to testing batches.

¹ In Wisconsin, between June 2004 and August 2009, 3570 DOMs were produced, of which 3422 shipped to the Pole. In Germany, between September 2004 and September 2008, 1190 DOMs were produced, of which 1168 shipped to the Pole. In Sweden, between September 2004 and September 2009, 1041 DOMs were produced, of which 979 shipped to the Pole.

During production, each DOM was labeled in three ways: by a production number (2 letters and 4 digits), a hexadecimal code mainboard ID number, and a nickname.² As a guideline to limit biased attachment, we avoided naming DOMs after people.

The DOMs within the freezer were set up on individual stands. Each stand consisted mainly of a black PVC pipe with diameter slightly smaller than the DOM diameter, and the freezer held two shelves of stands on each side of a central aisle (see Fig. 4.1). At the bottom of each stand was an optical fiber pointing up towards the DOM. The fibers in each freezer were bundled into a small circular region and all the illumination choices (including laser) were made to create a diffuse spot centered on that circle. Therefore the light getting to each fiber was similar and stable, but not identical. All the optical fibers connected to a single laser that could be pulsed quickly to provide light response tests. When the laser pulsed, it illuminated all DOMs, from which individual responses were read. During testing, each DOM and its stand was covered by a black plastic bag to further protect against light contamination. The electronics in this configuration mimicked the electronics for a single full deployed string in the ice.

Figure 4.1: Photos of DOMs set up for production tests. Each DOM is seated on a black stand and connected to the other DOMs.



(a) The entire inside of the DOM testing freezer in Madison. Two wooden shelves on each side of the central aisle each contain 15 DOMs on stands, for a total of 60 DOMs.



(b) Three DOMs set up for testing. The wooden supports of the upper shelf can be seen at the edges of the photo.

The many temperatures at which the DOMs were tested mirror the many situations in which the DOMs are expected to work well in the field. Tests at room temperature looked for initial operation or connection problems. Tests at -45°C and -20°C mimicked the temperatures at the top and bottom, respectively, of a deployed string. Tests at

² The nicknames had a theme for each testing set, which was chosen by people involved in DOM testing. The themes of the DOM names were varied and often silly. The purpose of this silliness was error prevention: if a particular DOM encountered problems during the testing process, and we discussed it in terms of a silly name, we were more likely to remember the particular name later on the ice if we encountered the DOM again. Some of the themes were Wisconsin breweries, phobias, Berlin S-Bahn stations, species from the cartoon series “Futurama,” places invented in the novel series Harry Potter, words and expressions describing rain used in Oregon, and megafauna from the Pleistocene era; this last theme included the nickname “DemonDuckOfDoom”.

–55° mimicked the winter surface temperatures at the Pole, where exposed IceTop DOMs operate. Each temperature was tested twice, once with temperatures ascending and once with temperatures descending. The only hysteresis noticed was in the freezer temperature, not in any temperature-dependent DOM process. The whole cycle (before data analysis) took 2-3 weeks.³

At each temperature, several tests were performed:

- Each DOM was pinged to make sure it was still connected and the operating system was awake and responding.
- The cable connectors and local coincidence (LC) (see §7.3) circuits were tested in chain through the freezer: light was pulsed onto each DOM, and the LC circuit on the nearest- and next-to-nearest neighbors were tested. This took about 10 minutes.
- A further suite of tests of the mainboard, called STF for **S**imple **T**esting **F**ramework, was performed. This tests individual components like the digitizer baselines, amplitudes, and pedestals; the pressure gauge; the voltages around the flasher LEDs; clock stability and sampling speed; and cross-talk between communications and PMT signals. STF took about an hour.
- The gain on each PMT was calibrated to ensure linearity with the input voltage. This took about two hours.
- The data acquisition system was tested by looking at pulses due to cosmic ray interactions, especially their shapes and the time differences between them. This took between 6 and 20 hours.
- The PMT optics were tested by flashing the laser, at a single photoelectron amplitude level, through the optical fibers to each DOM. This was the primary way of measuring the 2-3 ns resolution described in §3.4.1 and [79], and of measuring PMT afterpulsing. This took about two hours.
- Finally, the noise rates on each DOM were monitored for at least 24 hours with readout every second to test for any unexpected drift or unexplainable spikes in rate.⁴

The results of these tests were then compiled and analyzed by the tester, and an action was recommended for each DOM: ship to the Pole, repair and retest, or fail. These recommendations were reviewed and discussed among all DOM testers in weekly meetings. The measured performance characteristics were also stored in a DOM database.

³This process was greatly simplified by internet access to the freezer controls and DOM mainboards, allowing testing to continue while testers remained outside the lab, e.g. on campus.

⁴Some noise rate spikes were found to correlate with operation of a nearby synchrotron source [80].

4.2 South Pole Testing

The testing apparatus at the South Pole was in a tent at surface temperature (see Fig 4.2). DOMs were stacked on wooden palettes. The tent provided some light protection. The DOM tests were operated from within an adjacent heated shipping container. Many challenges of testing were products of South Pole conditions (e.g., cables becoming stiff from the cold temperatures) rather than unexpected results from tests.

The tests performed at the South Pole were very similar to the tests during production, but with a shorter time scale, especially during the noise monitoring. Specifically, the data acquisition system checks were skipped, the light-based PMT tests were skipped, and the noise monitoring time was reduced by at least a factor of 2.

In total, there were 5484 DOMs deployed. Of those, 5404 DOMs are still operational (98.5%). Specifically, 5081 of 5160 In-ice DOMs, and 323 of 324 IceTop DOMs are still operational. Of the DOM failures, 49 were during deployment or freeze-in (61%), and the rest failed during ongoing operation.

DOM testers were also responsible for helping with the deployment process as needed.

4.3 Specialty Testing

The purposes of these tests were very different from the purpose of the production testing. Whereas the production testing decided, for each DOM, whether the DOM should be deployed, the specialty tests used DOMs to measure a wider variety of quantities. One set of tests measured the dark noise rates of DOMs and PMTs at various temperatures (Sec. 4.3.1), and the other measured internal temperature variations within a DOM (Sec. 4.3.3).

Some of these DOMs will never be deployed. Specialty test DOMs did not have to pass all production tests; for example, a DOM that failed because of a bad flasher board could still be used for noise tests.

4.3.1 Dependence of Noise Rates on Temperature and Glass Presence

In the first set of specialty tests, the noise rates of the DOMS with high quantum efficiency PMTs were compared to the noise rates of the DOMs with standard quantum efficiency PMTs. These rates were measured at many more temperature points than the production tests to test the unexpected effect of higher noise rates at lower temperatures.

One set of tests investigated whether the new model of high quantum efficiency PMTs would be suitable for a large water Cherenkov experiment like LBNE [81]⁵. Specifically, we charted the noise rate as a function of temperature, paying close attention to the region close to and above freezing, where LBNE would operate. These tests used Hamamatsu model R7081-02, the PMTs in the DeepCore high quantum efficiency DOMs. This test was run on 57 DOMs immediately following production tests on the same DOMs.

The procedure for this round of tests was similar to the “monitoring” step of the production tests, with frequent temperature changes. At each temperature, noise rates were measured. Each temperature was measured twice: once

⁵At the time of these tests the project was called LBNE; it has since been renamed LBNEF.

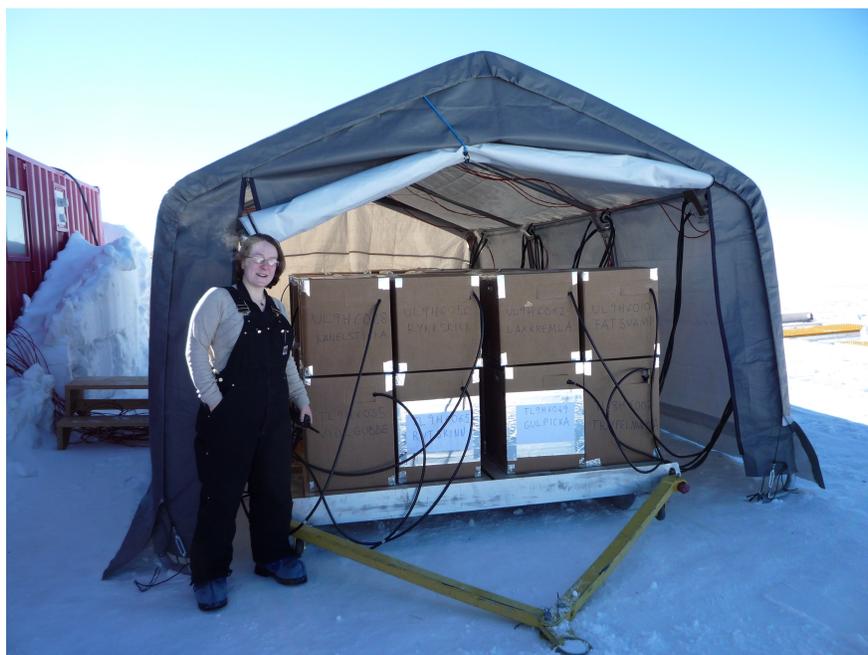


Figure 4.2: Photo of the South Pole DOM testing site. DOMs are contained one each in the cardboard boxes. Eight boxes sit on a wooden pallet. Eight pallets sit on a sled that can be towed out to the drill site. Cables, looped around the tent supports, connect each DOM to the testing computers housed in the shipping container, seen to the left in red. For scale (1.6 m) the author of this thesis is shown to the left.

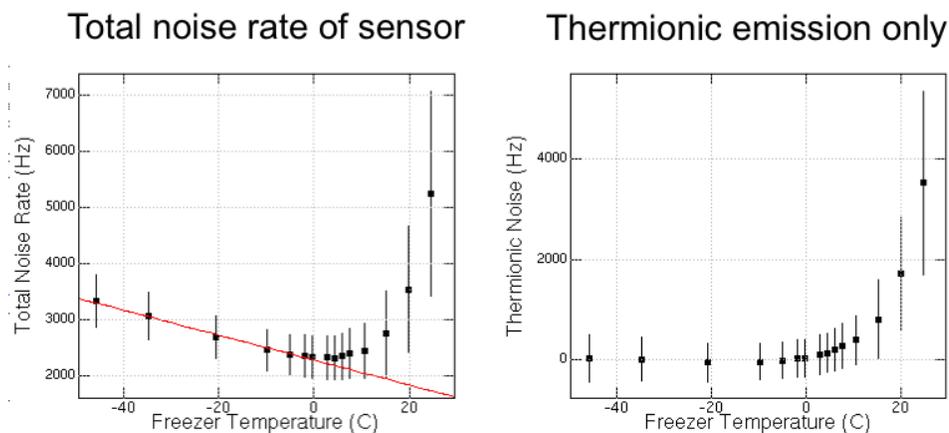


Figure 4.3: Average noise rates as a function of temperature for high quantum efficiency PMTs within DOMs

as part of a descending series and again as part of an ascending series. The results are averaged across DOMs and shown in the left plot of Fig. 4.3.

These results showed a rise in noise rates at low temperatures (below freezing). This rise had been observed also in the production tests, but not specifically studied. It was hypothesized that this effect was primarily a function of the glass pressure housing of the DOMs, and thus would not appear in LBNE-like situations where only PMTs were used. The right plot of Fig. 4.3 shows the observed noise rates with the low-energy noise (a linear fit to the low-temperature region) subtracted off.

We recommended to LBNE that the high quantum efficiency PMTs have 35% higher efficiency, with a manageable increase in noise rate. Balancing the higher efficiency with higher cost was left as a question for the larger community. These results were presented at Fermilab in July 2009.

A further set of tests investigated the low-temperature noise increase. We measured noise rates with varying deadtime after each initial noise pulse. For each temperature, noise rates were measured in 1 second samples for 5 minutes at each deadtime setting. With higher deadtimes, the low-temperature noise increase went away (see any subplot of Fig. 4.4). This indicates that the low-temperature noise comes in bursts, correlated with itself. The noise can thus be separated into above-freezing “thermionic” noise and below-freezing “correlated noise.”

The physical process underlying this correlated noise remains unknown. It is possible to speculate that the noise rates below freezing temperatures are caused by a single, temperature-independent process such as radioactivity in the glass, which is then observed at temperature-dependent rates because of variations in the transmission properties of the glass with temperature. Such speculations were not specifically addressed during the study.

The noise rates were compared between DOMs versus bare PMTs and high quantum efficiency versus normal PMTs (both contained in and outside of DOMs). For these tests, at least one of each type was used: a regular quantum

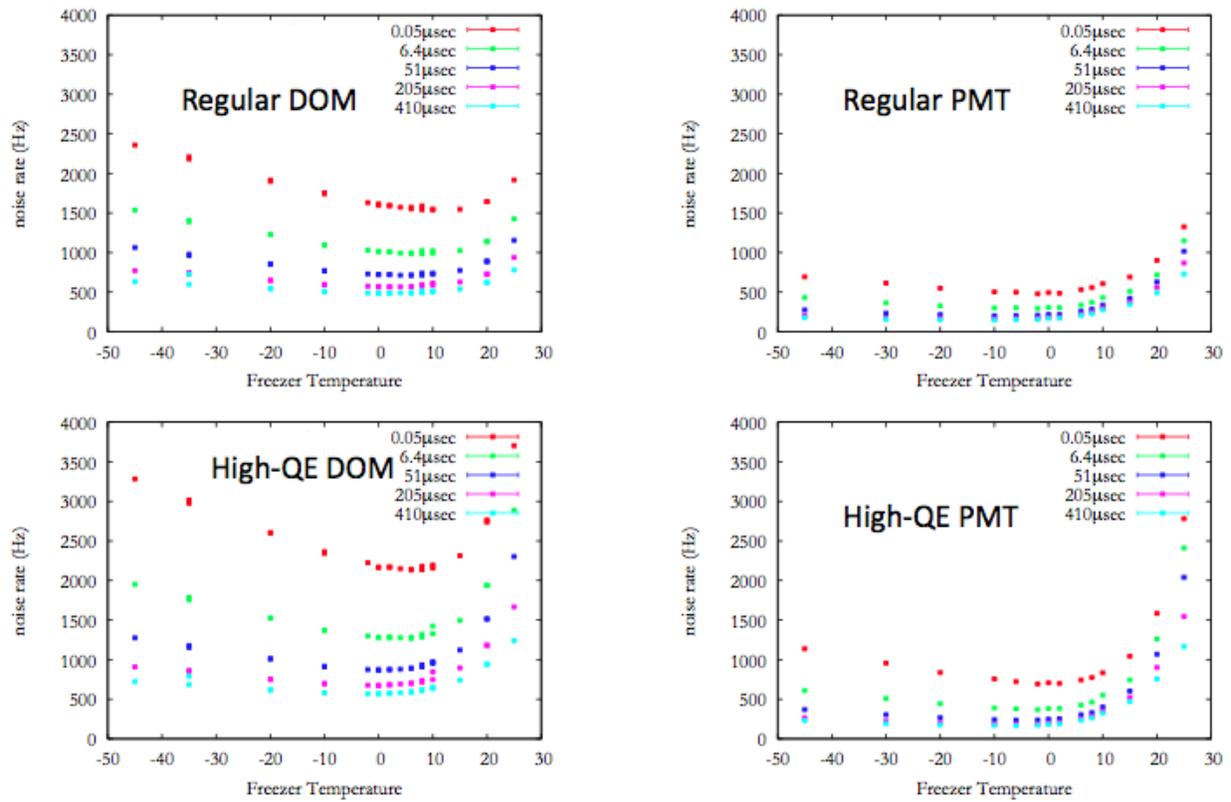


Figure 4.4: Dark noise rates as a function of temperature comparing standard versus high quantum efficiency PMTs and DOMs versus bare PMTs. The colors represent different deadtime settings; lower deadtime allows more clustering in noise pulses.

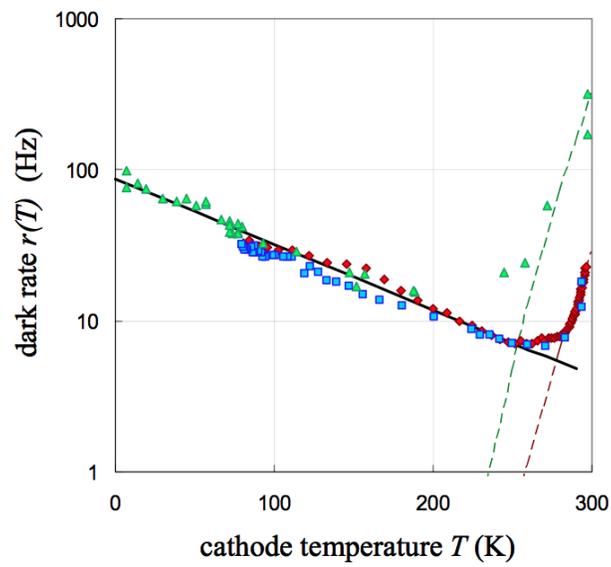


Figure 4.5: Dark noise rate tests on non-IceCube PMTs. The same temperature-dependent noise rate is seen below freezing here and in IceCube PMTs and DOMs. To quote the caption directly from [82]: “Dark rate for two tubes of the same model (R7725). Tube 1 was cooled to 4 K (triangles). Tube 2 was cooled to 80 K (squares) and then warmed back up to room temperature (diamonds). The solid line corresponds to [an exponential fit found in [82]]. The dashed curves indicate the temperature dependence expected for thermionic emission.”

efficiency PMT, a high quantum efficiency PMT, a DOM containing a regular quantum efficiency PMT, and a DOM containing a high quantum efficiency PMT. These tests confirmed that the increasing noise effect is stronger both with high quantum efficiency and with the full DOM housing, and the results are plotted in Fig. 4.4.

A similar effect was documented with much smaller PMTs and a larger energy range in [82], the main figure of which is reproduced here in Fig. 4.5. This paper found that the effect correlated with PMT photocathode area, similar to the IceCube finding that the effect correlated with the amount of glass close to the PMT.

4.3.2 The Vuvuzela Noise Model

Based on the tests in the previous section and other investigations of detector noise, a new software package for modeling noise was developed. The new package was called “vuvuzela;” the old package was called “noise-generator.” The name “vuvuzela” refers to the plastic trumpet-like noisemakers made famous during the 2010 World Cup in South Africa. The major improvement was that vuvuzela includes better modeling of hits correlated with each other in time, i.e., noise that is not purely Poisson distributed in time. While the specific parameters of the vuvuzela model are still undergoing precise calibration at the time of this thesis, the overall improvement is clear: the vuvuzela model matches the experimental data better than the noise-generator model.

Vuvuzela development is documented in [83]. Whenever possible, simulated data sets used for the analysis presented in subsequent chapters use vuvuzela noise.

4.3.3 Internal DOM Temperature Tests

The second set of specialty tests considered whether the noise increase in DOMS could be from differences between the reported temperature in the freezer, the reported temperature on the DOM mainboards, and the real temperature at various points in the DOM. This was investigated by measuring the temperature at various points within a single DOM. The DOM mainboards typically reported temperatures $\sim 10^\circ\text{C}$ warmer than the freezer settings. This study was performed as a cross-check of the study in the previous subsection. Both addressed the question of whether high quantum efficiency PMTs would be useful for LBNE, specifically considering noise rates close to the expected operating temperatures of LBNE. While the study in the previous subsection addressed the temperature dependence of the noise, this study looked for any global shift in the temperature on the PMT that might have been related to the elevated temperatures observed on the mainboards.

The method for the main study was to place 5 temperature probes within a DOM, then wire the signals through the standard DOM communications channels to a circuit board outside the freezer. The voltages could then be read externally, giving a temperature reading for each probe. These readings were then compared to the temperature setting of the freezer, and the results are plotted in Fig. 4.7.

The new equipment for this study was a set of 5 temperature probes. The probes used were model AD590 2-Terminal IC Temperature Transducer from Analog Devices. These probes are manufactured to have a linear current

output of $1 \mu\text{A/K}$. The current was measured with a standard voltmeter reading across resistors on a small circuit board.

These probes were calibrated in a series of measurements before the freezer tests began. They were used to measure a wide range of temperatures, in parallel with measurements from a mercury thermometer and an alcohol thermometer. The two thermometers agreed with each other to within their readout errors, $\sim 0.5^\circ\text{C}$. The temperatures tested are shown in Table 4.1. The resulting measured voltages were fit with a linear function of temperature. The results found were that after calibration the temperature probe reports agreed with the thermometers and each other to within 1%. This best fit linear calibration function was used to run the main freezer tests. The fit we found was consistent to within half a percent with the predicted linear temperature dependence specified by the manufacturer.

Temp $^\circ\text{C}$	Circumstance
-31	small DOM freezer in Chamberlin Hall
-19	food freezer
23	room temperature in the hallway
25	room temperature in the office
30	room temperature in the lab with servers
40	beaker of sand heated in a microwave then stirred
93	beaker of sand sitting on a hotplate

Table 4.1: Temperatures at which temperature probes were tested, with the circumstances used to create those temperatures.

The probes were installed in 3 locations within the DOM (see Fig. 4.6): next to the temperature probe on the DOM mainboard, touching the PMT, and touching the glass housing within the potting gel. Additionally, one probe was fastened to the outside of the glass housing, and one probe was allowed to hang free within the freezer. These 5 probes were wired to the standard DOM communications channels.

The freezer temperature was varied and measured. The results are shown in Fig. 4.7. As expected, the probe on the mainboard read consistently higher temperatures than the probes in other locations. The differences between the measurements on the DOMs and free-hanging measurements were larger at lower temperatures. The expected result was that the temperature would vary monotonically from the mainboard out to the free-hanging probe. An unexpected result was found: the probe touching the PMT was the closest in temperature to the free-hanging probe. The other three probes (on the mainboard and either side of the DOM glass) followed the expected trend.

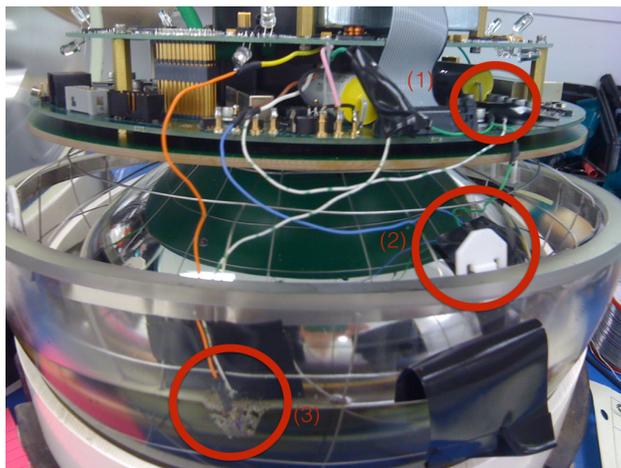


Figure 4.6: A photograph of the temperature probe locations within the DOM. The locations of the three internal probes are marked with red circles: (1) on the DOM mainboard, (2) on the PMT and (3) within the potting gel touching the glass housing.

Conveniently, this means that the PMT temperature is closer to the temperature of the freezer than to any other temperature within the DOM, and so measurements taken at various freezer settings on many DOMs are directly applicable to other experiments (e.g., LBNE) that use PMTs but not DOMs.

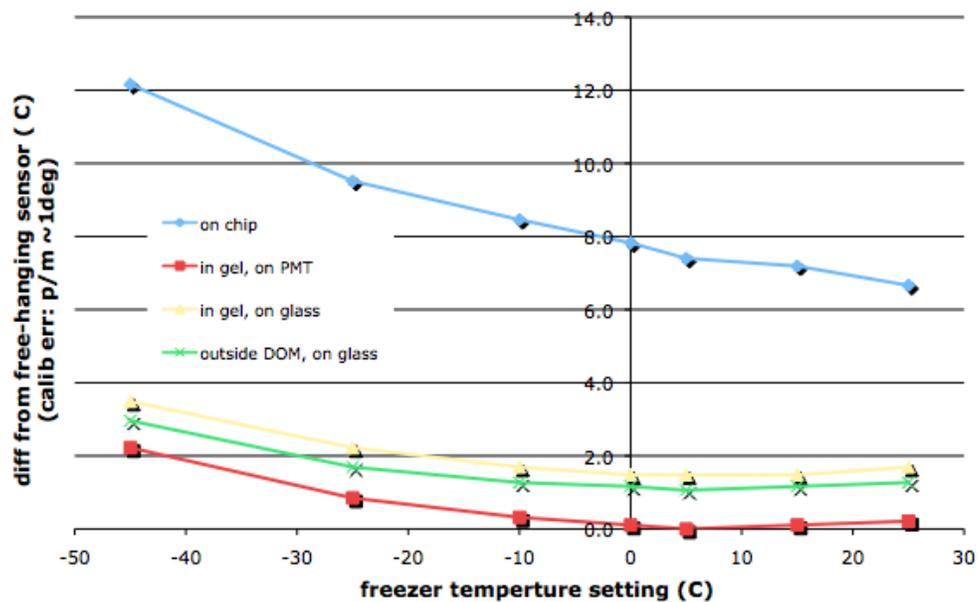


Figure 4.7: Temperature dependence of the temperature measurement deviations between various points on a DOM compared to the ambient freezer temperature.

Chapter 5

Reconstruction Terminology and Methods

This chapter describes all the reconstruction algorithms used in this thesis. They are organized roughly by how long it takes to run each algorithm, with the fastest first. This is also the order in which the reconstructions are often run, with quality cuts to reject some events in between each reconstruction algorithm. Thus, each cut level (running a reconstruction and then cutting on the results) takes a similar order-of-magnitude amount of processing time and/or disk space. This reference of some common reconstructions may be useful to future IceCube scientists; reconstruction or veto algorithms specific only to this analysis are discussed in Ch. 7. For brevity, reconstructions not used in this thesis analysis have been omitted, without intending any offense to their authors and maintainers.

Software to run all of these reconstructions can be found within the IceCube software repository¹, along with code for fitting the resulting sample for neutrino oscillations searches. These work within the IceTray software system [84]. Large simulated data sets are generated with the IceProd [85] software system.

5.1 Basic Terminology

This section explains some common basic variables used to describe IceCube events. As a charged particle moves through the ice relativistically, it emits Cherenkov light. This light is detected by PMTs which output “waveforms”, which are then digitized on DOMs to make the signal into “pulses,” each of which has an integral measured in “charge”. The signal is fit using a least-squares fit, using a template of a single photoelectron pulse as the basis function. The number of basis functions needed to fit the waveform has only a tenuous relationship with the amplitude of the original waveform, and so while “NPulses” is a commonly used variable [86], it is not in general physical useful. Total reconstructed charge or “Q_{Tot}” is both more accurately correlated with the waveform integral and the total energy of the event.

Information from many DOMs is combined into an “event” which typically lasts several microseconds, determined by the time window in which several DOMs have signals. Each event is saved as a discrete set of data called a “frame,” which incorporates the event-specific detector responses with detector geometry, calibration, and status information.

¹<http://code.icecube.wisc.edu/svn/sandbox/gladstone/OscillationFitter>. The reconstructions described in this chapter were developed through broad collaboration effort; even though the software cited here calls many reconstructions, the author does not claim credit for having developed them. However, the author does claim credit for modularizing and commenting the scripts, with the hope that others will not have to reproduce the effort of researching reconstruction settings.

The true interaction is the event, and the digital datum is the frame. Part of IceCube software development is confirming that each frame contains one event, and describing how often and what ways a frame contains multiple events or no events.

The number of DOMs that have signal over threshold in an event is called “NCh.” The same term can be used to describe only the DOMs with local coincident (see §7.3) hits (most common at high energies) or all DOMs, including isolated hits (most common at lower energies). It is usually preferred that one specify which hit cleaning, if any, was used when reporting an NCh value.

The word “hits” is often used in an ambiguous way, and is best understood as a casual term. It can sometimes refer to the number of hit DOMs (NCh), sometimes the number of pulses, sometimes the sum of charge on each DOM rounded to the nearest integer and added up across the detector. It’s best to specify with each usage which meaning is intended, or interpret it as a non-specific term related to NCh.

5.2 “I3Particle” Definition

A common datatype within the IceCube software is the “I3Particle”, which is composed of several fields. Because these fields are commonly used together, the I3Particle is used for many different things: usually to represent particles within the simulation, or to represent the outputs of various reconstruction algorithms. Different applications may use the fields with different meanings. For example, “Time” may mean the neutrino interaction time in the ice for a reconstruction, or the interaction time of the parent cosmic ray shower in parts of the simulation. This definition is provided here as a template for the output of the various reconstructions that follow in this section.

For reference, the data fields of an I3Particle are:

- Position: (x, y, z) measured in meters from roughly the center of the detector.
- Direction: (zenith, azimuth) measured in radians.
- Time: usually measured in nanoseconds, with varying zero-point.
- Energy: measured in GeV.
- Length: measured in meters, but not always filled out.
- Speed: canonically measured in meters per nanosecond, but not always filled out.
- Shape: This is used to tell some reconstructions how to treat different classes of events differently. Examples include Cascade, Track, and Contained Track.

- `particleType` and `pdgEncoding`: these are used to describe particle classes within the simulation. Examples include muon, positron, helium nucleus, and electron anti-neutrino. The “pdg” part refers to the Particle Data Group list of particle types [87].
- `locationType`: Examples include `InIce`, `IceTop`, and `Anywhere`.
- `Major` and `Minor ID`: these numbers are used to track individual particles within the simulation.

5.3 Hit Cleaning: SRT and TW

Hit cleaning algorithms are included before or within many reconstructions. The idea of hit cleaning is to make a guess about which hits were created by the same particle interaction, and are thus interesting to analyze together. Hit cleaning can be run extremely quickly.

There are two kinds of hit cleaning used in the event selection of this analysis. The first is “time window” cleaning, abbreviated as TW. This cleaning, as the name suggests, selects hits within a specified time window before and after the trigger condition is met. For example, a low-energy event might have many microseconds before and after the SMT3 trigger condition is met, but that time is mostly filled with noise. To improve the signal to noise ratio, TW cleaning cuts down on the time before and after the trigger.

The second common kind of hit cleaning is called “RT” cleaning. The “R” stands for “radius;” the “T” stands for “time.” This cleaning uses the idea that hits clustered closely together in time and space were probably created by the same process (in that sense, RT cleaning offline is similar to the LC condition used in online triggering). The algorithm works by searching for any pairs of hits that occur both within time T of each other and radius R. Depending on the configuration values of R and T, this usually lets in more hits than local coincidence (see §7.3) cleaning.

The results of two hit cleaning algorithms are shown in Fig. 5.1.

The danger of hit cleaning, especially RT cleaning, is that it can never be perfect. It is often useful, but it can never be perfect. Consider the case of a high-level reconstruction that has noise built into its likelihood function. The input to a fit is a “hit series”, a collection of pulses at specific times on specific DOMs. The analyzer has the choice to feed a hit series to the reconstruction either before or after hit cleaning. If the choice is to use hit cleaning, the assumption is that there is no noise remaining in the hit series, and the likelihood should reflect that. If there is no hit cleaning, or only some hit cleaning that is easy to build into the likelihood function (like TW cleaning), then noise can be built explicitly into the likelihood, as it is in Millipede-based reconstructions (for a full discussion of Millipede, see the following §5.11). If noise is cleaned out imperfectly and the likelihood has no knowledge of that noise, it will interpret those extra hits as signal, and return misreconstructed energies. Having fewer total hits also contributes to a less smooth likelihood space, making minimization more difficult. Any of these outcomes is less precise and more error-prone than giving the likelihood a hit series with some noise intact.

Despite that caveat, many simple reconstructions are greatly improved by including some hit cleaning; see for example Improved Line Fit, §5.6.

5.4 Topological Trigger

The Topological Trigger splitter algorithm [89][90][91], or “TTrigger,” is a next-generation form of hit cleaning. It attempts to split every event into sub-events. This is an attempt to eliminate a particular class of mis-reconstructed events: those which are composed of multiple downgoing events, precisely timed in a way that confounds directional fits.

The TTrigger algorithm works as follows. Each hit within an event is considered in relation to all other hits. Two hits are considered “topologically connected” if the pair meets several criteria:

- both hits are in the in-ice portion of the detector
- if the hits are on the same string, they are within 30 DOMs of each other
- if they are on separate strings, those strings are within 500 m of each other
- the residual time between them ($\delta t - \delta r/c$) must be less than 1000 ns

Topological connection is considered transitive. Topologically separate groups of at least 4 hits each are considered “split” by TTrigger into separate events, and each is passed to the regular trigger system. This splitting system is an effective, fast way to search for and split coincident events.

5.5 Noise Engine

Noise Engine [83] is an algorithm designed to identify and throw out events that probably include DOM noise only, and not any signal hits. It was one of the first projects developed using the “vuvuzela” noise model (see §4.3.2).

The algorithm is based on the older “track engine” software. Both are based on the idea that hits from a traveling particle should have net direction between them, but random combinations of noise should not. The algorithm constructs a vector from each pair of hits, adds all the vectors together, and then examines the magnitude of the resulting sum. If the result is close to zero, it is more likely the event was pure noise, and Noise Engine cuts that event. Noise Engine also includes several checks to avoid throwing out large or otherwise interesting events, and many of the parameters are configurable. This means that even though it was originally optimized to be run as a very low-level cut for noise-only triggers, it may also be useful in some analyses at higher levels.

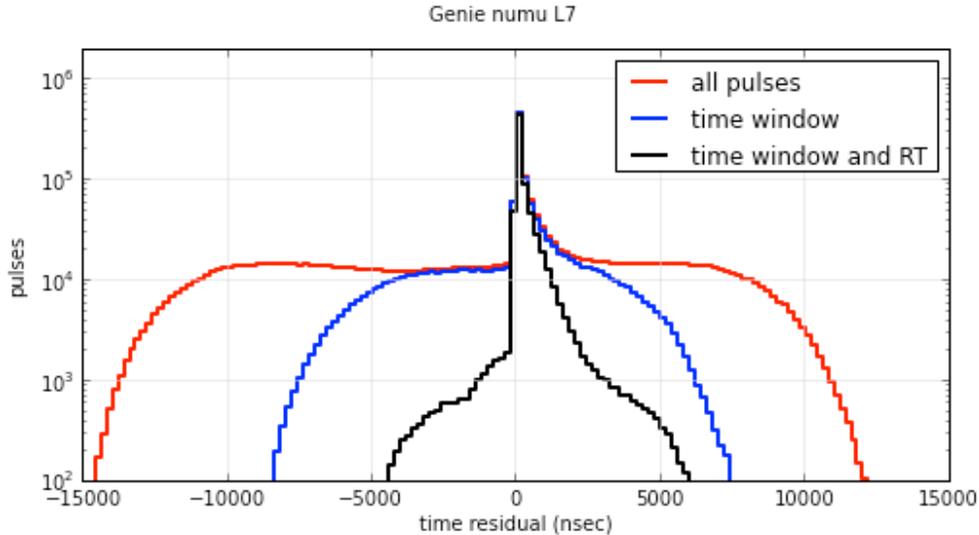


Figure 5.1: The rates of reconstructed pulses, with various cleanings applied, versus the *time residual*. The data sample used here contains charged current ν_μ events generated with Genie Monte Carlo generator ([88], see §7.2) and then processed to the L7 set of cuts defined in Ch.7.

To calculate the residual, consider the Monte Carlo true position of the muon emitting light. We can calculate the minimum time for light to travel along the path of the muon (starting at the true Monte Carlo interaction time), depart at the Cherenkov angle, and arrive at the position of the DOM where the pulse was observed. The difference between that minimum arrival time and the time the pulse was observed is called the time residual, and is assumed to arise from scattering in the ice. For pulses that are truly emitted by the muon, the time residual must be positive, unless the muon position is incorrect, which we ignore here because this plot uses Monte Carlo truth. Pulses from random noise have uniformly distributed residuals.

The three series shown here in different colors represent the results of different hit cleaning algorithms:

- The red curve, with the highest rates, shows the results with no hit cleaning; as expected, this has broad shoulders showing a significant contribution from random noise.
- The middle curve, in blue, shows the results of applying a time window cleaning, which cuts away all pulses outside of a specific time window relative to the trigger time. The time window was the one used in this analysis (for the pulse series used in the “monopod” reconstruction algorithm): $4 \mu\text{sec}$ before the trigger until $6 \mu\text{sec}$ after the trigger.
- The lowest, black curve additionally uses an “RT” cleaning, which keeps only pulses within a certain radius and time of another pulse. The radius and time settings were the ones used in this analysis for cuts requiring a minimum number of hits: $R = 150\text{m}$ and $t = 1\mu\text{sec}$.

5.6 Linefit and Improved Linefit

This is one of the simplest fitting algorithms used in IceCube. Linefit is often used as a seed for more computationally intensive fits, and works best for extended tracks.

Linefit uses a least squares minimization to find the track that minimizes the sum of the squares of the distances between the track and the hits. This is also the best time, position, and direction for a plane wave passing through the hits in the detector. Considering the 4 dimensions of x , y , z , and time, this uses the same math as fitting a line to points on the 2D graph, except that the 2D graph is 4D, and the line is a plane wave. This has the downside of being fairly inaccurate except in the very simplest of cases, and the upside of being extremely fast and easy to understand. Making a new software implementation of Linefit has become a classic end-of-bootcamp software exercise for incoming students, generally with great success.

Linefit has a well-defined meaning of speed: the speed at which the plane wave moves through the detector. This speed has proved a useful cut variable: muons travel on average close to the speed of light, while contained or starting events average lower speeds (because their light travels outward in many directions).

For many online applications, we use a modification called “Improved Line Fit” which includes rigorously studied hit-cleaning and other optimizations [89].

5.7 CLast

The CLast algorithm was originally designed as a fast seed for cascade-type fits. It combines several quick estimates, providing an estimate for each field of the software type “I3Particle”. The name refers to the fact that it is an improvement over the CFirst cascade seed algorithm, specifically changing the method for estimating the time. The name “CFirst” was chosen as a shorthand for a first-guess cascade reconstruction.

CLast returns an I3Particle with the following fields:

- Position (x , y , z): the charge-weighted average of all observed pulse positions. This kind of charge-weighted average is often called the Center of Gravity, or “CoG”.
- Direction (zenith, azimuth): the major axis of a tensor of inertia (ToI) of the input pulses.
- Time: the most interesting field in CLast, as described in the remainder of this section.
- Energy: a polynomial function of the total charge, with coefficients optimized at some point (and configurable).
- Shape: always set to Cascade.
- particleType: always set to Unknown.
- FitStatus: always set to OK.

- Length: not set.
- Speed: not set.

The method for determining an interaction time is as follows. Three time calculations are attempted. If the first succeeds, it is returned as the result; if not, the algorithm looks for the second. The third and last method always succeeds, even if the first two do not. All of them start with the hypothesis that the CoG position was the interaction vertex.

The third choice, called the “Earliest Vertex Time” is the simplest. It always exists, so it is the last backup. For a hit DOM (then looping over all of them), the earliest pulse time is taken. Based on the distance between that DOM and the CoG, that time is moved back to what time the light would have left the CoG to arrive at the DOM: that is the “vertex time” for that DOM. Having looped over all DOMs, take the earliest vertex time.

The first choice is called “Earliest Threshold Time”. Similar to Earliest Vertex Time, it finds a vertex time for each DOM that had hits. The variation is that, out of those various times, it only considers times from DOMs that met a certain threshold condition. That condition is a minimum number of direct hits. The specific number is configurable by the user (the default is 3). This ensures that the candidate vertex times meet a certain minimum fit quality before they are considered, and is also the earliest time that meets that threshold quality.

The second choice is called the “Time of Max Direct”, which is also an extremum of the set of vertex times based on the first time measured in each DOM. In this case, it selects the vertex time with the most direct hits.

5.8 Describing Ice with Photonics Tables

The most advanced reconstructions incorporate detailed knowledge of the properties of the ice and how light travels through the ice. More discussion about how we measured the ice properties can be found in §3.3 and [71][72]; this section describes how we use that description in software. This software, which creates, describes, and accesses “photonics tables,” is the interface between ice measurements and event reconstructions. Given a point in the detector where we hypothesize the light came from, and a DOM we hypothesize the light arrived at, we can read the probability of that hypothesis from the photonics tables. These tables translate the ice property models described in §3.3 into a format we can use for simulation and reconstruction.

Tables are generated by simulating many photons and tracking them individually until they either are detected, absorb into the ice, or leave an extended volume around the detector. We run enough photons to populate even the far bins of the tables, typically 7.7×10^{10} . This process takes about 35,000 CPU-hours, and we try to run it on a cluster with several hundred cores. The NPX4 cluster that IceCube runs in Madison has about 2400 cores, of which any one user can expect a consistent few hundred, so generating tables takes at least a week. Using the standard Hofstadter’s law [92] expansion, this is more accurately predicted, with debugging, as a month.

These tables are set up with 6 dimensions, describing the position of the emitting particle and the detecting DOM. Those dimensions are:

- depth of the emitting particle
- zenith angle direction at which the particle emits
- polar angle between the emitting direction and the detecting DOM
- azimuth angle, relative to the emitting direction, of the detecting DOM
- distance from the emitter to the detecting DOM
- time between emission and detection

This makes the tables quite large: up to 4 TB for a full set, including cascade and muon light-emission models and fine binning. To run a fit (historically) one needed to load the full tables into the computer's RAM, and few computers exist with that much RAM. We have one in Madison, and used it to execute the fit described in the subsequent paragraphs. Additionally, as a historical note, these tables also had interpolation problems between bins in each of the six dimensions: any binning or linear interpolation between bins causes systematic errors. A more coarsely binned set might be only 20-50 GB, but these were typically crippled by interpolation errors. So, two problems need to be solved: the size of the tables in RAM needs to be reduced, and the interpolation between bins needs to be improved. These can both be addressed by using a fit to these tables instead of the tables themselves, and this fit was developed by Madison grad students during the time of this thesis. An operational overview of this fitting process is presented in the subsequent couple of paragraphs; a more complete and mathematical discussion can be found at [93][94].

The fitting process that was settled on is called "spline fit", or "splining" for short. This process uses a series of quadratic functions to fit the tables, with certain set points being fixed to the table values. The fixed points are called "knots". The etymology around splines derives from their original usage in engineering and navigation. "Spline" was originally a thin strip of flexible wood, and it was literally laid out of a physical plan or map, and pinned at specific points; the form it assumed was the form with least stress throughout the rest of the curve, because the wood would naturally minimize points of stress. The engineer could then simply trace the resulting curve. The numerical fitting algorithm is a digital version of this process.

The spline fit runs on a non-negative linear least squares fit, a common theme of the last few years at IceCube. The basis functions of the fit are the spline shapes, which vary to minimize the distance to the histogram data. The fit runs by loading a full set of photonics tables (which, as discussed earlier, could be multi-tera-byte, so fitting a single section requires 48 GB of RAM) and fitting using well-established algorithms that other people have optimized. When the fitting run is done, a graduate student, by hand, looks through several projections of the fit onto a single

dimension, ensuring the quality of the fit. The tools for these cross-checks were designed and first used by Jakob van Santen. In DeepCore and PINGU, where low energies are more important, the regions of the fit close to the DOM are more important, and more fitting knots were added by hand to improve the fits in those regions. These low-energy cross checks were recently performed by Moriah Tobin. The splining process, if there are no errors and no rerunning required, takes 4-6 hours per file, and runs on 1540 files, for a total of 7.7 kCPU-hours, or a couple of days on our cluster.

Making lookup calls to these tables is often the time-limiting step of a reconstruction algorithm, so a common method for optimizing the speed of a particular reconstruction is to strategically reduce the number of photonics lookups. For example, if we are reconstructing events that only occur within the DeepCore region, and they are known to have energies below around 100 GeV, we could use a reduced geometry to reduce the number of lookups. If we can assume there were no hits in the outer regions of the detector that were related to the event in question, we can ignore those regions, and avoid looking up those DOMs.

5.9 The Gulliver Software Suite

Gulliver is the framework in IceCube for running likelihood fits. It divides the algorithm of running a likelihood fit into several independently defined parts. By keeping these parts independent with well-defined interfaces between them, changing one part at a time is made easier. Thus, for example, a minimizer algorithm can be swapped out without fear of introducing typos into the likelihood function.

Many of the energy reconstructions described in this chapter use methods described in [95].

Gulliver gets its name from the literary character Gulliver the traveler, because it wanders around in the parameter space, minimizing things from – not kidding – the Lilliput library. The names, original code, and many years of maintenance are from IceCuber David Boersma.

A Gulliver fit has 5 parts:

- **Fitter Module:** this is the only one that is a module (the rest are services). The purpose is to tell the services about each other, and tell IceTray where to run the fit. Since this does not have much specific content, we can all use the same ones. I3SimpleFitter is the almost-exclusive standard; I3IterativeFitter and Paraboloid are common exceptions.
- **Minimizer Service:** the machinery that you probably do not want to mess with. The minimizer takes a series of input parameters and varies them until the output is maximized; the Simplex minimizer in Minuit2 [96] is what we usually use.

- **Seed Service:** provides a start to the minimization process. It starts with a name, and then uses the dictionary within each frame to translate between that name and a frame object. Types of frame objects include I3Particles, I3ParticleVectors, and I3Doubles. This process is quite simple.
- **Likelihood Service:** this is the first on the list that physicists could spend time developing. It translates an I3Particle-based configuration to a “goodness” number, e.g. a χ^2 . Examples all have “likelihood” in their names. This could use light propagation information, spline tables, or Pandel (see §5.10 or [97]) functions. Most of the physics that goes into describing a fit goes into the likelihood function.
- **Parameterization Service:** two functions to translate (a) back and (b) forth between an I3Particle-based configuration that the likelihood understands and a set of doubles that the minimizer can deal with. Minimizers only understand vectors of parameters, not I3 types. Examples: I3Simple parameterization (which only uses one particle), and MuMillipede (which uses a vector of particles; see §5.11.2).

The “parameters” which the minimizer sees can just be the reconstruction variables, e.g. $\text{par1}=x$, $\text{par2}=y$, $\text{par3}=z$, $\text{par4}=\text{zenith}$, $\text{par5}=\text{azimuth}$. However, it can be good to define the parameters with some transformation of the physics variables, such that $-\log(L)$ behaves more smoothly as function of the parameters, and/or the full range of the parameters can be used (no bounds). Examples: a conversion between E and $\log(E)$, projecting a hemisphere onto a flat map (for direction).

This could be useful if, for instance, the I3Particle-based configuration space has some poles you want to avoid. The method “hypothesis” translates from the seed (an I3Particle-based configuration) to the new variables, and the method “update physics variables” translates the other way: from the physics variables to the I3Particle-based configuration.

The process of iteratively minimizing can be abstracted to the process of Minuit2 iterating between the Likelihood and the Minimizer, with the Parameterization to translate between them. Maintaining the modularity of these fits makes it much easier for code maintainers to compare and debug code. Any hassle or additional debugging that a reconstruction author may be required to do while integrating a new reconstruction into Gulliver should be considered an investment towards future ease of maintenance and debugging.

5.10 SPE and MPE

SPE and MPE are two kinds of Gulliver-based likelihood fits that do not use the ice model directly. “SPE32” stands for **S**ingle **P**hoto **E**lectron reconstruction iterated **32** times. “MPE” is a modification of SPE; it stands for **M**ulti **P**hoto **E**lectron.

Both use the Pandel function to describe light propagation probabilities. This uses a parameterization [97] from the Baikal neutrino observatory [98]; the author of the parameterization lends his name to the function. Pandel functions

do not include a depth dependence or detailed information about the ice. This makes it unnecessary to repeatedly look up information in photonics tables, so Pandel-based fits run faster than the reconstructions described later. Knowledge of the ice is included in a general sense: the specific values of the Pandel functions used in IceCube [99][100] are fit to photonics tables, creating a depth-independent “bulk” ice model.

SPE considers the first observed photon from each DOM that had a signal. For each of these hits, it calculates the probability of a photon traveling from the hypothesis muon track at the Cherenkov angle and arriving at the DOM at that time, according to a Pandel likelihood function for a single photon. The product of all these probabilities is the likelihood of that hypothesis track. A minimizer is used to vary the track and find the most likely direction, interaction time, and interaction position. This process is repeated 32 times to avoid getting caught in local minima of the likelihood space. The seed directions for the 32 iterations are selected via a 2-dimensional Sobol sequence.

MPE is a modification of SPE. Specifically, it uses a modified Pandel likelihood function for a single photoelectron (PE). It first counts the total number of PEs that arrived on the DOM, and modifies the Pandel function to include that many PEs, still basing the start time of the function on the arrival time of the first hit. This modification shifts the most likely value of the Pandel function to later times. The rest of the reconstruction procedure continues similarly to SPE, except that a single iteration of the minimizer is used.

The “paraboloid errors” are a description of the likelihood space near the best-fit value of either fit. We used the paraboloid errors around MPE. The paraboloid algorithm uses a minimizer to map out the region around the best fit where the significance drops off by 1 sigma. This is based on the assumption that a region of the likelihood space with a sharper peak gives a higher confidence in the best fit within that peak. Once that region is obtained, the directional components (zenith and azimuth angles) of that region of likelihood space is described with an ellipse. That ellipse is the result. The magnitudes of the major axes of that ellipse can be used as a proxy for reconstruction quality.

There are two common misconceptions about SPE and MPE. One is that they depend on a depth-dependent ice model. They do not; the “bulk ice” model used in Pandel functions is an average over measured IceCube ice properties. The Pandel likelihood function is based on light propagation in Lake Baikal, and re-scaled to fit AMANDA scattering data, but it does not vary with the depth in the ice. The second misconception is that MPE considers the times of multiple hits from each DOM. Instead, it considers the total number of hits and time of the first hit. This is an improvement over SPE, but it still does not include all the data available to us. This conscious dismissal of available data marks the SPE and MPE algorithms as intermediate-level reconstructions, not to be considered final results.

5.11 The Millipede Software Library

Millipede is a software package for doing likelihood based reconstruction that tries to incorporate as much information as possible about light travel through the South Pole Ice. It is designed to be used in conjunction with the Gulliver suite, to facilitate swapping out different parts.

The central portion is a method for calculating likelihoods of observed charge given hypothesis source cascades. The configuration of the cascades that emit the light is left in a different part, which makes it more configurable; specific configurations are described in the following subsections.

The likelihood calculation works as follows. The charge on each DOM is split up into time bins based on the amount of charge, so that each bin contains a certain (configurable) amount of charge, and more bins are created on each DOM until no bin has more than that amount of charge. The first of the two main functions is “Get Response Matrix.” When this function is called, it constructs a matrix with the hypothesis sources as columns and the observed charge (split up by DOMs and DOM-time-bins) as rows. The problem of observed charge amplitude is thus reduced to a linear algebra problem: we know the observed charge \vec{q} and the matrix of energy losses M , and we want to find the energy loss at the sources $\vec{\epsilon}$:

$$\vec{q} = M\vec{\epsilon}$$

Solving for the energy losses is thus reduced to a simple matrix inversion, which is performed in the second of the two main functions, “SolveEnergyLosses.” The reader should note that the matrix involved is quite large, especially because DOMs are included in the calculation whether or not they have observed charge. For a large event, the matrix would include 5,000 DOMs, perhaps half of which have 3-5 time bins, so we can approximate that as 10,000 rows in the matrix. A typical Millipede source configuration is one source every 10 meters along a track, so the matrix could easily have 100 columns. Externally optimized sparse matrix linear algebra libraries have been essential to making this fit work well.

A convenient coincidence of the Millipede physics configuration is that it is very similar to the configuration for reconstructing the results of Positron emission tomography (PET) medical scans, and thus has a vast literature addressing reconstruction techniques. PET scans detect pairs of gamma rays emitted by a positron-emitting radionuclide (tracer) within a biologically active molecule. The positions of the detecting machines are known, and precisely timed knowledge of the detected gamma rays allows the scanner to reconstruct the position of the emitting particle. This is algorithmically similar to Millipede, where we reconstruct the position of particles emitting light detected by DOMs of known position.

Specifically, PET scan literature has addressed the effect of the choice of minimizer used in reconstructions. By default, the minimizer Millipede uses is a preconditioned conjugate gradient (PCG) minimizer [101], which works faster and more reliably than the standard non-monotonic maximum likelihood minimizer. Specifically, PCG avoids spikiness in the likelihood search as a function of the number of iterations, so the circumstance of early minimizer truncation does not introduce unnecessary complications.

At the time of this thesis, Millipede is the best tested, most maintainable, and most accurate reconstruction suite available in IceCube.

5.11.1 Monopod

Monopod is a specific physics hypothesis used within the Millipede Software Library. The configuration is the simplest possible: a single point source of light, a good approximation of a cascade interaction. For tracks shorter than the vertex resolution, this can also be a fairly good fit. It is the only physics configuration within Millipede from which the energy can be read directly without further translation, although some scaling may be necessary depending on the kind of interaction.

This is the fit used to obtain the observed energy variable in this thesis. In this configuration, the CLast (see §5.7) algorithm was used as a seed for Monopod. Monopod is particularly good at reconstructing the energy of events which are fully contained within the detector, as is the case for nearly all oscillation-energy events.

5.11.2 MuMillipede Fit

MuMillipede is the muon reconstruction within the Millipede library, and the fit for which Millipede gets its name. The image to which the name refers is the arthropod millipede, which has a segmented body with legs extending from each segment; this looks similar to a chain of cascades within the detector, each with a Cherenkov cone extending from it.

The physics hypothesis in MuMillipede is a series of cascades spaced along a track. Specifically, given a track direction and cascade spacing parameter n , MuMillipede places a cascade every n meters along the track. For the energy range relevant to IceCube, many energy losses from muons are stochastic (e.g. bremsstrahlung), so a fit with independent energy losses along a track is often more accurate than a continuous loss “light saber” type model that IceCube has sometimes used in the past.

A modification of MuMillipede was used for the high energy starting events (HESE) analysis [36][102]. The same cascade configuration hypothesis was used. The change was that instead of letting a minimizer vary the direction to find the best fit, all directions were specified one at a time in a scan across the whole sky. The result of this was a likelihood map for all directions, rather than a single best-fit direction. One advantage of such a fit is full confidence that the reported maximum likelihood is indeed the true global maximum; another advantage is knowledge without further calculation of the angular confidence contours of the directional fit. The disadvantage is that this approach is CPU intensive: each event takes 3 to 10 CPU years to reconstruct.

5.11.3 Igelfit

Igelfit is an extension of MuMillipede. It runs MuMillipede several times with fixed directions, then compares the likelihood of each direction and returns the one with the best reduced likelihood. It automates and somewhat simplifies the search process undertaken by the HESE analysis. In particular, it has a configurable number of search directions:

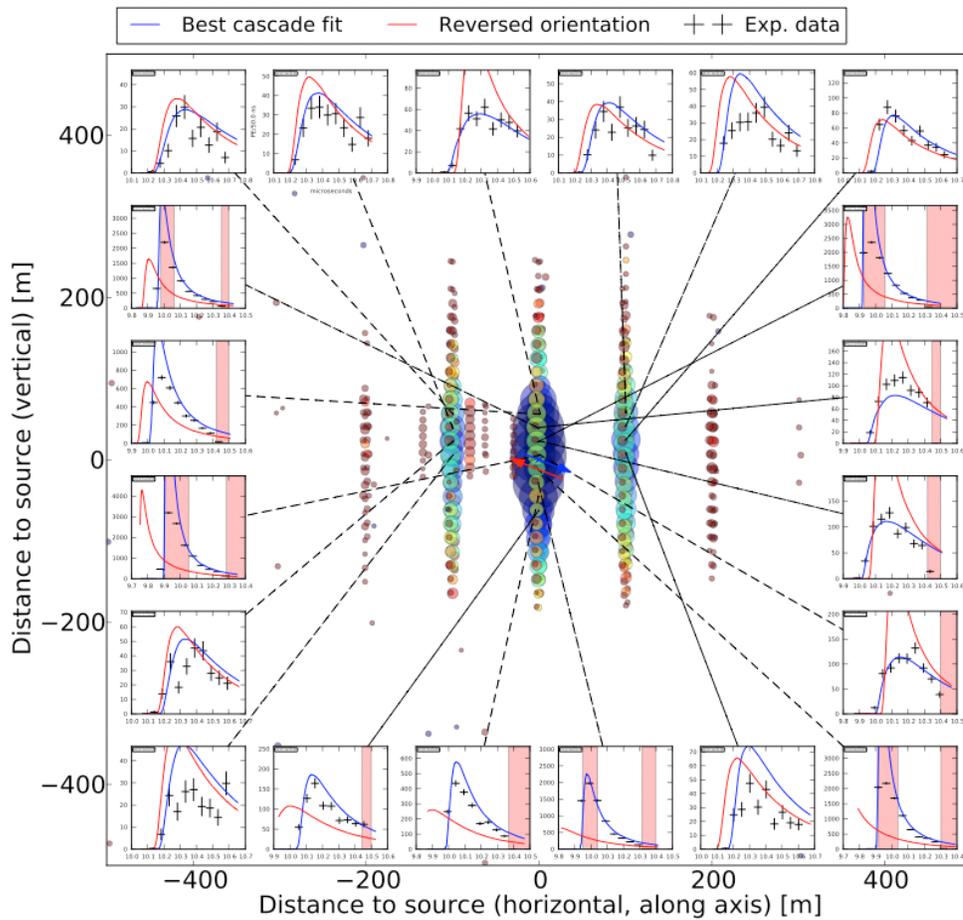


Figure 5.2: A multi-panel plot showing how we combine information from many PMT waveforms to reconstruct the direction of cascade-type events. The central panel shows a dot for each DOM that registered a signal for this event. The size of the dot is proportional to the amount of charge the DOM registered; the color shows the time that the signal started, with blue being early and red being late. For some DOMs, a line connects the DOM dot to a plot on the outer edges, showing the waveform from that DOM. Within those plots, the black crosses show the observed data. The red and blue curves show two different predictions for the light arriving at the DOM. Blue shows the best-fit direction; red shows the exact opposite direction. Notice that the blue fits the black better than the red does, when considering all DOMs together (even though some DOM signals are not perfectly predicted by either direction). This figure is meant to demonstrate that it is possible to find a direction of travel from a cascade-type event, even though it is not as easy as for a track-type event.

128, 512, or 1024. The interaction vertex position and time are given by the seed vertex position and time, and not allowed to vary.

In studying the various input parameters for Igelfit, each number of trial directions was tested, and no improvement was found by adding more test points. This indicated some deeper pathology, and was one of the factors that led to the project of hand-tuning the spline fits of photonics tables for low energies and small distances (the region which is of particular interest for low-energy fits).

Whereas Igelfit does not let the direction of the hypothesis track vary, a test was performed using Igelfit as a seed for MuMillipede, which does allow the track direction to vary. No improvement was observed; this was an unexpected result.

Igelfit is named for the german word “igel,” which means “hedgehog.” The many test directions radiating from a single point looks like a hedgehog (see Fig. 5.3) instead of a millipede.



Figure 5.3: A hedgehog, or “igel”, as rendered by the artist Abigail Rindo in a drawing depicting the croquet scene from Lewis Carol’s “Alice in Wonderland.” [103]

5.11.4 Other Millipede Configurations

The modular design of the Millipede software library makes it easy to create fits around new physics hypotheses. Some other fits that are in creation or testing include a fit for tau events and a fit for low energy cascades combined with tracks.

The tau fit is called “Taupede,” in keeping with the “-pede” theme. The physics hypothesis it uses is two cascades connected by a track whose length is related to the total observed energy. This fits the hypothesis of a tau neutrino interacting in the ice, creating an initial hadronic cascade and a tau charged lepton, and then the tau decaying in a second cascade a short distance later.

The low energy fit is called “Biped” because it includes two particles: a cascade and a track, with the same starting vertex and position but independent directions and energies. This fit is currently under development.

Chapter 6

The Shadow of the Moon in Cosmic Ray Muons

*Finally, the sun and the moon must cast a “shadow”
in the flux of high-energy primary cosmic rays
and observations of this shadow effect might give new
information about the magnetic fields of these bodies.*

— GEORGE W. CLARK, 1957 [104]

The Moon blocks a fraction of cosmic rays from reaching the Earth, and this deficit is useful for calibrating any detector sensitive to cosmic rays. This chapter describes some of the IceCube Moon shadow searches. Observing the Moon shadow validates IceCube pointing capabilities. It validates that there are no large shifts in the coordinate transform systems, and that the timing is accurate on a large scale.

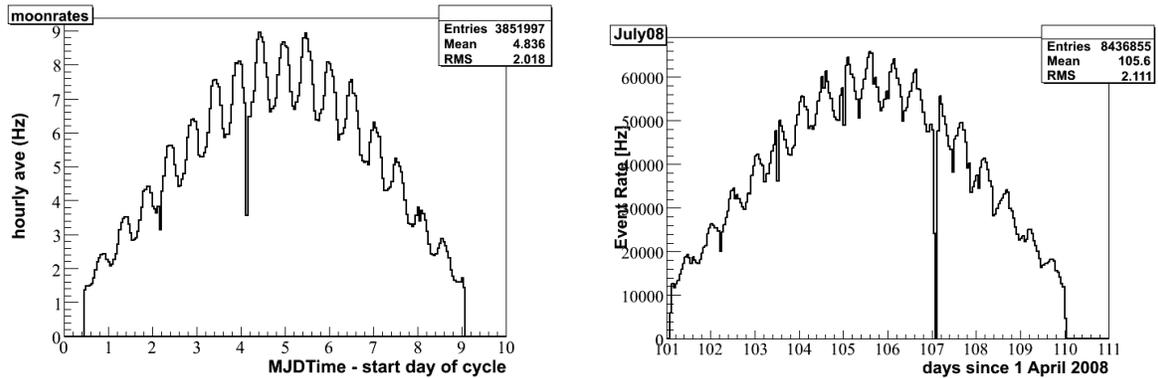
The Shadow of the Moon has been used as a calibration analysis in many gamma-ray telescopes, for example HAWC [105], ARGO-YBJ [106][107], MAGIC [108], Milagro [109], and TIBET-III [110]. Other Moon shadow-observing particle physics experiments include: MINOS [111], Antares [112], and MACRO [113]. It was first proposed in 1957 [114] as an offhanded comment.

This analysis has been tried several times in IceCube as the detector improved. The analyses presented here (based on the IC22 (§6.2) and IC40 (§6.3) detector years) show the transition from the years of unsuccessful searches to the current era of increasing precision and significance, plus the transition from binned to unbinned searches (§6.4). The final IceCube Moon shadow paper [115] is included in its entirety as Appendix D.

6.1 The Moon Filter

As described in §3.5, IceCube can only transfer a limited amount of data to the North, and thus uses a system of filters to define which events will most likely be interesting in physics analysis. There is a specific filter dedicated to the Moon shadow analysis. This filter reconstructs event directions, then compares the direction to the current position of the Moon in the sky. If the event is close to the Moon, or within a defined off-Moon control region, it passes the filter. The exact meaning of “close” changes from year to year, as the filters are reviewed and sometimes redefined.

For IC22, we used a rectangular window of 90° in azimuth and 10° in zenith, in both directions away from the Moon. For IC40, we used a rectangular window of $40^\circ / \cos(\delta)$ in azimuth and 10° in zenith, in both directions away from the Moon.



(a) The Moon filter rates from July 2007, a cycle from the IC22 detector year.

(b) The Moon filter rates from July 2008, a cycle from the IC40 detector year. The feature around 107 is missing data from a dropped detector run.

Figure 6.1: MoonShadow filter rates as a function of time, over two one-month cycles. The filter turns on when the Moon reaches 15° above the horizon at the Pole. Higher rates from the filter correlate with higher zenith angles of the Moon.

For both IC22 and IC40, the reconstruction algorithm used was a single (i.e. not iterated) SPE fit (see §5.10). Additionally, during both years small events were discarded: events with hits on fewer than 12 DOMs or 3 strings.

The window is so much larger in azimuth because, for several reasons, we wanted the off-source region to be at the same zenith angle as the Moon, so we needed more sample space in azimuth than zenith. There is a strong zenith angle dependence of the events in IceCube: around 10^6 more down-going events than upgoing events trigger the detector. The Moon filter rate is so dramatically dependent on the declination angle of the Moon position because the cosmic ray muon flux has a strong zenith angle dependence: a higher event rate is observed from directions close to the zenith than close to horizon. As the Moon rises and sets relative to the South Pole, this means that the filter stream varies drastically. The highest rates come from the times when the Moon is highest in the sky (viewed from the ice), and the filter shuts off by design if the Moon is less than 15° above the horizon. Below that, the data rate is so much lower that it was not worth the computing time, for a first analysis, to save the data.

An additional effect causing expected variations in the filter rate comes from the azimuthal angle between the detector and the Moon. Events traveling close to several strings of the detector are more likely to be detected and reconstructed, so when the aligned strings point to the Moon, the event rate increases. This can be seen in the peaked structure of Fig. 6.2, which carries over into peaks in the rates from typical months, shown in Fig. 6.1. The peaked structure is much stronger in Fig. 6.2 than Fig. 6.1. because the azimuthal structure averages out as the detector rotates relative to the Moon. The livetime of each Moon cycle is not long enough for the peaks to completely average out, so peaks are still present in Fig. 6.1.

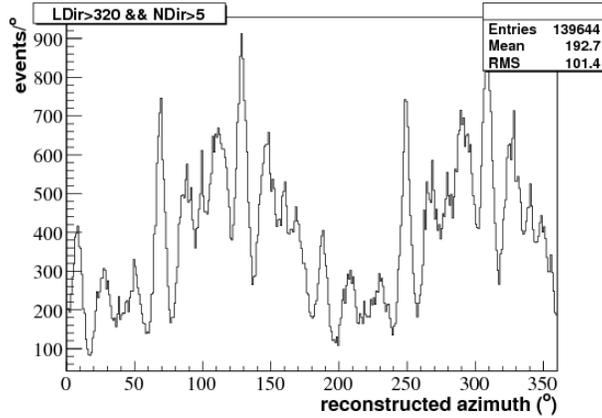


Figure 6.2: Moon filter rates versus reconstructed azimuth, showing a peaked structure. This peaked structure causes the 12-hour peaks in Fig. 6.1. These data were taken using the IC22 detector configuration.

Consider the various peaks of Fig. 6.2, which represents the data from a single 8-day Moon cycle (the 8 days within a single lunar month when the Moon was high enough above the horizon that the Moon filter took data). There are two sizes of structure: the narrow peaks $\sim 20^\circ$ wide and the wider peaks, of which there are two. The wide peaks correspond to the times when the direction of the Moon aligned with the long axis of the asymmetrical IC22 detector configuration; the corresponding two wide troughs likewise correspond to the times when the direction of the Moon aligned with the short axis of the detector. The narrower peaks arise from alignments of specific strings. Because of the extreme discreteness of the IC22 detector setup relative to the total detector size, azimuth plots from IC22 (e.g. Fig. 6.2) have larger variations in rate than similar plots from later and larger detector configurations.

6.2 The IC22 Binned Analysis

The main new idea of the IC22 Moon shadow analysis was an attempt to use off-source regions from different zenith angle bands relative to the Moon position.

Events were selected to improve the median angular resolution while still preserving adequate statistics. An approximation argument was used to find a significance indicator function, as follows. In this discussion, “signal” is understood to mean “deficit,” since the signal of interest is a lack of events from the direction of the Moon.

The final search bin contains events from both signal and background; the significance will scale roughly as signal over square-root of background:

$$\text{significance} \propto \frac{\text{signal}}{\sqrt{\text{background}}}$$

The search bin size should be proportional to the median angular resolution $\Delta\Psi$ of the sample:

$$r \propto \Delta\Psi$$

The number of background events will scale directly with the area of the search bin and the total efficiency of the cuts:

$$\text{background} \propto \text{efficiency} * (\Delta\Psi)^2$$

The number of signal events, assuming that the Moon is a point-like source relative to the angular resolution of this sample, scales only with the efficiency of the cuts:

$$\text{signal} \propto \text{efficiency}$$

Putting those together, we see:

$$\begin{aligned} \text{significance} &\propto \frac{\text{signal}}{\sqrt{\text{background}}} \\ &\propto \frac{\text{efficiency}}{\sqrt{\text{efficiency} * (\Delta\Psi)^2}} \\ &\propto \frac{\sqrt{\text{efficiency}}}{\Delta\Psi} \end{aligned}$$

Cuts were designed to maximize this function.

The following event selection cuts were used. Events with hits on fewer than 12 DOMs or 3 strings were discarded; these cuts were applied at filter level. Using the reconstruction algorithm SPE8 (see §5.10), events with a direct length of less than 340 m or fewer than 6 direct hits were discarded. These cuts produced a sample which, in cosmic ray muon simulation samples, had a median angular resolution of 1.80° with a passing rate of 50.6%. Using a similar argument, the bin size was optimized to 1.56° . A search bin of this size should contain $\sim 45\%$ of the signal (in this case, the deficit).

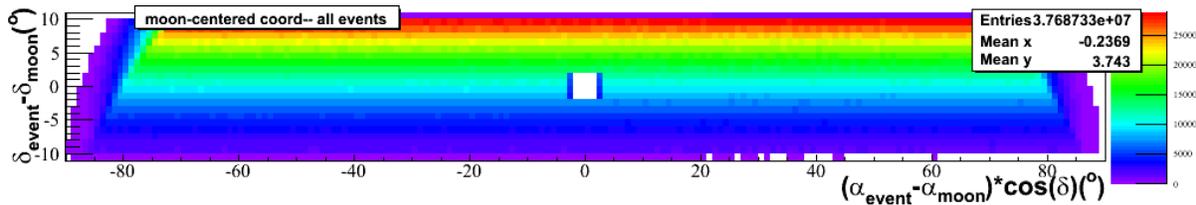


Figure 6.3: The total IC22 Moon shadow filter sample, after analysis cuts. The color indicates event rate. The y-axis shows the declination (similar to zenith, for IceCube) difference between each reconstructed event direction and the position of the Moon at the time the event was recorded. The x-axis shows the right ascension (similar to azimuth, for IceCube) difference, with the scaling factor to account for geometrical differences in height. The central bins have been excluded in this figure to preserve a blind analysis during the early stages of analysis development.

To search for the deficit at the Moon’s position, the reconstructed directions of all events were plotted relative to the position of the Moon at the time each event was recorded. A trigonometric scaling function of the Moon’s current declination $\cos(\delta)$ is added to account for the spherical geometry. The total rates are shown in Fig. 6.3, with the central Moon position blanked out. The zenith angle dependence of the flux is strongly visible.

This phase space relative to the Moon was split up into bins. Rates from bins in the same declination band were similar, with some remaining structure in azimuth. This structure can be seen in Fig. 6.4.

To account for this structure, an assumption was made that the right ascension structure was the same in adjacent declination bands. A normalization function was formed based on the rates of the two declination bands immediately adjacent to the central, Moon band. Applying this normalization, the corrected Moon band is shown in Fig. 6.5. The Moon is in a bin with a 1.7σ downward fluctuation from the average. This 1.7σ is considered the final significance of the IC22 Moon shadow search, and is not publicly documented elsewhere.

6.3 The IC40 Binned Analysis

The IC40 binned analysis used similar techniques to the IC22 analysis, with some notable differences highlighted here, including a more significant result, and abandoning the declination-band normalization procedure. A fuller documentation can be found in [116], which discusses an analysis of 8 of the 14 lunar months of IC40.

The reconstruction algorithm on which the analysis is based changed from SPE8 to SPE32 (see §5.10). This, in addition to improvements in the detector itself, improved the angular resolution of the data sample.

The event selection cuts were re-optimized, using the same significance estimator function. Events were discarded if they had an estimated reconstruction error more than 1.6° . No cuts were made on the direct length or number of direct hits. The resulting median angular resolution was 1.27° , and the search bin size used was 1.25° square, which was found from a radial search bin size of 0.8° .

The significance for each bin was calculated using [117], which takes into account the size of the background sample:

$$\text{significance} = \frac{N_{\text{on}} - \alpha N_{\text{off}}}{\sqrt{\alpha(N_{\text{on}} + \alpha N_{\text{off}})}}$$

where

$$\alpha = \frac{N_{\text{on}}}{N_{\text{off}}}$$

This significance was calculated for each bin of Fig. 6.7 in the following way. The data are plotted relative to the Moon, correcting for projection effects. The bin size is the optimized $1.25^\circ \times 1.25^\circ$ ¹. Each bin successively is considered an on-source region, and 20 off-source regions are taken from the same declination band, starting 3 bins away (i.e., skipping two bins). Ten of the off-source regions are right ascension angles larger than the Moon position; ten

¹Recall that the filter window extends for 40° on both sides of the Moon position, and that the search bins are 1.25° wide. That leaves 32 bins on each side of of the Moon available for analysis.

are at angles less than the Moon position. To test the hypothesis that the background fluctuates randomly, these significances are plotted in Fig. 6.8, and are found to be Gaussian distributed around 0 with a width of 1, which confirms the hypothesis and validates the analysis method.

The final significance found for the Moon shadow in IC40 with a binned method was 7.56σ .

A second, similar search was performed with a different arrangement of on-source and off-source regions. This is shown in Fig. 6.9. The significance of the result from this search was similar: 7.9σ .

6.4 The IC40 and IC59 Unbinned Analyses

An independent analysis was performed on the IC40 data sample and continued with IC59 sample using an unbinned likelihood technique similar to IceCube point source searches [118](but without an energy-dependent term). A summary of the IC40 analysis can be found in [119], and more extensive treatments of both can be found in [120] and [121]. This section briefly describes the differences between binned and unbinned searches as they relate to the Moon shadow search.

A binned search is inherently limited by a balanced choice of bin size: if the bins are large, statistical errors in each bin are small but angular sensitivity is lacking. If the bins are small, statistical errors dominate. An unbinned technique lessens the impact of this choice.

A binned search, also sometimes called a “cut and count” analysis, defines hard edges of an on-source bin. A single event is either counted as in or out, regardless of how close to the edge of the bin it is. By contrast, an unbinned search considers the probability that any event was pointing in a given direction, based on the reconstruction error estimate of each event. This way, no information is thrown out; the reconstructed position and error from each event is included in the final analysis, without averaging together the errors of the whole dataset.

This is implemented through a likelihood function:

$$L(\vec{x}_s, n_s) = \sum_i^N \log \left(\frac{n_s}{N} S_i + \left(1 - \frac{n_s}{N}\right) B_i \right)$$

where \vec{x}_s is a directional position (in this case, relative to the Moon’s position), n_s is the number of signal events, N is the total number of events, and S_i and B_i are the expected shapes of the signal and background, respectively.

The shapes of S_i and B_i were obtained by constructing a position distribution of all events: each event contributed a Gaussian centered at its reconstructed position, with a width of its reconstruction error estimate. This provides a smooth function that can be evaluated on an arbitrarily fine grid; in particular, the distribution can be evaluated on a grid much finer than the angular reconstruction error of the sample, so the measurement is not limited by the granularity of the grid.

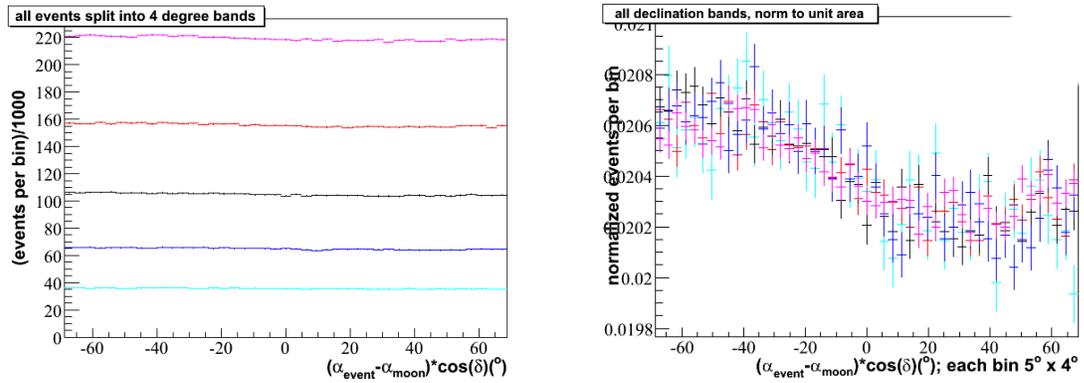
The signal region is centered, as always, on the Moon. The background shape B_i was obtained similarly to the binned searches: by looking at regions of the sky adjacent to but not containing the Moon.

At each point on the grid, using these S_i and B_i , the likelihood function is optimized by varying n_s . Because the Moon produces a deficit, negative values of n_s were expected in the region of the Moon.

For IC40, the best fit value of n_s for the deepest deficit was 5326. To calculate the significance of this, the fluctuations of the two background regions were measured. These were fit with to a Gaussian distribution, which was found to have widths of 544 and 498 events, both centered at 0. Taking this width as a standard deviation, the deficit has a significance of $10 - 11\sigma$. The similar analysis of the IC59 dataset found a deficit of 8660 events, for a significance of $13 - 15\sigma$. Both of these results are described in detail in Appendix D.

6.5 Future Studies

In the future, it could be possible to observe the Moon shadow every month, and use it as a form of monitoring. It could be possible to automate certain elements of this analysis to the extent that the angular size of the shadow, the depth relative to the surrounding rate, and the shift from magnetic fields could all be used to monitor changes in the detector. Currently, we still take data with a Moon filter, and scientists from the IceCube cosmic-ray group continue to study the output.



(a) Rate versus right ascension relative to the Moon, for 5 declination bands relative to the Moon.

(b) The same bands as (a), normalized to equal area. Even before correcting for the global structure, the central Moon bin can be seen as a slight outlier.

Figure 6.4: The IC22 declination bands relative to the Moon. The band centered on the Moon is shown in black. Bands above the Moon are shown in red and magenta; bands below the Moon are shown in blue and cyan. A similar structure can be seen in all bands.

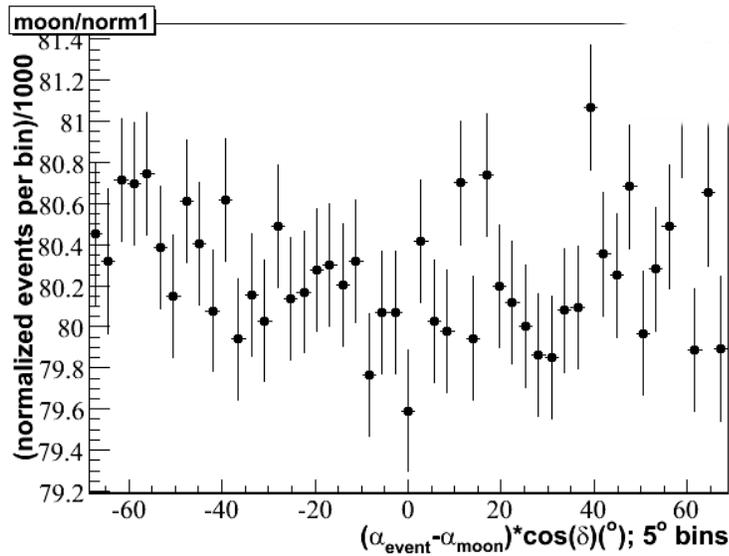


Figure 6.5: The IC22 Moon shadow final histogram. The central bin, containing the Moon, shows a 1.7σ deficit.

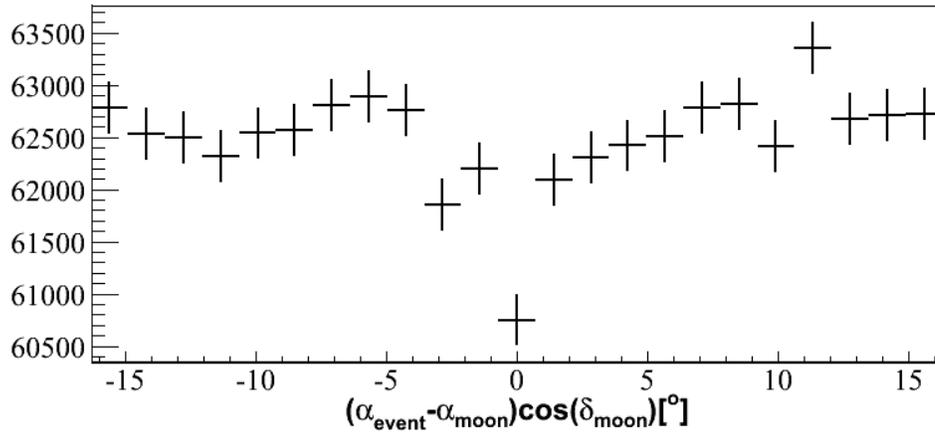


Figure 6.6: A histogram showing the rate of events in a declination band around the Moon, using IC40 data. A clear deficit can be seen in the direction of the Moon, at 0 on the x-axis. The errors shown here are \sqrt{N} . This figure is the IC40 version of Fig. 6.5.

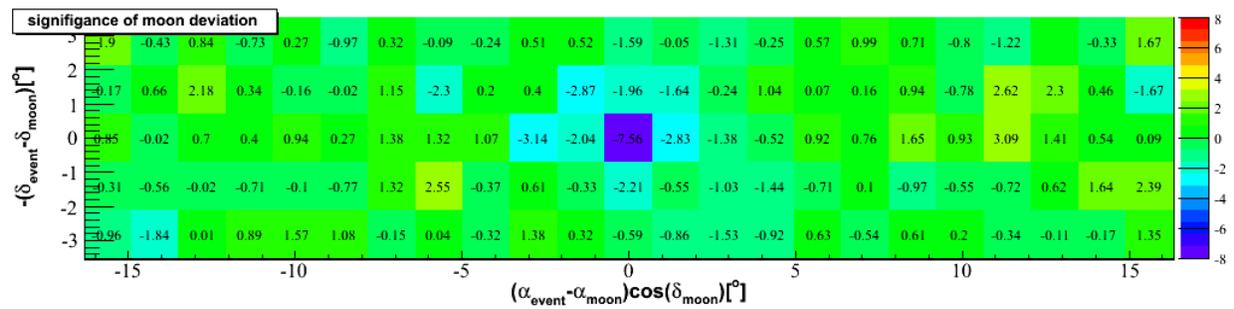


Figure 6.7: A histogram showing the significance of deviations in rates in the region around the Moon, using IC40 data. See text for significance calculation.

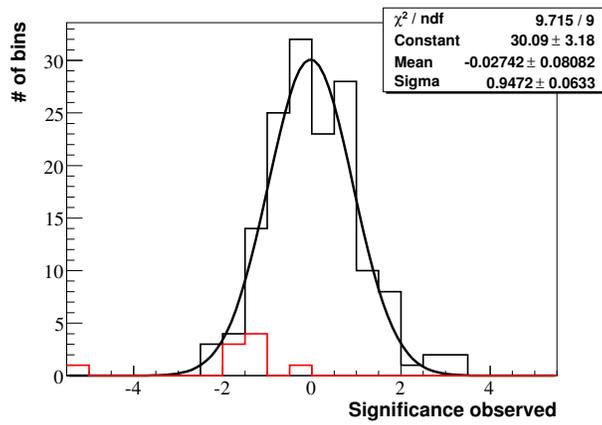


Figure 6.8: Gaussian fit to the significances of deviations in the region around the IC40 Moon. The 9 bins at and adjacent to the Moon are shown in red; all show deficits. The remaining bins make up the larger, black histogram. A Gaussian distribution is fit to those, and its width is found to be consistent with 1.00. This indicates that the non-Moon fluctuations are properly random. Caveat: this figure was made based on 8, not the full 14, lunar months of IC40.

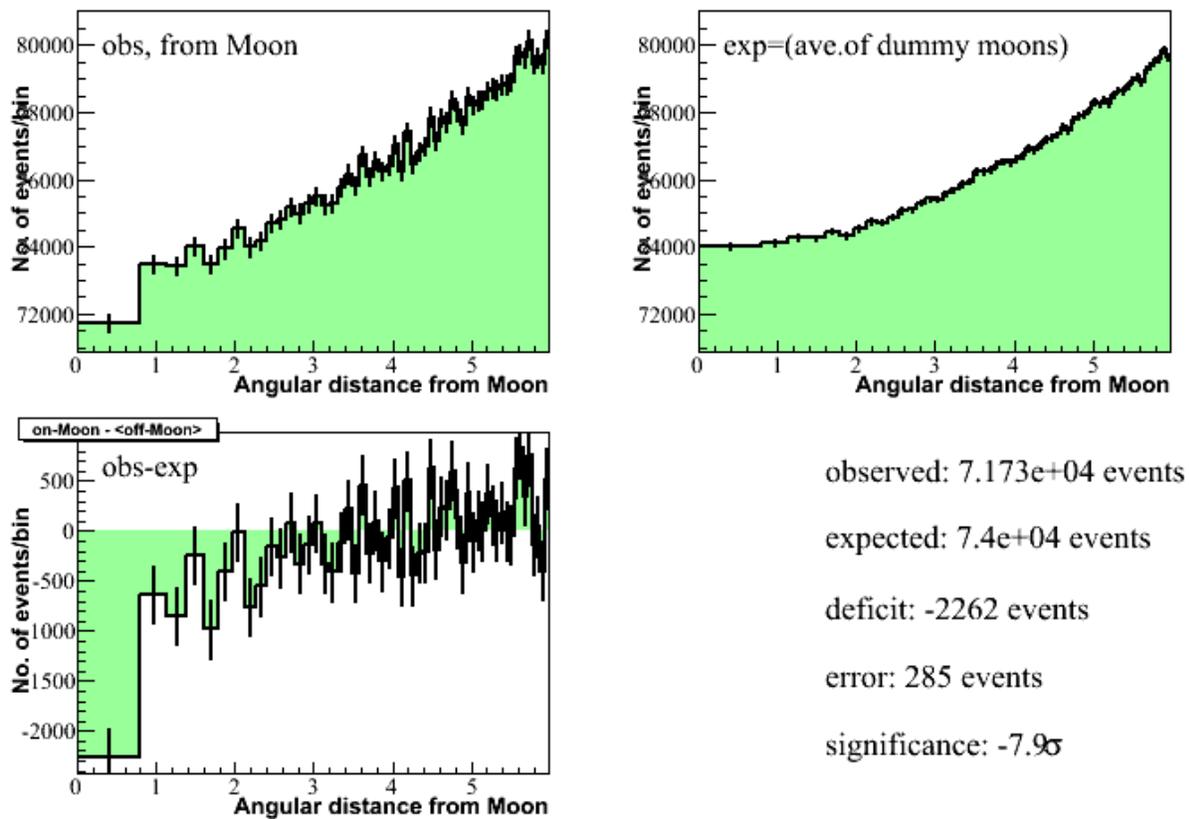


Figure 6.9: An alternate Moon shadow search using annular bins around the Moon position. Instead of binning events rectangularly, this search uses the angular distance between the Moon and the reconstructed event direction. The bin size decreases for larger distances so that each bin represents the same solid area. There is an overall increase at larger radii because of the strong zenith angle dependence of the cosmic ray flux; if this flux were linear or quadratic with zenith, this method would have smoothed it away. The top-left plot shows a basic histogram of number of events relative to the Moon. The top-right plot shows a compiled histogram, averaging over 8 constructed off-source histograms. Each off-source histogram, or “dummy Moon”, was constructed by considering a point away from the actual Moon but in the same declination band. The various off-source histograms are spaced far enough apart that they do not overlap. The bottom-left plot shows the difference between the two top plots. The deficit from the Moon shows as the first bin of the bottom-left plot. Various statistics about the plots are shown in the bottom-right. The final significance, using Li and Ma [117] statistics on the central bins, is 7.9σ .

Chapter 7

Event Selection

Hofstadter's Law: It always takes longer than you expect, even when you take into account Hofstadter's Law.

— DOUGLAS HOFSTADTER [92]

This chapter describes the event selection algorithm used to isolate a sample of upgoing neutrinos for oscillation analysis. The specific cuts are described here; for a more general introduction to the algorithms, see Ch.5. The cuts fall broadly into three categories: cuts performed online for all DeepCore analyses (§7.3), cuts designed by previous oscillation analyses (§7.4), and cuts designed specifically for this analysis (§7.5). Additional introduction to the simulation specific to this analysis is in §7.2, and comparisons between cut levels in both data and simulation are in §7.6.

This analysis focuses on atmospheric neutrino events that have traveled up through the Earth. Atmospheric neutrinos are created isotropically around the Earth (specifically, at all zenith angles relative to IceCube), and we focus on the upgoing neutrinos for two reasons.

First, we know most accurately that they are neutrinos and not other cosmic-ray induced air shower products, since neutrinos are the only particles that travel through the Earth without decaying or interacting.

Second, the most upgoing neutrinos give us the best view of oscillations because of the length over which they travel. The IceCube detector is most suited to $o(1)$ TeV or higher energy interactions, but with DeepCore some events are reconstructable as low as 10 GeV, with limited resolution. Given the limitations of IceCube, oscillation effects are most easily observed at the highest energies they are known to exist.

To determine the most relevant energies, consider the two-neutrino approximation of the oscillation formula:

$$P(\nu_\mu \rightarrow \nu_\tau | L(\text{km}), E(\text{GeV})) = \sin^2 2\theta_{23} \sin^2 \left(1.27 \frac{\Delta m_{23}^2 L}{E} \right) \quad (7.1)$$

Note especially that the oscillation period varies with L/E . To see a given peak with the highest possible energy events, we need to give them the longest possible length over which to oscillate. For atmospheric neutrinos, that means looking directly to the other side of the Earth, which means the events arrive in our detector going directly upwards. Assuming an oscillation length of the diameter of the Earth (~ 12000 km), the oscillation maximum at the highest energy occurs between 20 and 30 GeV; this approaches the energy range visible to Super-Kamiokande.

A further aim of the event selection presented here is to isolate muon neutrino charged current interactions from electron and tau neutrino charged current interactions and any neutral current interactions. It would be particularly helpful to distinguish charged current ν_μ from ν_τ , because the oscillation effect in question produces a disappearance in the ν_μ rate and an appearance in the ν_τ rate at the same energy and oscillation length. However, no part of this analysis was specifically designed to distinguish charged current muon neutrino events, because they already dominate the sample of well-reconstructed contained upgoing events, and no method known within IceCube can separate the various flavors of events in the 10 GeV region accurately enough to further improve the purity of the data sample while preserving an adequately large sample.

The goal of this event selection is to identify a sample, as large and pure as possible, of upward traveling events in this 20-30 GeV energy range.

7.1 Limiting Bias Using Blind Analysis

Blind analysis is a technique for limiting unconscious bias, commonly used within particle physics [122]. In IceCube, the specific method of blind analysis we use restricts the experimental data sample to only 10% of the total data while developing an analysis. The complete dataset is then used after an “unblinding review”, during which the collaboration approves a defined post-unblinding procedure. The unblinding review incorporates an analysis approval process from the IceCube collaboration.

Within this analysis, the event selection was designed using a reduced sample of data. The sample was defined by selecting runs with run identification numbers ending in “0”; this approximates 10% of the total sample. The livetime of this reduced sample, called the “burnsample,” was 781.19 hours. The livetime of the full sample was 7500.71 hours; the burnsample represented 10.4% of the full sample. Both parts of the event selection (both the selection presented in §7.4 and §7.5) were designed this way. However, only the analysis in §7.4 and [123] defined an unblinding procedure; after unblinding the data once, a second unblinding was considered redundant. The fact that the limiting factor of this analysis was not the statistical size of the sample (see §10) further reduced the importance of blind analysis in this context.

7.2 Monte Carlo Simulation

The IceCube simulation uses several different Monte Carlo generators to produce different kinds of simulated data. Cosmic ray air showers, a main background to this analysis, are generated with the CORSIKA software package (Cosmic Ray Simulation KASCADE) [124]. High energy neutrinos (and originally, all signal neutrinos) are generated with a software package called NuGen, based on ANIS (All Neutrino Interaction Generator) [125]¹. The third generator is Genie (Generates Events for Neutrino Interaction Experiments) [88], which is used for sub-TeV energy neutrinos.

¹This citation [126] brings this thesis to an above-average number of citations for a PhD thesis in engineering and technology.

NuGen and Genie each have advantages and disadvantages for this analysis. While NuGen generates neutrinos in the wide energy range from 10 GeV to 1 ZeV, it is most useful in a narrower range. NuGen does not include interaction processes other than deep inelastic scattering, and so it is best suited to energies above several TeV. NuGen does have the advantage of including information about propagation through the Earth based on the Preliminary Reference Earth Model [127], which is particularly important for oscillation analyses. Genie includes more detailed interaction information and cross sections at energies relevant to oscillations, and is reviewed, maintained, and verified by the wider oscillation community, especially accelerator neutrino experiments. The disadvantage of Genie is that it is currently tested only at the lower end of the IceCube energy range. The process of integrating these two generators is underway.

This analysis uses an energy-dependent mix of NuGen and Genie neutrinos as a signal sample. The method for mixing these two sets was developed for a previous IceCube oscillation analysis [123]. At low energies, only Genie is used. At high energies, only NuGen is used. In a transition region, both are used, with a linear transition from one to the other. The transition energy range is from 50 GeV to 190 GeV. The weighting modification is displayed in Fig. 7.1. An example of a distribution split up into all types of simulation is shown in Fig. 7.2.

7.3 The Online Filter

IceCube uses a system of triggers and filters (see §3.5) to perform the earliest stages of all event selections. This analysis starts with a filter stream designed to select events starting in DeepCore. Events passing at least one filter are saved in files which, for historical reasons, are called “Level 2” files; this is the start of the numbering scheme for labeling higher cut levels.

The trigger [58] works in the following way. IceCube uses a system of “hard local coincidences” (HLC) over a pre-defined trigger time window to decide when to read out the DOMs. The HLC condition is satisfied for a single DOM if it crosses a voltage threshold and then, within $\pm 1\mu\text{sec}$, another DOM out of its nearest or next-to-nearest neighbors on the same string also crosses the voltage threshold (then both pass HLC). The DeepCore filter uses data from two simple majority triggers: the first requires 8 HLC DOMs within $5\mu\text{sec}$ anywhere in the detector, and is called SMT8. The second, called SMT3, requires 3 HLC DOMs within $2.5\mu\text{sec}$, and those DOMs must be from within the DeepCore fiducial region. That region is defined as the high quantum efficiency strings and their neighbors, using only the DOMs below 2100 m depth. The SMT3 trigger adds less than 10 Hz to the total trigger rate.

In addition to the HLC hits used for triggering, information is recorded about all hits within $\pm 10\mu\text{sec}$ of the trigger time, but if an event did not pass HLC we don’t save as much detailed information about its PMT response; we save only the time and amplitude of the first pulse, on the assumption that there was only one pulse in the PMT waveform. These hits are called “soft local coincidence” because they coincide with triggering DOMs that may be relatively far away within the detector, but are still “local” in time. Soft local coincidence hits can be useful for further vetoing and event reconstruction.

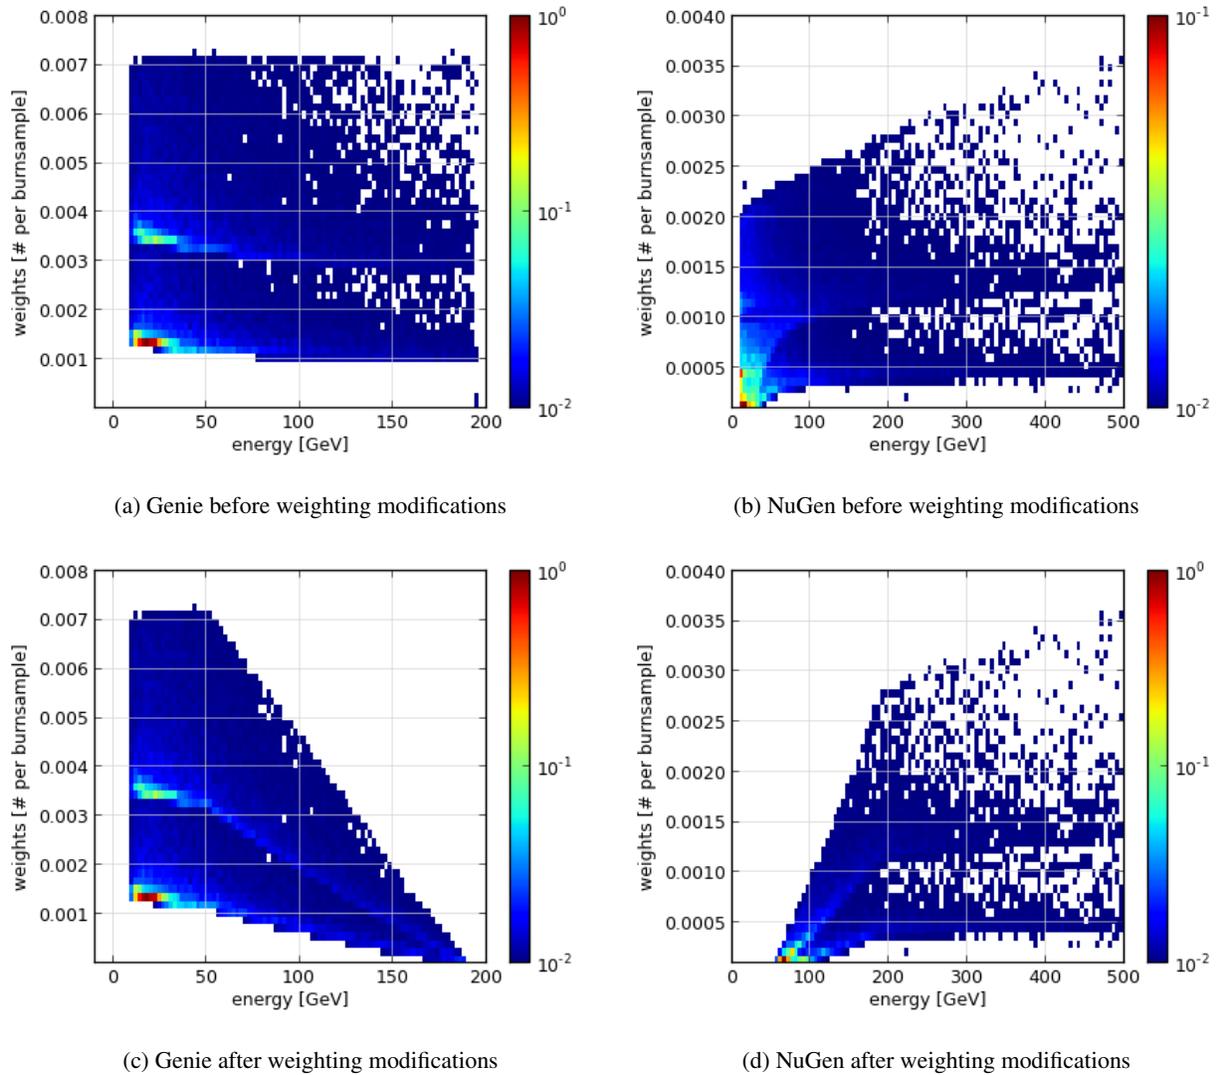


Figure 7.1: The Genie-NuGen cross-weighting method. The weights shown include oscillation effects, assuming maximal oscillations. The y-axis (weight of each event in the simulation) is different between Genie and NuGen because of different numbers of files generated: when more files are generated to represent the same livetime, each event has a smaller weight. The z-axis, shown as color, shows the number of generated events with a specific weight and true energy. This figure is meant to illustrate the effect of the cross-weighting method, which is seen in the difference between the upper panels and the lower panels. In the upper panels, the cutoff of the colored region reflects the full energy spectrum of simulated events. In the lower panel, the weights have been reduced in select energy regions: high energies for Genie, low energies for NuGen. In all plots, the clustering of events creating the streak-like structure is caused by pileup of similar interaction types: charged and neutral current interactions.

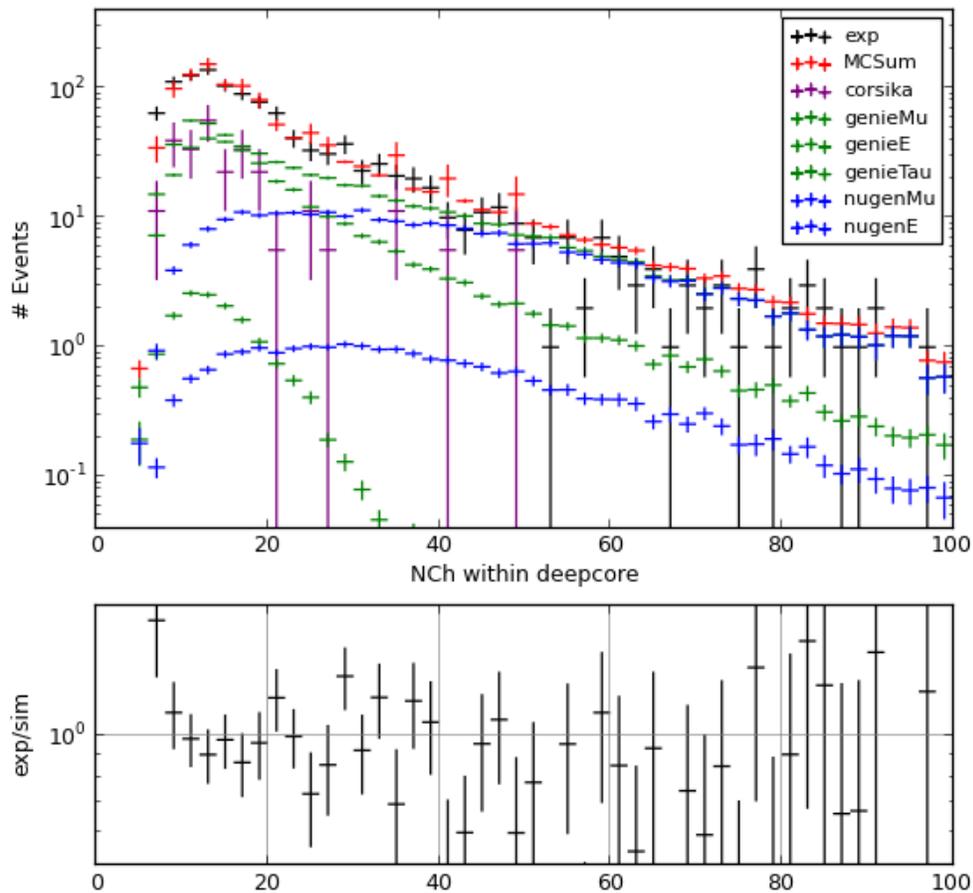


Figure 7.2: NCh versus rate at cut level 10, showing all simulation components. The neutrino simulations are split up by flavor. Notice that NuGen dominates at higher energies, and Genie at lower energies. The flavors are listed in the plot legend in the same order in which they appear in the plot: mu on top, then e, then tau. NuGen neglects tau flavor neutrinos. The sum of all simulation components is shown in red, with data in black for comparison.

The lower plot shows the ratio of data to total simulation; in terms of the upper plot, this would be the ratio of black to red.

The main purpose of the DeepCore filter is to distinguish interactions that started within the detector from particles arriving from outside, which were most likely created in cosmic ray showers unrelated to the neutrino physics in question. The assumption here is that if an event started outside the detector, it probably deposited some light on the way in. At DeepCore energies, an algorithm based on this assumption reduces the cosmic ray background by a factor of 3 to 4 orders of magnitude, but since there are 10^6 times more background than signal to start out with, further event selection is still needed.²

The DeepCore filter is based on the following specific algorithm. It starts with events that pass the aforementioned SMT3 trigger. Hits within an event are split into those in the fiducial region (DeepCore and the center of the IceCube strings) and in the veto region. Considering hits in the fiducial region, an approximate time and position of the interaction is estimated by taking the amplitude-weighted average of all the hits. Because of the weighting involved, this is often referred to internally as the center of gravity or CoG, even though the weighting is by charge amplitude and not by any mass. This CoG time and position can then be compared to the time and position of each pulse within the veto region (considering one pulse at a time) to check for causality. The speed between these two space-time points is calculated: the spatial distance divided by the time difference. If the speed looks like it is close to the speed of light, the interaction was probably from a cosmic ray muon starting outside the detector. The speeds, calculated in this way, are plotted in Fig. 7.3 for a sample of atmospheric neutrinos and a sample of cosmic ray muon background.

Specifically, after calculating the speeds between the CoG of the fiducial hits and each of the veto hits, the event passes the DeepCore filter if 0 or 1³ of the veto hits falls within the speed window 0.25 to 0.4 m/nsec. This cuts down the rates about 15 times relative to the SMT3 trigger, to ~ 16 Hz.

7.4 Preliminary Cut Levels

This section describes the cuts designed for a previous oscillation analysis [123], which (for the levels described here) were also used for this analysis.

7.4.1 Cut Level 3: First Fits

Level 3 is the first cut level specifically designed for oscillation analyses. It includes cuts on 4 quantities, each of which can be calculated quickly.

²Selecting contained events in an effort to reduce background is a technique common to many IceCube analyses. In this analysis, containment eliminates most throughgoing cosmic ray muons. At higher energies, where muons deposit more light in the detector, a larger fraction of the throughgoing cosmic ray muons can be eliminated. At the highest energies (100 TeV to PeV and above), effectively all muons from cosmic ray interactions are eliminated, as well as neutrinos from cosmic ray interactions in the atmosphere [128], leaving astrophysical neutrinos as the only remaining source. This was the basis for the High Energy Starting Event [36] or “muppet”[102] analysis designed in 2012.

³In some previous IceCube works, this was set to 0, 1, or 2; for this analysis, the cut value was consciously lowered to help reduce cosmic ray muon background.

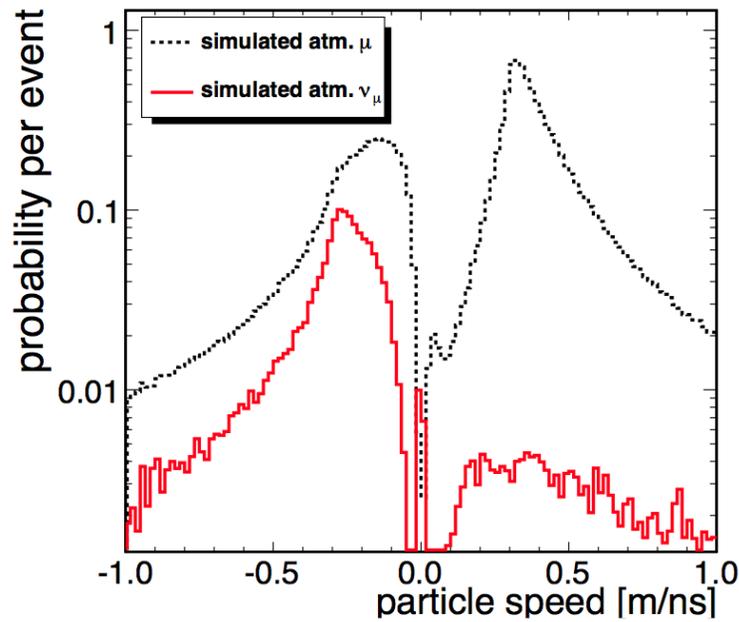


Figure 7.3: Particle speed probabilities per event for simulated muons from cosmic-ray interactions (black dashed line) and simulated muons from atmospheric neutrinos inside DeepCore (red solid line). The speed is defined to be positive if the hit occurred before the COG time (see text) and negative if it appeared after. Hits in the veto region are generally expected to have a speed close to $c \simeq +0.3\text{m/ns}$. Smaller speeds occur for light delayed by scattering. Larger speeds are in principle acausal, but since the COG time represents the start of a DeepCore event, whereas the COG position defines its center, the particle speeds for early hits are slightly overestimated. Events with a hit within a particle speed window between $+0.25$ and $+0.4$ m/ns are rejected. This is Figure 11 from [58].

A cut is applied on the number of DOMs that register a signal, after applying hit cleanings (see §5.3). The time window for cleaning is $[-4000,+6000]$ nsec relative to the trigger time. The time and radius values for Seeded RT cleaning are $R=150$ m, $T=1000$ nsec. After these hit cleanings, a minimum of 6 hits is required.

A cut is applied on a Linefit reconstruction (see §5.6) performed on only hits within the DeepCore fiducial region. The zenith angle of the resulting fit must be roughly upgoing or close to the horizon. The cut value is at 0.2 in $\cos(\text{zenith angle})$, which cuts out the most downgoing events.

Similarly, a cut is applied on an SPE reconstruction (see §5.10) performed on only hits within the DeepCore fiducial region. The cut value is also 0.2 in $\cos(\text{zenith angle})$, cutting out the most downgoing events.

Finally a cut is applied on the output of NoiseEngine (see §5.5), which eliminates noise-only triggered events. The output of NoiseEngine is a binary pass-fail, with parameters having been optimized during the algorithm's development.

7.4.2 Cut Level 4: NCh NVeto

The Level 4 cuts are based on two ways of counting the number of DOMs that register signals. The first cut requires 6 or more DOMs with hits in the DeepCore fiducial region, without any other hit cleaning. The second cut requires no more than 1 DOM hit in the veto region, that is, any of IceCube outside of the DeepCore fiducial region. The veto count only includes DOMs hit before the initial trigger (3 DOMs hit in rough coincidence with each other). The veto count also has hit cleaning applied to it: both Time Window cleaning and SeededRT cleaning (see §5.3).

7.4.3 Cut Level 5: TTrigger Splitter

The Level 5 cuts are designed to eliminate a particular class of mis-reconstructed events: pairs of downgoing cosmic rays that can confuse the timing of basic reconstructions, which return an upgoing fit based on a single particle hypothesis. This is accomplished using the TTrigger splitting algorithm (see §5.4).

Most of the specific settings for running TTrigger were optimized for astrophysical point source searches at higher (e.g., 10-100 TeV) energies. The exceptions were optimized for a DeepCore oscillation search [123]; those exceptions were the XY distance, the Z distance, and the Time Cone.

Using the output of TTrigger, the analysis further required at least one set of topologically connected hits. If a second set was found, i.e., if the event could be split successfully into two sub-events, the second sub-event was required to have zero hits. This is equivalent to requiring one and only one sub-event.

The resulting sample has ~ 544 k events predicted from simulation scaled to the total livetime of IC79, with a neutrino to cosmic ray muon ratio of 41 to 503.

7.4.4 Cut Level 6: Improved Linefit and SPE32

The Level 6 cuts use more sophisticated directional fits to further reduce the background of downgoing cosmic ray muons by cutting out all events reconstructed as downgoing. The fits used are Improved Linefit (§5.6 and [89]) and SPE32 (§5.10). The cut values are at the horizon ($\cos(\text{zenith angle}) = 0$). Both of these fits are run only on the hits after event splitting, that is, only on one of the sub-events.

The resulting sample has $\sim 133\text{k}$ events predicted from simulation scaled to the total livetime of IC79, with a neutrino to cosmic ray muon ratio of 27 to 106.

7.4.5 Cut Level 7: Causality-based veto

The Level 7 cut looks for causal relationships between hits in the veto region and the triggering fiducial hits. This cut was designed specifically for DeepCore oscillation analyses [123], and may be of interest to the wider DeepCore group.

The algorithm compares the time and position of the pulses which first trigger the detector: at least 3 DOMs hit within the DeepCore region and within 2 DOMs of each other. A hypothesis is tested that a throughgoing muon created hits in the outer region of the detector as well at the inner region. This is tested by looking at the speed between the outer hits and the inner hits. A speed consistent with the speed of light indicates the possibility of a sub-threshold muon. This algorithm is discussed in more detail in [123].

The resulting sample has $\sim 15.3\text{k}$ events predicted from simulation scaled to the total livetime of IC79, with a neutrino to cosmic ray muon ratio of 8.6 to 6.7.

7.5 Newly Optimized Cuts

The cuts at Level 8 and higher were optimized specifically for this thesis (in contrast to the previous section, which was optimized for [123]).

7.5.1 Cut Levels 8: Monopod

The Level 8 processing centers around the particle reconstruction algorithm Monopod, described in more detail in §5.11.1. The Level 8 cuts are all cuts on Monopod output and quality. As this reconstruction includes many technical improvements on previous reconstructions, it was expected to have greatly increased energy resolution, and the cuts were designed to optimize this resolution.

One cut required a minimum reconstructed energy of 10 GeV. At energies lower than this, there was no simulation available to evaluate the quality of the sample. Additionally, a non-linearity was observed in the rate at these energies, and without simulation this was difficult to investigate, so the region was excluded.

Another cut restricted the sample based on the reconstructed interaction vertex: it was restricted to the DeepCore fiducial region below the dust layer, specifically to -500 to -200 m in IceCube coordinates. The vertex also had to be within 150 m of the central string, String 36. Well outside of this region, the event rate was already low because of cuts at previous levels designed to veto events starting in the outer regions of the detector. Close to these cuts, a distortion in the reconstructed energy was observed (see Fig. 7.4). This distortion can be understood by a decreased light detection in those regions: if a particle is mistakenly placed in those regions by a reconstruction, the reconstruction will assume that the same amount of observed light was created by a brighter particle, thus misreconstructing the energy as too high. To maintain reconstructed energy accuracy, these regions were cut.

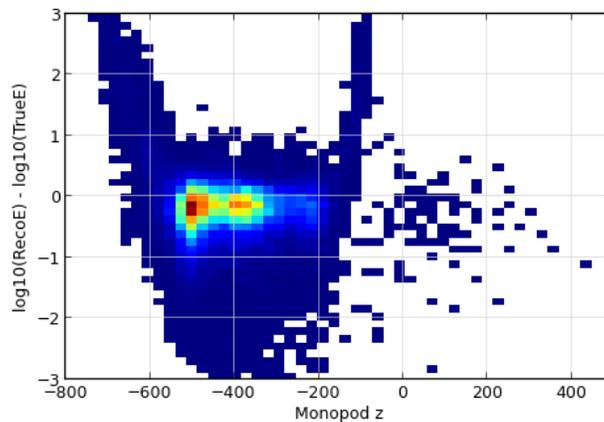


Figure 7.4: Monopod reconstructed depth versus energy resolution. The lowest DOMs in the detector are at -500 m; the highest are at +500 m. The dust layer (see Fig. 3.5) is at -100 to 0 m in these coordinates. The cuts for this analysis, as described in the text, cut out the regions in the dust layer and below the detector where the reconstructed energy is distorted.

The last Level 8 cut required a minimum of 5 direct hits. A direct hit is defined in terms of the “time residual:” the difference between the actual light arrival time and the expected light arrival time according to a particular emission position hypothesis. A hit is considered “direct” if it has a residual between -15 nsec and +75 nsec. A larger number of direct hits within an event indicates a high quality reconstruction, so requiring many direct hits cuts out events with bad energy and zenith angle resolution. On the other hand, it also simply indicates an event with a lot of hits, so cutting out events with few direct hits cuts the oscillation signal, which is at low energies. A balance was achieved by cutting the region dominated by cosmic ray muons.

The resulting sample has ~ 6.2 k events predicted from simulation scaled to the total livetime of IC79, with a neutrino to cosmic ray muon ratio of 4.9 to 1.3.

7.5.2 Cut Level 9: Igelfit

The final level of processing was Level 9, which centered on the Igelfit reconstruction (see §5.11.3). Igelfit was chosen for its accurate zenith angle resolution. No additional cuts were performed at Level 9.

7.5.3 Cut Level 10:

The final selection level includes cuts on reconstructed energy and zenith; this cut level places the events in a 2-dimensional histogram for analysis. The overflow and underflow bins are not included in the analysis, thus defining the edges of the histogram amounts to defining a cut. The placement of these cuts was studied with round trip trials (Ch.9), considering the full analysis chain and optimized for significance and limiting the influence of systematics-dominated regions at the lowest energies and most downgoing angles. Further discussion of this optimization is deferred to Ch. 9. The allowed events were between 10 GeV and 1 TeV, and had upgoing reconstructed zenith angles cutting out the horizon: $\cos(\text{zenith}) < -0.2$.

The resulting sample has $\sim 3.9\text{k}$ events predicted from simulation scaled to the total livetime of IC79, with a neutrino to cosmic ray muon ratio of 3638 to 296.

As Level 10 is the final cut level, this is the point to include information about the final resolution of the energy and zenith angle reconstructions. These are included as Fig. 7.6 and Fig. 7.5. The energy resolution was measured at 35%, which is 0.15 when measured in the logarithm. The median zenith angle resolution was measured at 11.8° .

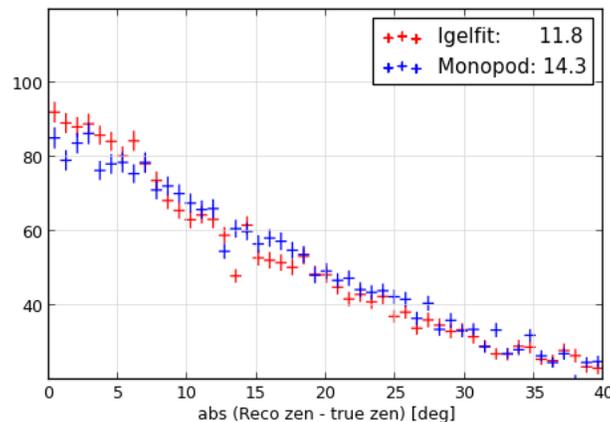


Figure 7.5: Event rate versus zenith angle reconstruction error for Igelfit and Monopod. The median of each sample is included in the legend: 11.8° for Igelfit, and 14.3° for Monopod. As expected, Igelfit performs better than Monopod.

7.6 Event Selection Summary

The remainder of this chapter contains summary views of several variables plotted at all available cut levels.

The true and reconstructed zenith angle and neutrino energy are particularly relevant because they are the closest to the physics being probed; recall that oscillation probability is a function of neutrino energy and length traveled. In our case, the length traveled through the Earth from atmospheric production is a direct function (cosine) of the arrival angle of the neutrino. The neutrino energy is analyzed within a histogram with a logarithmic x-axis to emphasize the low energy (20-30 GeV) region where the oscillation signal is the strongest, while still including data from higher energies more accessible to IceCube.

The variable NCh (the number of DOMs with hits) is shown because it correlates with energy, and is available in both the simulated and experimental sample at all cut levels. It shows basic agreement between the experimental and simulated data, with better agreement in the signal region at higher cut levels.

More variables from each cut level are available in Appendices A-C.

Finally, the predicted rates for each type of simulated data are shown at each cut level. While the final cuts (especially the choice of where to place the final zenith angle cut) are tested with Round Trip Trials (Ch 9) and not through methods described here, the results are included here for completeness.

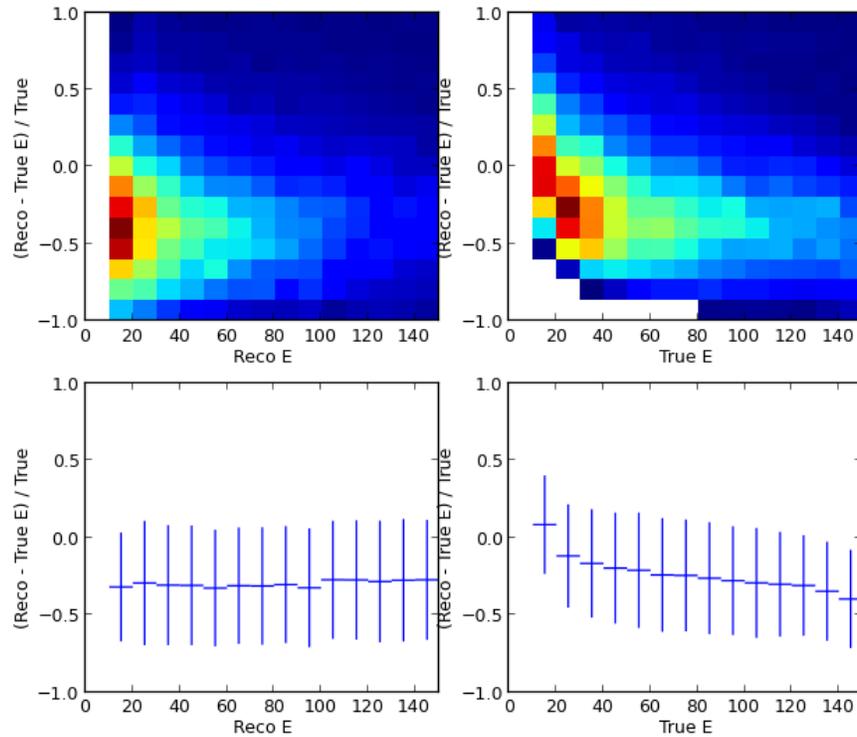


Figure 7.6: Event rate versus energy reconstruction error for Monopod. The y-axis on all the sub-plots shows the fractional energy error, i.e. $(\text{reconstructed} - \text{true})/\text{true}$. The left and right upper plots show the reconstructed and true energies, respectively. Each of these is worth considering, for different reasons. The plot of fractional energy error versus true energy tells us if there are any problems in isolated energy regions of the simulation. The plot of fractional error versus reconstructed energy tells us how much to trust the values returned from the reconstruction, which is of paramount importance to later analysis. The two lower plots are profiles of the two upper plots; the size of their error bars is the RMS of each column of the upper plots. The error bars in the lower-left plot give us the resolution of the sample as a whole: they are all between 0.32 and 0.37, centered on 0.35. We report 0.35 as the fractional energy error, and use this (35%) later as the energy bin size.

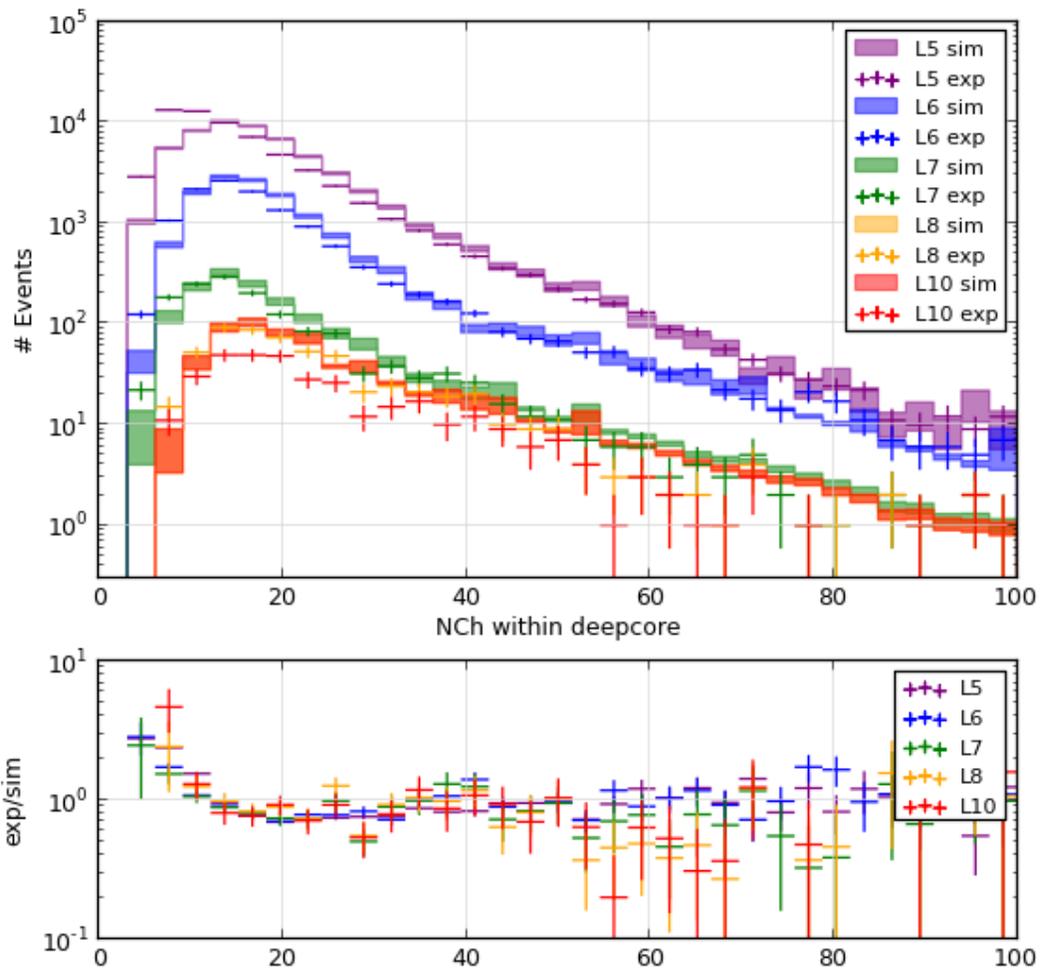


Figure 7.7: NCh at all cut levels. The cut levels are defined throughout the preceding chapter. Hits are only considered within the DeepCore fiducial region, but no other hit cleaning is applied.

Figure 7.8: Monte Carlo true energy at all cut levels. The cut levels are defined throughout the preceding chapter.

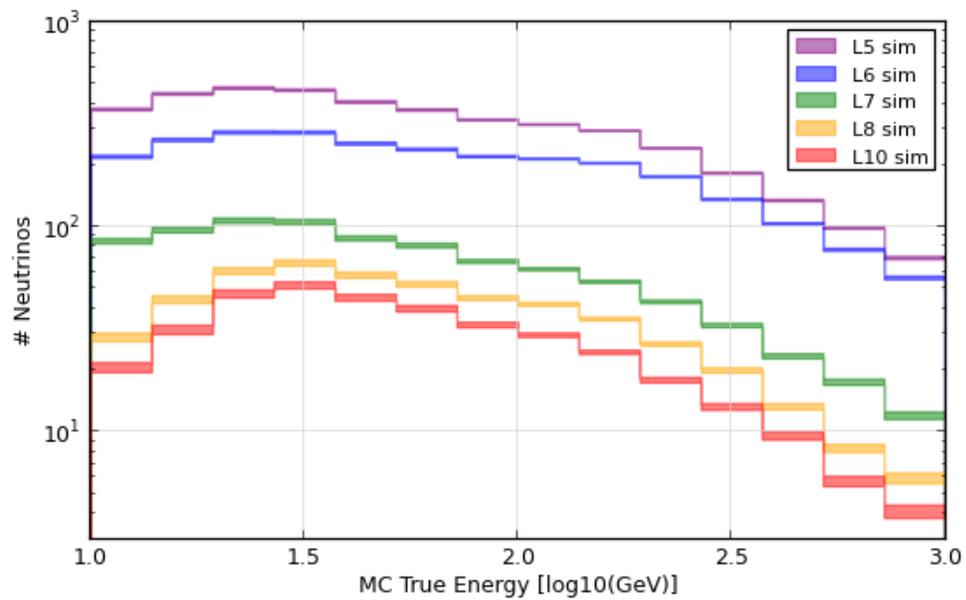
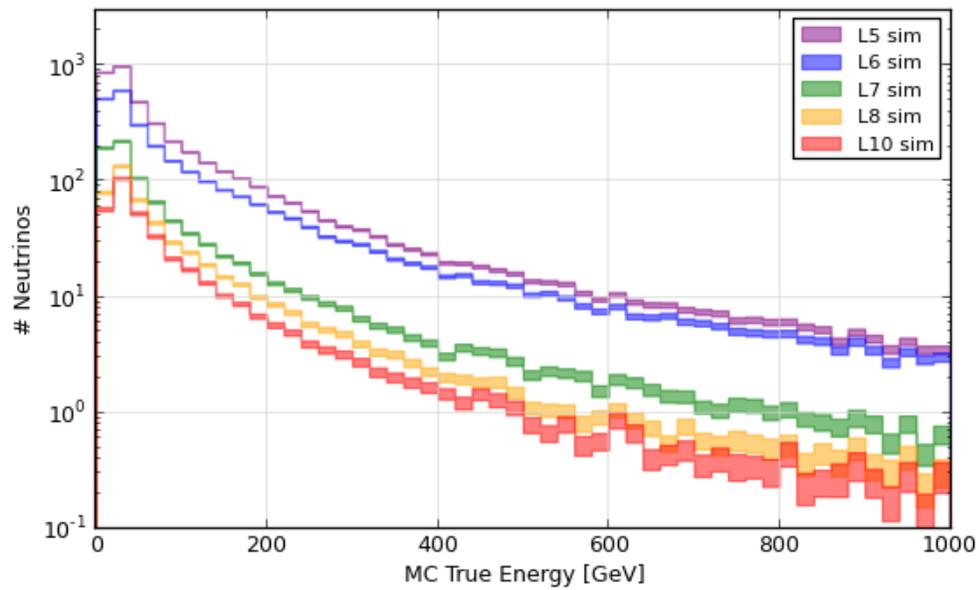


Figure 7.9: Monte Carlo true zenith angle at all cut levels. The cut levels are defined throughout the preceding chapter.

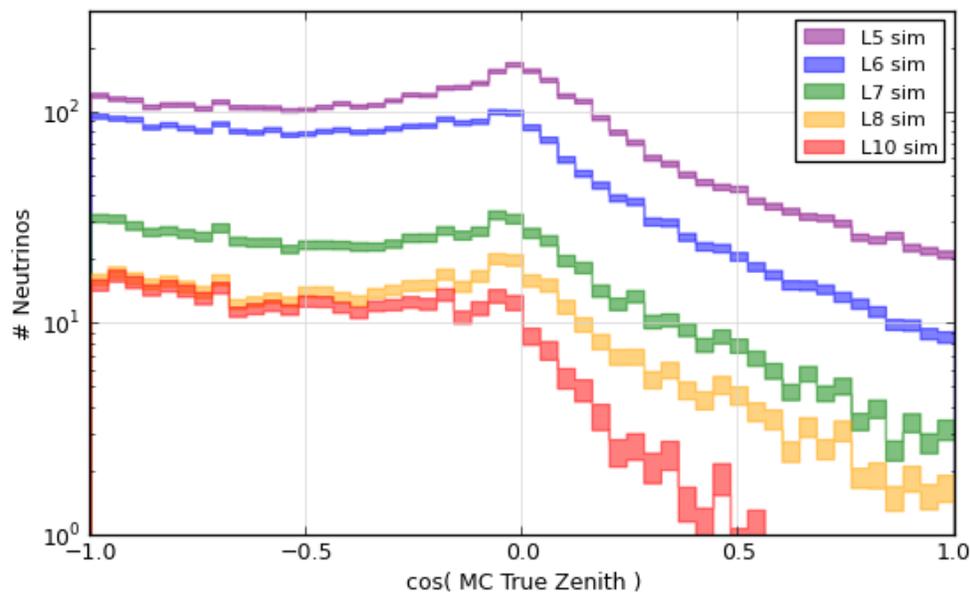
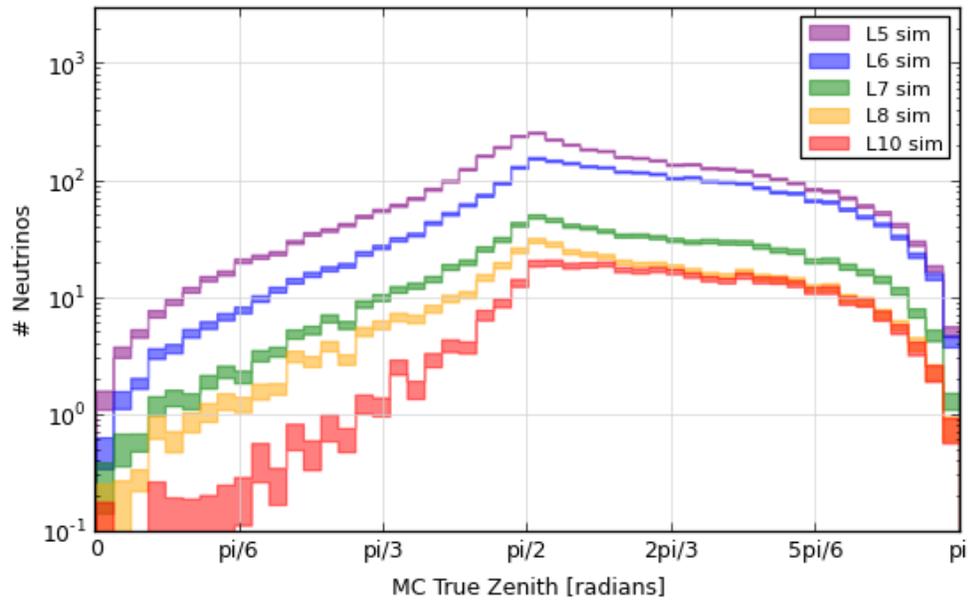


Figure 7.10: Reconstructed energy at all available cut levels. The cut levels are defined throughout the preceding chapter. The reconstruction used here is Monopod.

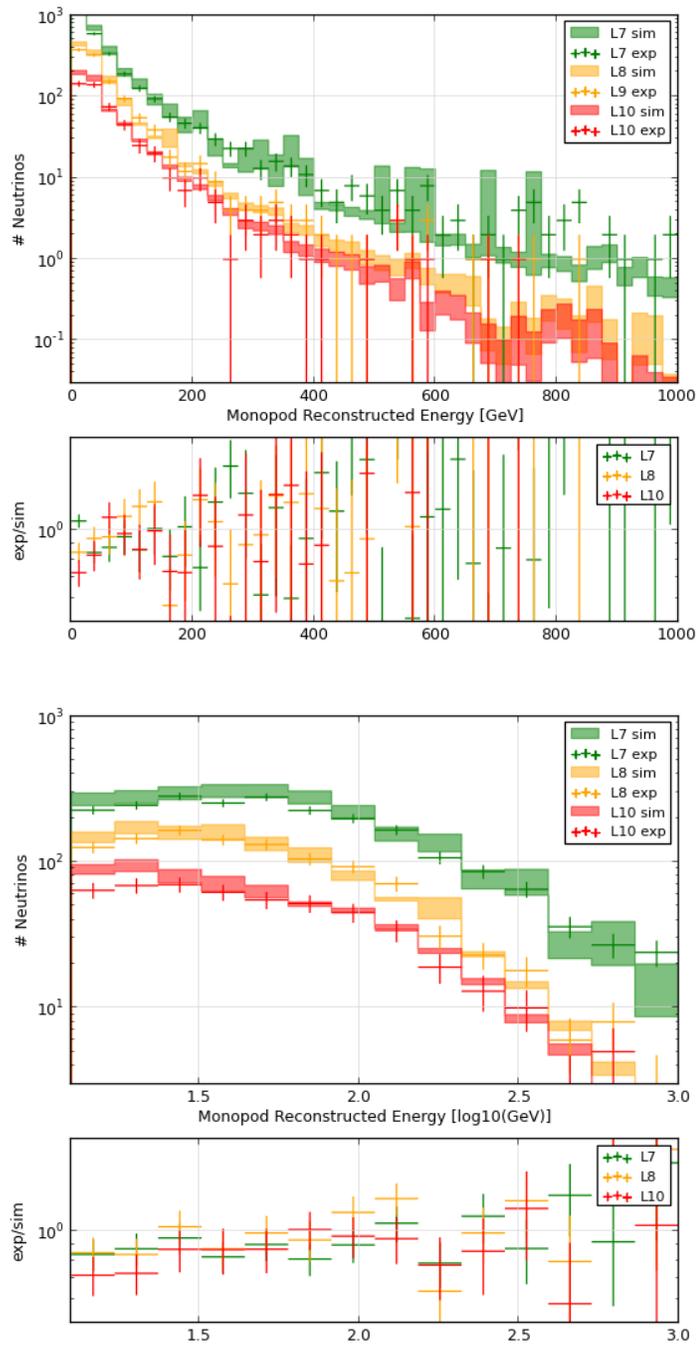


Figure 7.11: Reconstructed true zenith angle at all available cut levels. The cut levels are defined throughout the preceding chapter. The reconstruction used here is Igelfit.

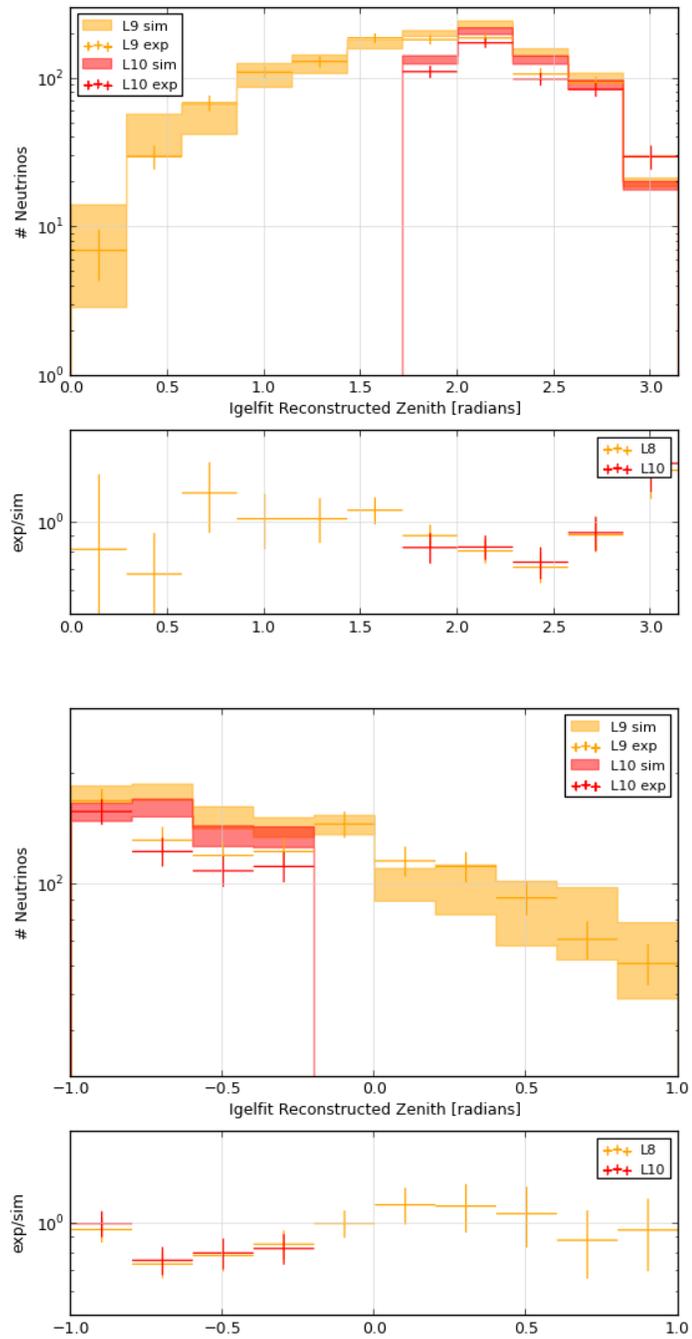


Table 7.1: Event rates passing each cut level for each type of data. For simulation, the rates are weighted to an atmospheric energy spectrum, and total weights are scaled to the livetime of the data: 7500.71 hours.

Cut Level	Genie ν_μ	NuGen ν_μ	Sum ν_μ	Genie ν_e	NuGen ν_e	Sum ν_e	Genie ν_τ	CORSIKA μ	Exp
L5	14897	13452	28348	11214	958	12172	900	501899	603218
L6	8773	9901	18674	6602	636	7237	592	105578	118830
L7	2096	2389	4485	2269	198	2467	159	6673	14140
L8	1175	1267	2442	1299	74	1374	71	1264	5247
L9	1175	1267	2442	1299	74	1374	71	1264	4227
L10	688	800	1488	647	30	677	58	242	2085

Table 7.2: Number of events passing each cut level for each type of data. These are listed to give the reader an approximate sense of the statistical errors in the simulation. For example: the cosmic ray muon background sample (CORSIKA) has much greater generated livetime than the signal, but still has the largest statistical error at the higher cut levels. To translate these values into precise statistical errors, one would also need the energy-dependent relative weighting of each event.

Cut Level	Genie ν_μ	NuGen ν_μ	Genie ν_e	NuGen ν_e	Genie ν_τ	CORSIKA μ	Exp
L5	570937	491160	260552	260552	107478	18654	603218
L6	360276	324155	162593	162593	64422	3924	118830
L7	86866	99814	53857	53857	16737	248	14140
L8	43617	58125	31607	31607	8780	47	5247
L9	43617	58125	31607	31607	8780	47	4227
L10	30270	39894	18497	18497	5387	9	2085

Chapter 8

Analysis Method

This chapter describes the process of using a sample of events collected by IceCube to study neutrino oscillation parameters. The rough outline of this analysis is that we compare the experimental data to simulated data, and by varying the simulation input of the oscillation parameters, we find which values produce the model that most closely matches the data.

The comparison between each model and the data is performed using a maximum likelihood technique (§8.1). This likelihood fit includes nuisance parameters specific to this analysis (§8.2). Once a likelihood is computed for each hypothesis test point of the oscillation parameters, a significance is found for the likelihood at each point: the process of figuring out the size of the error bars. This is discussed first in the general context of Frequentist confidence regions (§8.3), and then in the more specific case of Feldman-Cousins statistics (§8.4), which treats the statistical errors of a signal with low statistics and physical bounds on the allowed region more formally. Finally, some specifics are discussed (§8.5) about the software used to perform this analysis and how it has improved on previous analysis software.

8.1 Maximum Likelihood Technique

The technique of maximizing a likelihood function to find the best-fit simulation for a given experimental dataset is ubiquitous in modern experimental physics. A general likelihood function combines observed data with a hypothesis fit that could explain the data: in IceCube, one example is explaining the rate of data in each bin of a histogram with predictions for that rate from a simulation-generated histogram.¹ The likelihood is extremized for the best possible fit. In general, the value of a single likelihood does not mean much (it depends a lot on your binning choice), but comparing (by ratio) two different likelihoods tells you if your hypothesis changes are improving or degrading the quality of the fit.

More specifically, the usual situation for Poisson likelihood functions is the following: take two histograms, one for data and one for Monte Carlo (one may imagine this in a single dimension for simplicity, but the main analysis of this thesis uses two). For each bin, take the bin value in the MC histogram as the center of a Poisson distribution.

¹Another common example of likelihood fitting within Icecube is in particle reconstruction: a hypothesis particle could explain the observed light. See more on reconstruction in Ch. 5

The probability for that bin is the probability of getting the data value as a Poisson fluctuation around the Monte Carlo value. Do that for every bin, and then multiply the probabilities together. If you like, you can do that all in log space to make the computations easier. That's your likelihood function. If you evaluate that with two different Monte Carlo hypotheses, you can take a ratio to see which hypothesis describes the data better.

The analysis in this thesis uses a Poisson likelihood comparison of a two-dimensional histogram in data and simulation. The two dimensions are reconstructed zenith angle and reconstructed energy (more specifically, $\cos(\text{zenith})$ and $\log_{10}(E)$). The comparison uses the Poisson probability mass function:

$$P(\text{exp}|\text{sim}) \equiv P(k|\lambda) = \frac{\lambda^k}{k!} \cdot e^{-\lambda}$$

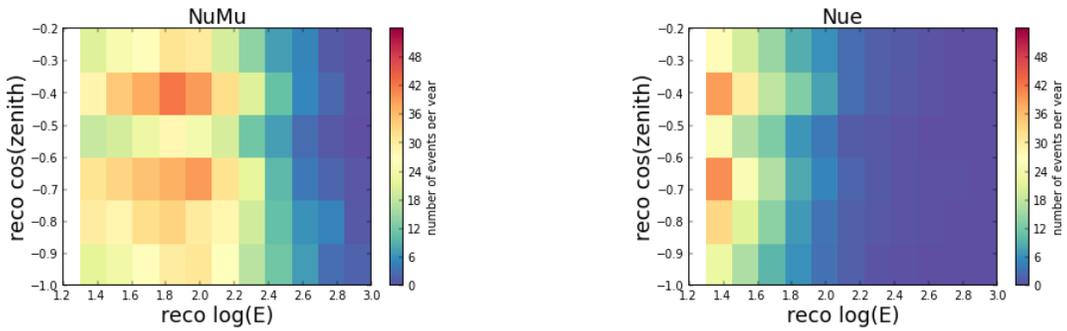
where, for each bin, the expected mean λ is given by the Monte Carlo number of expected events. The probability is evaluated at k , the experimentally observed number of events. This method noticeably does not include a term for the statistics on the Monte Carlo: if we generate double the Monte Carlo statistics, we get a better description of the physics input to the simulation, but that improvement is not reflected in the likelihood function. Some discussion of a method to include such statistics can be found at [129] or [130], but my analysis uses the standard Poisson likelihood.

Evaluating this probability at each bin of the zenith-energy histogram, the likelihood function is the product of all these probabilities. For ease of computation, we work with the \log_{10} of the likelihood function, so instead of multiplying the probabilities together we can add them.²

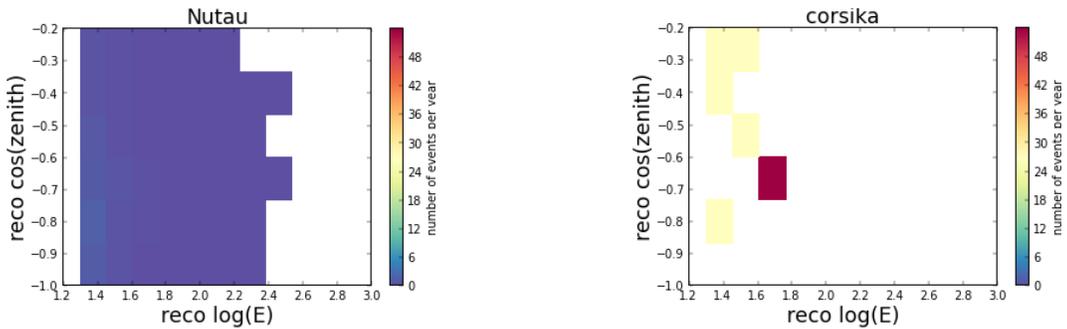
$$LLH = k \ln(\lambda) - \lambda - \ln(k!)$$

The prediction for the Monte Carlo rates histogram is constructed as a sum of several components, representing our best estimate of the true rates. This includes (as described in §7.2) contributions from ν_μ , ν_e , ν_τ , and cosmic ray induced μ . Each of these contributions is plotted in Fig. 8.1. Note that to make these histograms, we must choose appropriate limits in zenith angle and energy; the choice of where to put these limits includes the full analysis chain described in this chapter, and is discussed in more detail in Ch. 9. The histograms as shown in Fig. 8.1 also include an implicit choice of oscillation parameters. Generally, events weights are calculated assuming maximal oscillations, then the weights are reduced to reflect various hypotheses of oscillation parameters. The oscillation survival probability formula used in this analysis assumes two-flavor oscillations and neglects matter effects. The analysis was also tested with three flavor oscillations; the result was negligibly different but the calculation time was much greater. For illustration in this chapter, oscillation parameters are set to the global best fit [21] atmospheric values: $\sin^2(2\theta) = 1.00$ and $\Delta m^2 = 0.0024$. The sum of these four components is the Monte Carlo prediction for the case of those oscillation parameters.

²A note on software optimization relative to this formula: consider for a moment the $k!$ term. This is an expensive computation, and for the number of repetitions required, it's worth optimizing. One trick is to use an optimized function to do the log and the factorial together, e.g. the `scipy.special.gammaln()` function, or in equivalent time, the `scipy.stats.poisson.logpmf()` function. An even better solution is to skip that whole term, because it only depends on the data. Since we always compare Monte Carlo variations to a single experimental dataset, we do not need to evaluate a term on the dataset again for every variation.



(a) NuMu contributions to the total rate prediction; the numu contribution is 1271 events. (b) NuE contributions to the total rate prediction; the nue contribution is 512 events.



(c) NuTau contributions to the total rate prediction; the nutau contribution is 10 events. (d) Cosmic ray muon contributions to the total rate prediction; the contribution is 189 predicted events, from 7 simulated events

Figure 8.1: Component histograms of the simulated rates prediction. The neutrino samples each included a mix of Genie and NuGen generated data (as described in §7.2) and the cosmic ray muon sample is generated by CORSIKA. The livetime assumed for these figures is the full IC79 livetime. The oscillation parameters assumed are $\sin^2(2\theta) = 1.00$ and $\Delta m^2 = 0.0024$. The sum of these four components is the predicted rate histogram. The white bins contain no simulated events.

8.2 Nuisance Parameters: Physics and Systematics

The previous section described the method for computing the output of an single likelihood evaluation; the process of maximizing the likelihood function is described in this section. This is achieved by varying the input parameters to the Monte Carlo. These inputs are separated into two types: physics parameters (θ_{23} , Δm_{23}^2) and nuisance parameters (which describe systematic effects in the detector). The nuisance parameters can be further divided into two types: those that vary continuously and those that vary discretely.

The continuous variables are built into the likelihood function and varied by the same minimizer as the physics parameters. In this analysis, the continuous nuisance parameters include the relative normalizations of the various Monte Carlo samples: ν_μ , ν_e , ν_τ , and cosmic ray mu . The the slope (spectral index) of the input cosmic ray spectrum was also incorporated into the fitting software, but was not allowed to vary because it was not found to be a limiting error within the timeline of this thesis. The muon- and tau-flavor neutrino samples are constrained to the same normalization, and the electron-flavor neutrino sample contributes a Gaussian penalty to the likelihood function as it gets farther from the muon and tau neutrino normalization. The muon and tau neutrino, electron neutrino, and cosmic ray muon background samples are each constrained to stay between a factor of 0 and a factor of 2 of their original Monte Carlo expectation.

The discrete variables cannot be included in the likelihood function because it would create steps and walls in the likelihood space, which would confuse the minimizer and make it get lost and fail. Instead, we run a new simulation set for each variation of a discrete variable, evaluate the likelihood using that dataset, and then compare the likelihood outputs of the nominal simulation to all the variations, taking the best case when we calculate our limits.

The discrete systematics variables considered include: the ice model, the light collection efficiency of the DOMs, the relative efficiency of the high-quantum-efficiency DOMs and the regular DOMs, and the radial function describing light scattering around each string (the “hole ice”).

Simulating these several discrete-variable datasets is computationally intensive, so, like most of the simulation, this simulation step is performed centrally within the collaboration. We consider only one discrete systematic variation for each simulated dataset, so we do not simulate all the combinations. There is no simulation with a different ice model and also a varied hole ice scattering model. This did cause a problem when, partway through the systematics set generation, we found out that the nominal DOM efficiency was too low, and our data fit better to simulation if we used 110% of the nominal value.

Another complexity stems from the fact that a single Monte Carlo prediction requires input from several different Monte Carlo generators. Each discrete systematic dataset therefore needs to be simulated using each generator, with the neutrino generators running a separate set for each flavor of neutrino. For some systematic variations, it was not a high enough collaboration priority to simulate all the flavors for all the variations, so we asked for only numu files, and assumed the nominal rate for the other flavors and for background. This was especially true for background cosmic ray

samples, where simulating enough Monte Carlo to see an effect after event selection could take hundreds of thousands of times longer than it would for the signal (neutrino) samples. So we looked at numu with highest priority.

While many of these simulated data sets were generated, they were not included in the final analysis. At the time of their generation, the improvements to the noise model were not yet incorporated into the simulation, and the noise model was since confirmed as a larger systematic change to the data than any of the systematic variations that were simulated. New simulations incorporating both the new noise model and suites of systematic variations are currently underway.

The result of the computations described so far is a log likelihood map: a map of the parameter space in θ_{23} and Δm_{23}^2 of the likelihood of our Monte Carlo describing our data. At each point, we have varied the Monte Carlo systematics inputs (keeping the physics ($\Delta m_{23}^2, \theta_{23}$) constant at the test point values) to maximize the likelihoods. This includes checking all the discrete systematics datasets as well: we have to run a minimizer for the continuous nuisance parameters, save a likelihood value, and then run another minimizer over a different dataset for each discrete systematic. The best of these is the reported Llh value. Among all of these Llh values for all of these test points, the smallest one is our best fit point.

We can vary the step size of this minimization to balance two aims. First, we want the search to run quickly, with a minimum number of steps, which can be achieved by using a larger step size between trials. Second, in contrast, we want to arrive at the true minimum, which is achieved better by using smaller step sizes. To test this balance, I studied the effect of changing the step size on a single likelihood point, using a “data challenge” of simulated expected event rates combined with random variations. The physics parameter inputs were fixed at maximal mixing and world best-fit Δm^2 . When the step size became too large, the likelihood did not change at all from the likelihood of the seed parameter values; this is consistent with the step size being larger than the scale of the dip in the likelihood function around the true minimum. After a region of smooth transition, the likelihood difference leveled off at a single value (presumably the true minimum), retaining this value even as the step size decreased further by several orders of magnitude. The optimal value is right at the start of this leveling off, where the true minimum can be found, but the step sizes are still large enough to find the minimum as quickly as (accurately) possible. For this likelihood function and these parameters, that value was found at 0.1, as can be seen in Fig. 8.2.

8.3 Frequentist Confidence Region

From the previous sections, we get a map of likelihood values in the space of physics parameters θ_{23} and Δm_{23}^2 . The significances of these likelihood values, and thus limits on θ_{23} and Δm_{23}^2 , is the topic of this section. The limits are the “confidence interval” in the Frequentist model of statistics.³

³Frequentism, as a school of statistical thought, is often presented in contrast to the Bayesian approach [131]; this analysis uses only Frequentist constructions.

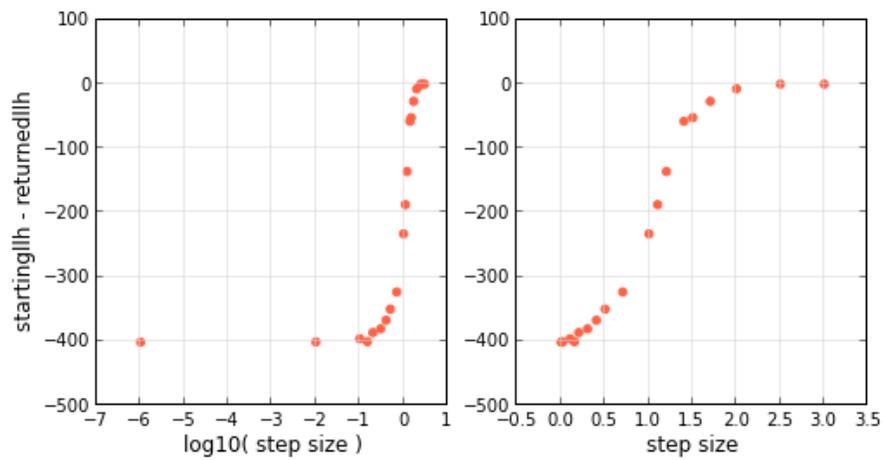


Figure 8.2: Likelihood minimization step size optimization. The x-axis in both plots shows the step size being tested; on the left this is plotted on a log scale, and on the right it is plotted on a linear scale. The data points are the same in both plots. The y-axis shows the improvement that the minimizer was able to achieve from the seed to the final output. Note that this improvement looks negative, which is expected because the likelihood function increases for “better” values (this is called, after all, the “maximum likelihood technique”) but the software package works by minimizing. In ambiguous cases, I try to use “better” in this text rather than “min” or “max”.

For a physics parameter θ , a frequentist 90% confidence interval (x_1, x_2) is the central interval that contains the true value of θ for 90% of many hypothetical repetitions of an experiment.

Often, such test statistics are distributed in predictable ways, so changes in the Llh function are often used to directly map the limit contours. Wilks theorem [132] states that in certain cases, a test statistic (our Llh) with n degrees of freedom will be distributed like a an n -dimensional chi-squared distribution, which is easily describable. Thus, if we can show that Wilks theorem applies (if we can show that our Llh is distributed overall like a 2-dimensional χ^2), we can simply draw contours around the best fit point through the points where the likelihood drops off by key values: 1σ where it has dropped by 2.3, 2σ where it has dropped by 6.2, etc. More of these values can be found in Table 8.1.

Note that this uses the standard convention for reporting significances of translating from p-values to “sigmas” using a Gaussian, normal distribution with width 1 sigma. Thus, we are interested in the 68% contours because a normal distribution contains 68% of its area within 1σ of its peak.

In the cases where it cannot be shown that the test statistic follows a χ^2 distribution, a more elaborate method is needed. This method is the subject of the following section.

Table 8.1: Conversions between significance in σ , central p-value, and (using Wilks’ theorem) difference in log-likelihood values.

Gaussian Significance	two-sided p-value	1 minus p-value	or 1 in every...	llh difference
1σ	0.68268949	3.173 e-01	3.15×10^0	2.2957
2σ	0.95449974	4.550 e-02	2.20×10^1	6.1801
3σ	0.99730020	2.699 e-03	3.71×10^2	11.829
4σ	0.99993666	6.334 e-05	1.58×10^4	19.334
5σ	0.99999943	5.733 e-07	1.74×10^6	28.744
6σ	0.99999999	1.973 e-09	5.07×10^8	40.087
7σ	0.99999999	2.559 e-12	3.91×10^{11}	53.382
8σ	0.99999999	1.221 e-15	8.19×10^{14}	68.677

8.4 Feldman-Cousins Error Contours

This section describes the standard method prescribed by Feldman and Cousins [133] for determining, independently of Wilks’[132] theorem, the frequentist confidence regions of a likelihood function. Conveniently, their paper

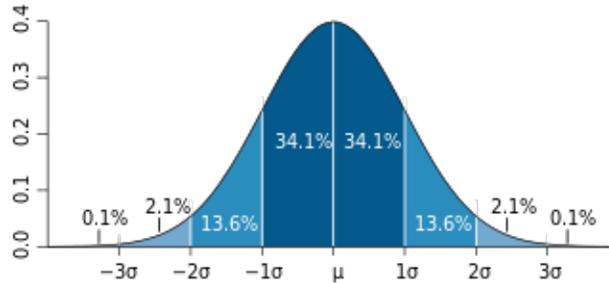


Figure 8.3: Normal distribution with significances labeled

was written specifically with neutrino oscillations in mind, and several subsequent oscillation papers include more detailed, lucid explanations of how to implement the method. For instance, I found [134] and [135] remarkably helpful while developing analysis software.

We have a log likelihood map⁴. We want to know how significant these likelihoods are, and whether a constant likelihood difference represents a constant significance of the result.⁵ For each test point, we generate many simulated trials: the heart of the Feldman-Cousins method is that, instead of assuming a distribution of the Llh function, we map it out with simulation.

In more detail, the procedure is as follows (for a more visual, pseudo-code version of this discussion, see Fig. 8.5). We call the global best fit to the data $(\Delta m^2, \theta)_{\text{globalbest}}$. In the process of finding the global best fit, we also found best fit nuisance parameters for that point: $(\text{nuisance}_{\text{globalbest}})$, which are not necessarily the best fits for any other test point, and a likelihood value $\text{Llh}_{\text{globalbest}}$. The central question is: for each other test point of physics parameters, how often is it possible that through random variations we could produce a likelihood as good as the one we found at the global best?

The phrasing of that question should suggest the following method. We loop over physics test points, choosing a particular one for discussion (without loss of generality): $(\Delta m^2, \theta)_{\text{test}}$. At that physics test point, we vary the nuisance parameters to maximize the likelihood; that defines $(\text{nuisances}_{\text{test}})$. Plugging these two into the Monte Carlo, $(\Delta m^2, \theta)_{\text{test}}$ and $(\text{nuisances}_{\text{test}})$ gives us a best predicted rate, an expected experimental histogram in reconstructed zenith angle and energy. The units of that histogram are total expected events per livetime (about 1 year).

We generate many pseudodata samples by taking random Poisson fluctuations of the rate in each bin of this histogram⁶. We run each pseudodata sample through the same software as we did to find the global best fit. That is,

⁴The likelihoods are represented as the colors of each point in e.g. Fig. 9.1.

⁵We want to know the significance of each point, and whether or not a constant color represents a constant significance.

⁶Note here the difference between this procedure (Poisson resampling) and the procedure described in [129][130]: the resampling described in these references considers the simulated weight of each event, and constructs a new random sample from a subset of the full simulation, drawing each event into the sample with probability according to its weight. By contrast, the procedure I used follows the more common practice of ignoring the statistical limitations of the simulation sample, and making variations by only considering the central value of the simulation prediction.

first we find a best fit to the pseudodata in all parameters $(\Delta m^2, \theta, \text{nuisances})_{\text{pseudobest}}$, and from that rate calculate a likelihood $Llh_{\text{pseudobest}}$. Second, we calculate the rates and likelihood from the pseudodata taking physics parameters of the test point $(\Delta m^2, \theta)_{\text{test}}$ and letting the nuisance parameters float to maximize the likelihood.

$$r_i = \left(\frac{Llh(\text{pseudodata}, (\Delta m^2, \theta)_{\text{test}}, \text{nuisances})}{Llh(\text{pseudodata}, (\Delta m^2, \theta, \text{nuisances})_{\text{pseudobest}})} \right)$$

This likelihood ratio is one trial of one test point; with many trials of that test point we can make a histogram representing the probability distribution that true physics parameters from the test point could be mimicking the data we observed. So, we do many pseudodata trials, returning a likelihood ratio (more precisely, a log likelihood difference) for each trial. We histogram those differences, taking special note of where one value falls in the distribution:

$$R_i = \left(\frac{Llh(\text{globaldata}, (\Delta m^2, \theta, \text{nuisances})_{\text{test}})}{Llh(\text{globaldata}, (\Delta m^2, \theta, \text{nuisances})_{\text{globalbest}})} \right)$$

Within the distribution of the many r_i , the placement of R_i tells us the p-value of our test point. An example of such a distribution (from another experiment) is shown in Fig. 8.4. The fraction of r_i with worse likelihoods than the point we are considering gives the p-value of that point. Most often, we're interested in finding and reporting the points with p-values (listed in Table 8.1) that represent integer sigma values using the Gaussian sigma p-value convention. If exactly 68% of the r_i are smaller than R_i (if exactly 68% of the pseudodata trials gave worse likelihoods than our global best fit null hypothesis), then our test point $(\text{nuisances}_{\text{test}})$ is on the 1σ contour line. If the p-value is smaller, our test point is outside of the 1-sigma contour; if the p-value is larger, our test point is within the 1-sigma contour.

By testing many points using this procedure, we can map out the significance contours on the likelihood map. For many oscillation analyses, a contour map of this style is the preferred method for reporting final results. The location of the 1-sigma or 90% confidence contours in one particular variable is also often quoted as a range, so the results do not have to be purely graphical. In the early stages of each experiment, eliminating the no-oscillations hypothesis with high significance is also common.

For a comparison between results using Feldman-Cousins and Wilks statistics, see §9.6.

8.5 Software Specifics

The procedure described over the preceding several sections gives a general idea of how to construct a likelihood analysis with significance contour results, with minimal specific description of the analysis in this thesis. This section describes more details specific to this analysis.

We then repeat this procedure for all interesting test points, taking advantage of a fantastic set of cluster computers. The number of pseudodata trials at each test point determines how long the job takes to run: 150 trials on my laptop take about 10 minutes, and for each point we run between 1 thousand and 10 thousand trials. We determine how many trials to run dynamically (after the minimum is reached): we keep track of the number of trials producing an r_i in the tail past R_i , and when that passes a configurable number (defaulting to 10), we stop simulating. This means that test

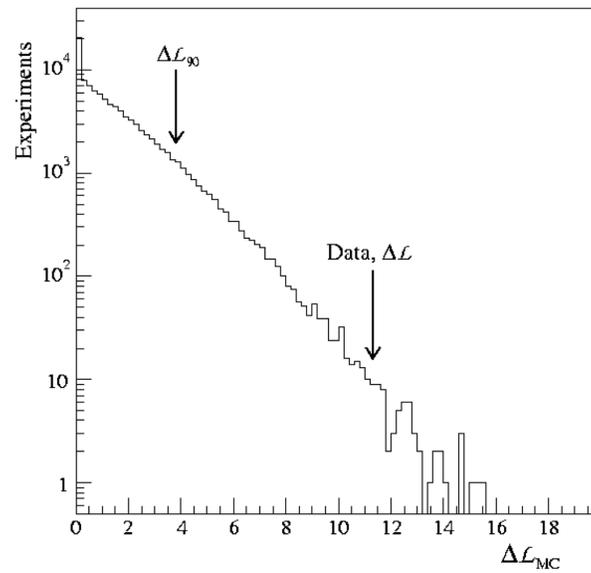


Figure 8.4: Example histogram of pseudodata trials outcomes for Feldman-Cousins contour construction. This histogram considers the no-oscillation case as a test point. The fact that such a small fraction of the trials falls beyond the likelihood difference of the data (marked “Data, $\Delta\mathcal{L}$ ”) indicates that the data support, with high significance, the existence of oscillations. From [134].

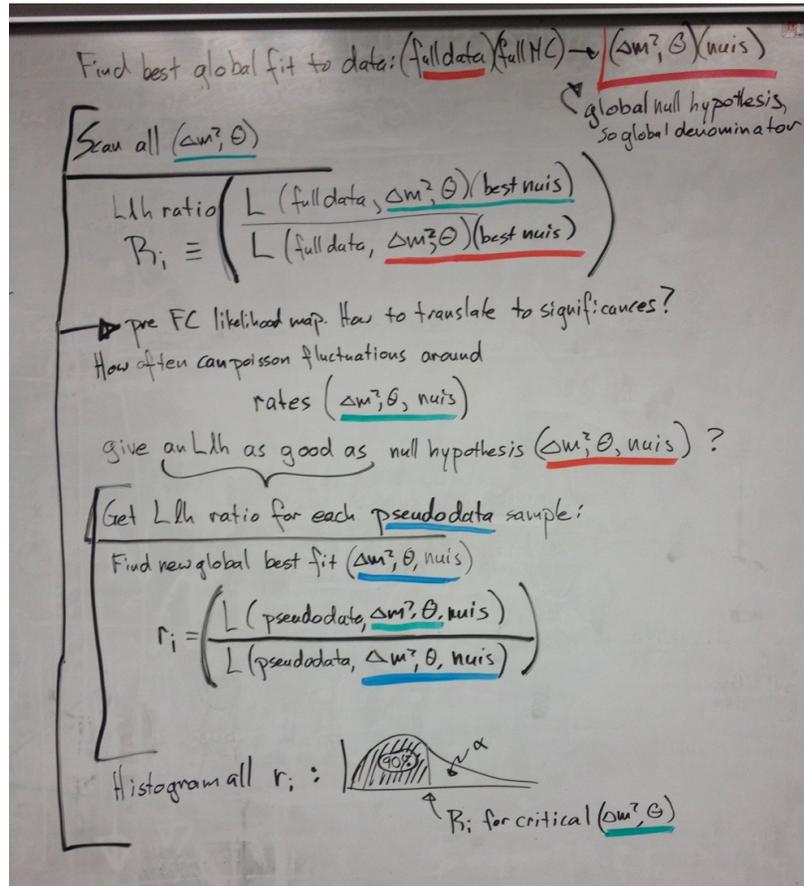


Figure 8.5: Example pseudocode explaining the Feldman Cousins statistical method. Colors are used to distinguish different points at which the oscillation parameters and nuisance parameters are evaluated. The first colored point is red, showing the global best fit point of the observed data compared to all options in the simulation; this is the point around which we are trying to draw error contours. The next color is green, showing a test point in the oscillation parameters. At each green test point, a likelihood can be evaluated and compared to the best fit point, but this is not yet a significance. To find the significance of this green point (again, there would be many green points tested), we assemble a histogram of pseudo-data samples, each made from Poisson variations around the “true” predicted rates using the green points. The best fit of each pseudo-data sample is marked in blue. For each blue pseudo-data best fit point, the likelihood ratio r_i is calculated and becomes an entry in the histogram. The likelihood ratio for the (green) test point is marked (big) R_i ; the fraction α of pseudo-data trials beyond that critical R_i is the significance of the (green) test point. By testing many green points, the experimenter can trace out contours, in oscillation space, of any desired significance, limited only by computing time. As a cross-check for the reader, if you have understood this correctly, you should see that the nested nature of these searches means that many, many trials are required to make a contour plot.

points further from the best point take longer to test; this is one motivation for making a close-cropped plot around the contour lines of primary interest. We do run one test point at the no-oscillation null hypothesis, which takes forever to run. The results are shown in Chapter 10.

A note on computational speed: there are several tricks that can improve the implementation of this code. As a rule, whenever there is a matrix multiplication, it is always better to use someone else's method instead of implementing it yourself, because as physicists we are not in the business of optimizing code (you can get factors of thousands in speed from this). Specifically, most loops or double-loops over arrays can be replaced by matrix multiplication, dot products, or some other function from the package "numpy", which have been extensively optimized, tested, and documented in online forums. Other software could also be optimized; at the time of this thesis, python was the best option.

Another trick was to do as much pre-processing of the data as at all possible before looping over it. This means consolidating it to one (or a few) files, applying all cuts in advance, and keeping only the final analysis variables. We went as far as to make histograms of the input simulation data, and work only with the histograms, without further reference to the original files. This makes file reading time negligible relative to Feldman-Cousins fitting and calculation time, which is as it should be.

The minimizer we should use is L-BFGS-B, where the L stands for low-memory (it only saves information about the last few iterations), BFGS are the initials of the authors, and B stands for "boundaries" (it handles boundary information better than a previous version.) The advantages of this method are discussed in [101] and §5.11.

Chapter 9

Round Trip Simulation Trials

A round trip trial using simulation is a technique for testing the limits of any given method, independent of external systematic effects. A regular likelihood test works by comparing two samples bin-by-bin; in a round-trip test, the “experimental” sample is produced by the experimenter, using known input parameters and the same machinery as the simulation sample. This is a way of testing the software against itself: if any problems arise, they must be from within the software being tested. If everything works well, a round-trip test provides a best case scenario knowledge of the system. More confounding effects may be introduced by factors not simulated in the test, but the experiment can never do better with experimental data than it did with round-trip simulation data.

In this chapter, round-trip tests are used to address several questions about the design of this analysis. First, as validations of the software itself, the input oscillation parameter point (§9.1) and statistical livetimes (§9.2) are varied. The next questions probe the quality of the event selection: where to put the final zenith angle (§9.3.1) and energy (§9.3.2) cuts, and how sensitive is the analysis to contamination from cosmic ray muons (§9.4) and tau neutrinos (§9.5) (contamination from electron neutrinos is not examined because the ν_e component normalization is free to float within the fit). Finally, the two statistical methods for determining significance are compared (§9.6), and found to give very similar results in the case of this analysis.

9.1 Input Oscillation Parameters

The central focus of this analysis software is its ability to find the true input values of the neutrino oscillation parameters Δm^2 and $\sin^2(2\theta)$. To test this, we vary the parameters used to produce input data for the round-trip test, then see how well the fitter can recover the varied values. In other experiments, the plots from these tests are often called “jellybean plots” [136] because of the appearance of small, oddly-shaped curves around many test points. Six trials are shown in Fig. 9.1.

These trials succeed, in the sense that each variation of the input point is recovered by the likelihood fitter and reflected in the shape of the likelihood space. One interesting thing to note about these plots is that as the input/best-fit point varies, the variations in the likelihood space are not the same, that is, the space not only shifts but also distorts. Specifically, the errors are smaller when the true (input) value of Δm^2 is larger, which makes the oscillation deficit

appear at higher energies. This makes sense, because higher energies have better energy resolution within IceCube, and better energy resolution translates to better limits on Δm^2 .

9.2 Statistical Increase

A second basic check is to confirm that increasing the statistical size of the sample increases sensitivity (decreasing errors). To test this, we can compare the size of the error bars from two round trip tests with different detector livetime but all other input parameters the same. Such a comparison is shown in Fig. 9.2. The two livetimes considered were 7500.71 hours and 1566.39 hours, representing 100% and 20% of the IC79 detector year, respectively. The precise value of the 20% sample livetime was chosen to correspond to the summed length of detector runs whose run numbers end in 0 or 1, which is not exactly 20% of the livetime.

The results of this test correspond to the expected result: the error contours are smaller with a higher statistics sample. As expected from a \sqrt{N} error approximation, the intervals are about twice as large when the data is five times smaller.

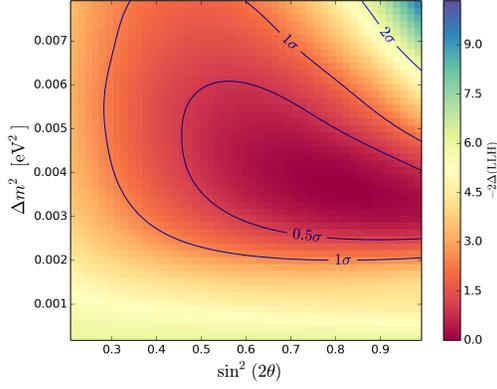
9.3 Placement of the L10 Cuts

The highest (last) level of cuts in this analysis is investigated in this section. These cuts are inherent in the definition of the two-dimensional histogram in zenith angle and energy used for comparing experimental to simulated data. The limits (and to a lesser extent, binning) of this histogram amount to a cut on zenith angle and energy. These cuts are optimized by using round-trip trials to maximize the sensitivity of the analysis. For both zenith angle and energy, the sizes of the bins are determined by the reconstruction resolution of the sample.

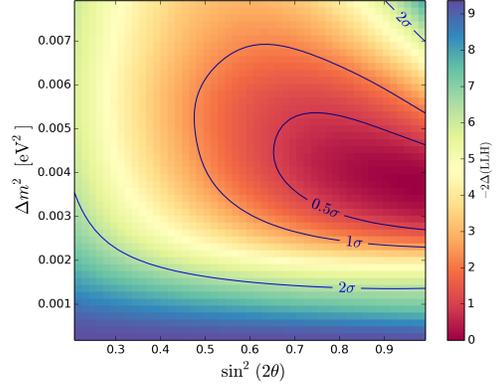
9.3.1 Zenith

The zenith angle cut is considered first. The zenith angle reconstruction, `Igelfit`, has an inherent resolution limit around 10° because of the discretization of the results it returns. The median error measured in simulation (§7.5.3) was 11.8° . The bin size should be related to the energy resolution, and because we use the whole range (a histogram rather than an on-source and off-source region) we use a bin size roughly equal to the zenith angle resolution. An additional complicating factor is that the bins are made constant in $\cos(\text{zenith})$ to give each bin equal solid angle area on the sky, which means the bin sizes are not constant in zenith. To resolve this, the smallest bin is taken as not less than 11.8° , and other bins may be larger.

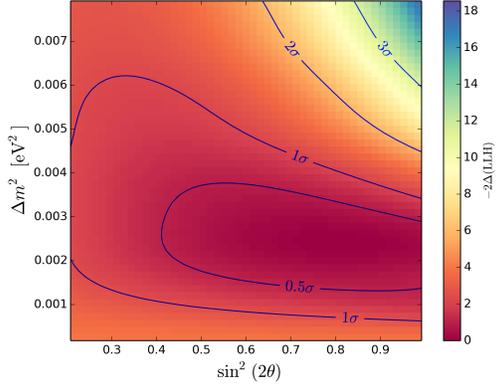
All configurations cut out the downgoing region, only considering events reconstructed as traveling up through the earth. The exact cutoff point was varied between $\cos(\text{zen}) = 0$ and $\cos(\text{zen}) = -0.7$ with step sizes of 0.1. Throughout, harsher cuts on zenith angle made the error bars in the likelihood scans uniformly larger; most likely the effect arose exclusively from the decrease in statistics. Based only on this effect, the ideal configuration includes the



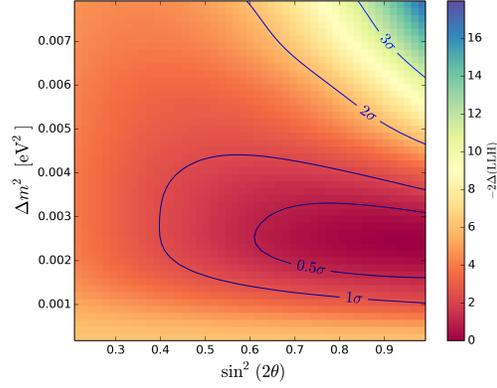
(a) $\Delta m^2 = 0036, \sin^2(2\theta) = 0.80$



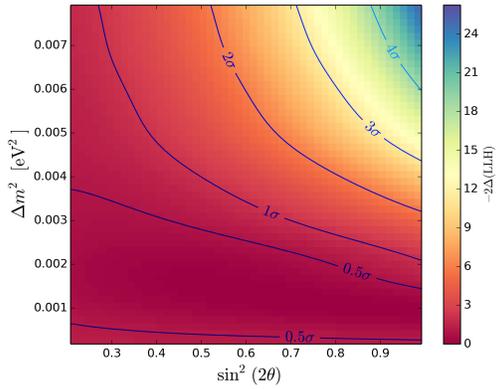
(b) $\Delta m^2 = 0.0036, \sin^2(2\theta) = 0.99$



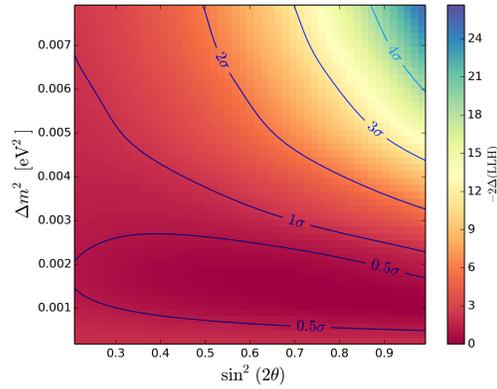
(c) $\Delta m^2 = 0.00235, \sin^2(2\theta) = 0.80$



(d) Expected [21] point: $\Delta m^2 = 0.00235, \sin^2(2\theta) = 0.99$



(e) $\Delta m^2 = 0.001, \sin^2(2\theta) = 0.80$



(f) $\Delta m^2 = 0.001, \sin^2(2\theta) = 0.99$

Figure 9.1: Round trip trials varying the input oscillation parameters. The input values are chosen to approximate 1σ errors. The rate for each sample is scaled to simulate the full IC79 detector year.

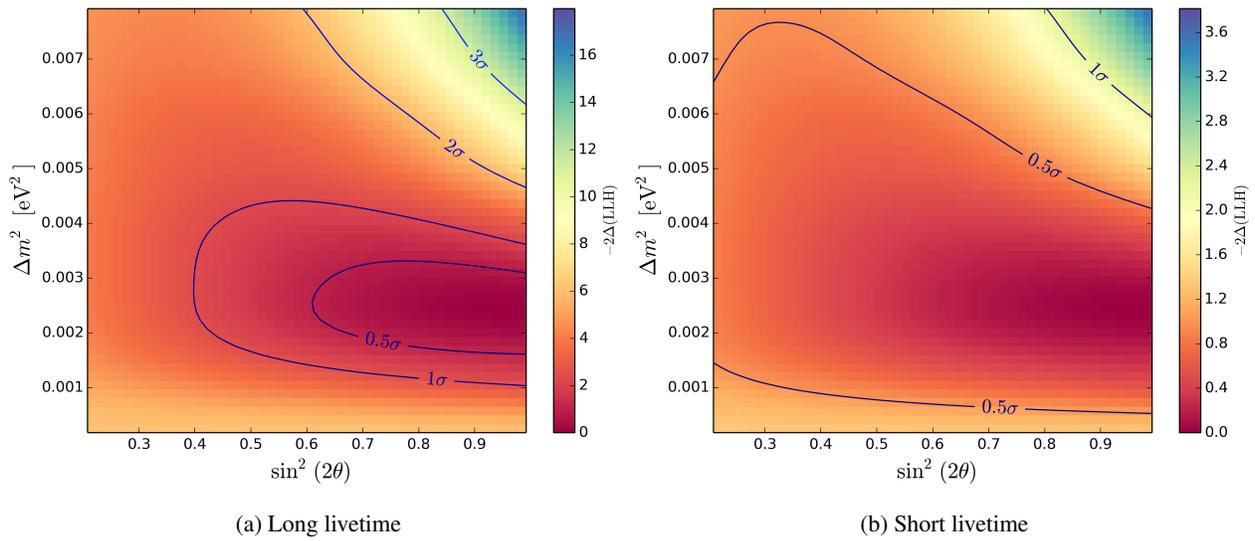
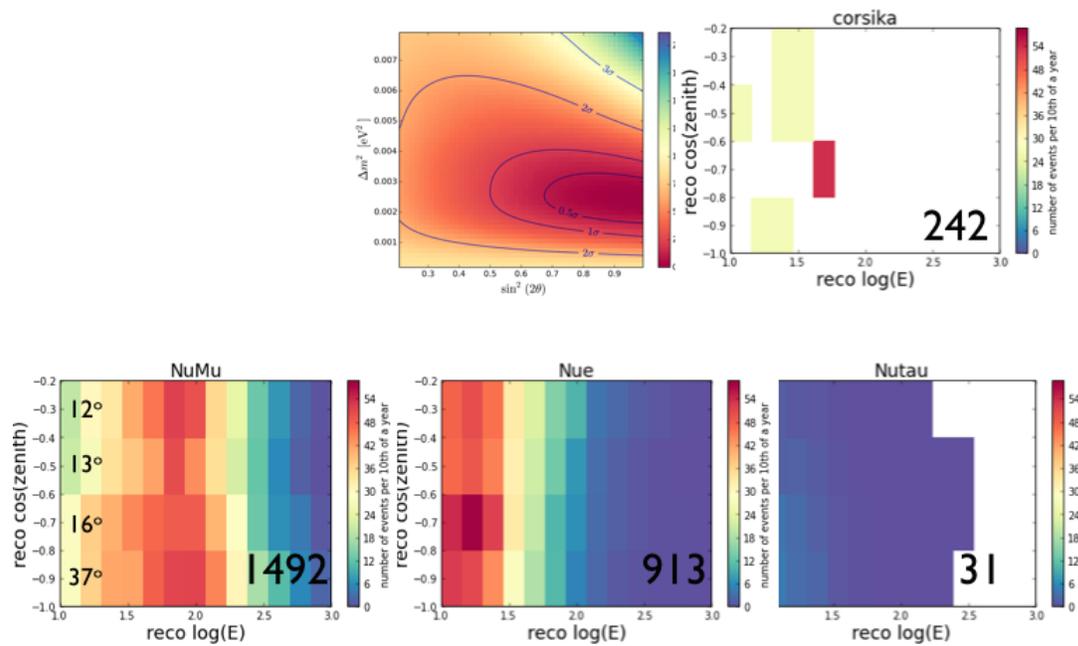
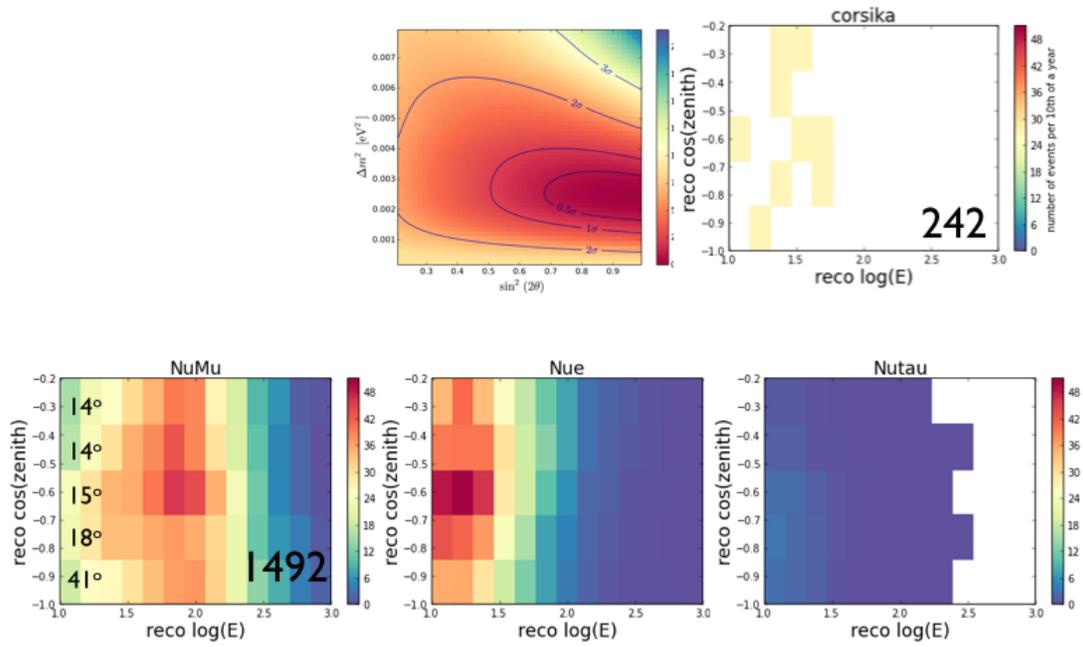


Figure 9.2: Results of two round-trip trials comparing long (7500.71 hours) and short (1566.39 hours) livetime.

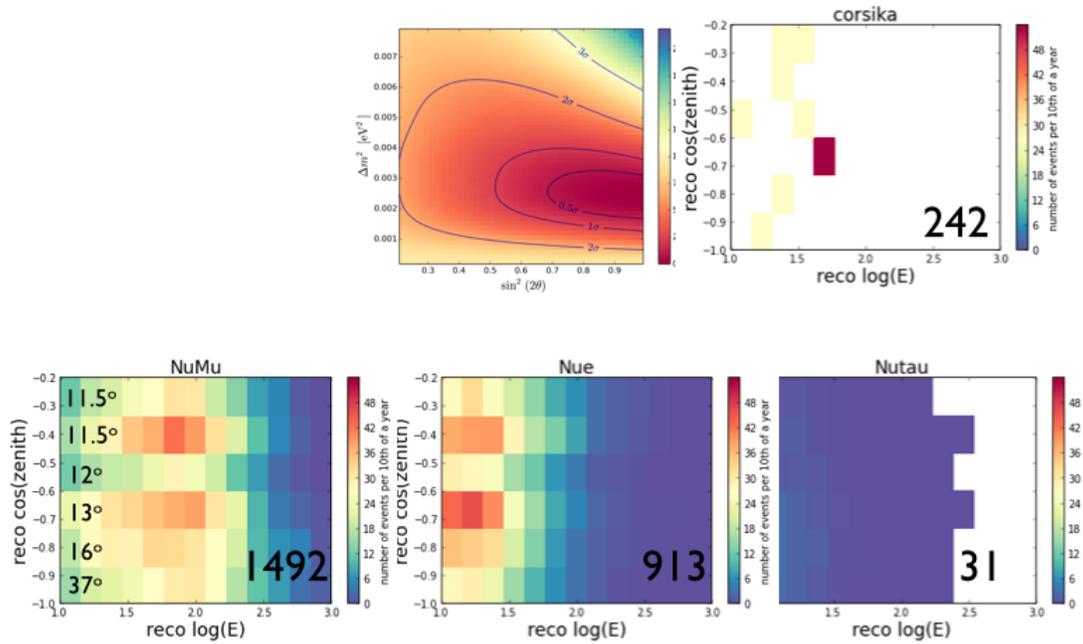


(a) Zenith angle binning variation: 4 bins

Figure 9.3: Results of round-trip trials comparing different binning configurations for zenith. The results from two additional configurations are shown in the following figure. For each configuration, the four components of the simulation rate (ν_μ , ν_e , ν_τ, μ) are shown at expected [21] oscillation values. Each component plot is labeled with the total number of events from that component; because the binning does not change the total number of events, these numbers are constant between configuration. Within the ν_μ plot in each configuration, the bin sizes are labeled in degrees, as a reminder that the bin size is not constant in degrees (but is still constant in solid angle). A scan of the oscillation values is also shown, with Wilks' theorem error bars. The livetime assumed is 7500.71 hours.



(a) Zenith angle binning variation: 5 bins



(b) Zenith angle binning variation: 6 bins

Figure 9.4: Results of round-trip trials considering additional binning configurations for zenith. See previous figure caption for more information.

largest possible region, i.e. cut at the horizon and not any lower. However, there is increased disagreement between experimental and simulated data in several variables at the horizon, so the horizon region is cut out.

Within the restricted zenith angle range of $\cos(\text{zen}) = -0.2$ to -1.0 , tests were performed to investigate the effect of varying the number of bins in zenith. These trials are shown in Fig. 9.3. The likelihood errors were smaller monotonically with more bins. The number of bins was thus set to the maximum number, 6, allowed without making bins smaller than the zenith angle resolution.

9.3.2 Energy

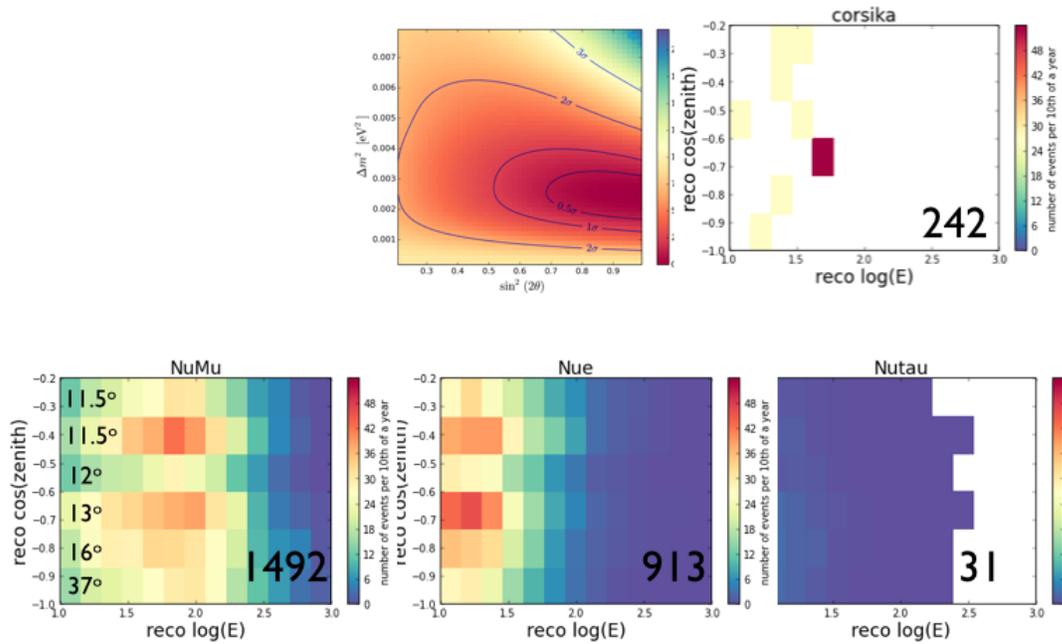
The test varying the energy cut uses logic similar to the test varying the zenith angle cut and binning. The energy resolution measured in simulation was 35%, or 0.15 in $\log_{10}E$. Based on that, the bin size in energy was 0.15. The low energy cutoff was already set to be at least 10 GeV because of the simulation generation cutoff and experimental/simulation disagreement below 10 GeV. The high energy cutoff was already set to be 1 TeV because the event selection emphasized low-energy events, so the contribution above 1 TeV is negligible. The parameter remaining to test is the option of pushing the low-energy cut higher. Three cases were tested, each one taking off another bin in energy: cutting at 1.0, cutting at 1.15, and cutting at 1.30. The results are shown in Fig 9.5. From this test, it is concluded that (similar to cutting in zenith) a harsher cut diminishes the statistics of the sample faster than it improves the quality, so leaving off the cut is the best course.

9.4 Variations in Cosmic Ray μ Rates

Cosmic ray muons are the major background for atmospheric neutrino studies in IceCube, so understanding the effect of muon contamination to our sample is particularly important. There are two ways used here to examine contamination using round trip tests, each of which is described and performed in this section. In any likelihood comparison, a hypothesis is generated by the fitter software, and compared against the data (in the case of round-trip tests, the “data” is simulated).

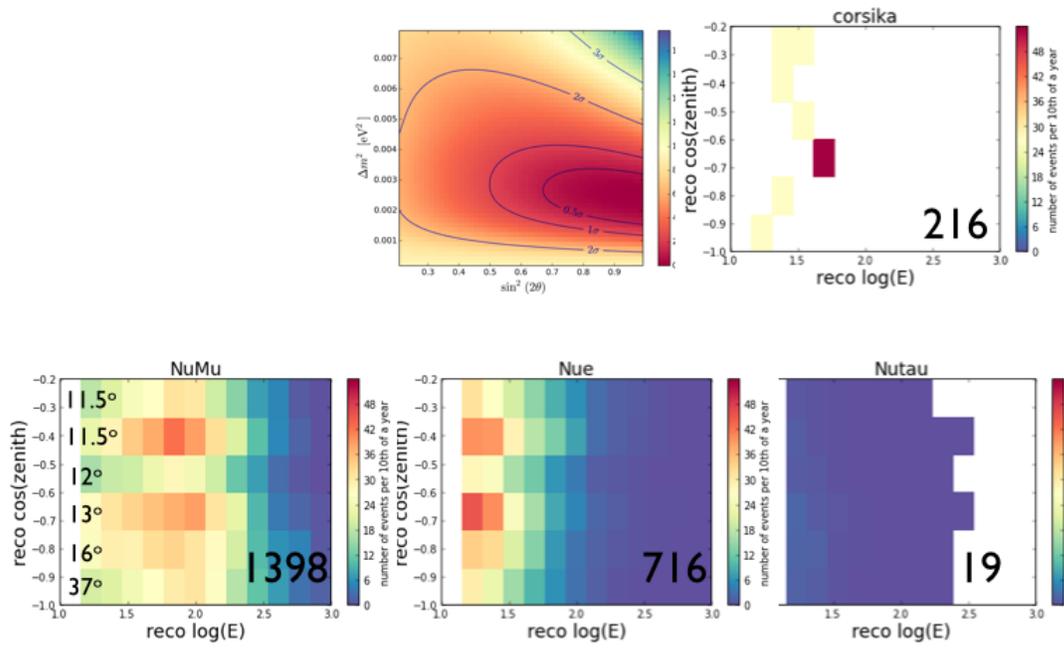
One way of performing round-trip contamination tests is to vary the cosmic ray muon background contribution to both the fitter hypothesis and the data; in this case, the fitter “knows” about the changes to the data, and data-simulation agreement is preserved. This test is useful for determining whether the level of contamination is acceptable for a good result. A set of tests using this method is shown in Fig. 9.7. The result of these tests is that cosmic ray muon rate variation has a minimal effect on the sensitivity of the analysis: the effect is only noticeable when the rate changes by an order of magnitude. From this we learn that cutting out all the muon background makes only a small difference to the sensitivity.

Another test focusses on how accurately we can describe the cosmic ray muon contamination. For this test, we use a different muon rate in the pseudodata than the fitter, and see how it affects the fit. The fitter varies each component of the fit, including muons, so it should be able to account for small discrepancies in the rate. But there are limits within

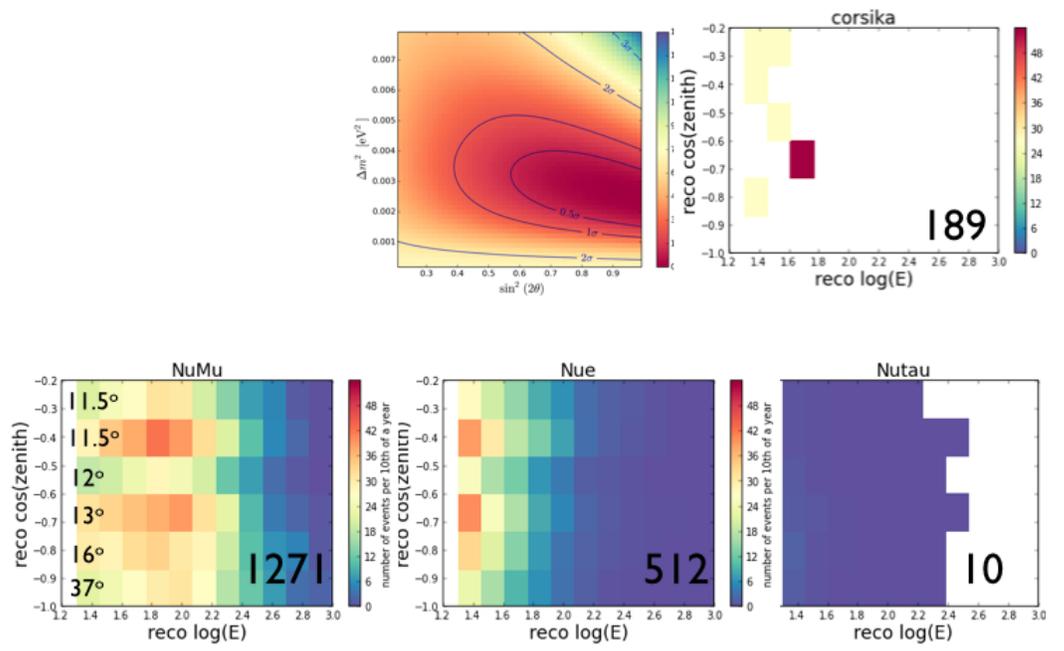


(a) Energy binning variation: 13 bins between 1 and 3 in $\log_{10}(E[\text{GeV}])$

Figure 9.5: Results of round-trip trials comparing different binning configuration for energy. The results from two additional configurations are shown in the following figure. For each configuration, the four components of the simulation rate ($\nu_\mu, \nu_e, \nu_\tau, \mu$) are shown at expected [21] oscillation values. Each component plot is labeled with the total number of events from that component. Within the ν_μ plot in each configuration, the bin sizes are labeled in degrees, as a reminder that the binsize is not constant in degrees (but is still constant in solid angle). A scan of the oscillation values is also shown, with Wilks' theorem error bars. The livetime assumed is 7500.71 hours.

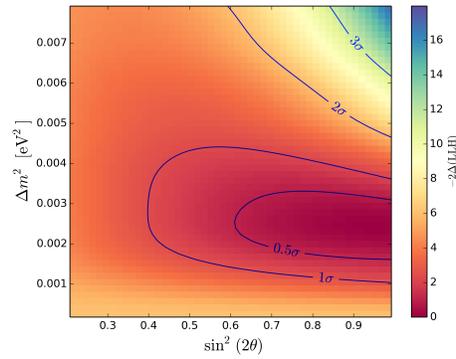


(a) Energy binning variation: 12 bins between 1.15 and 3 in $\log_{10}(E[GeV])$

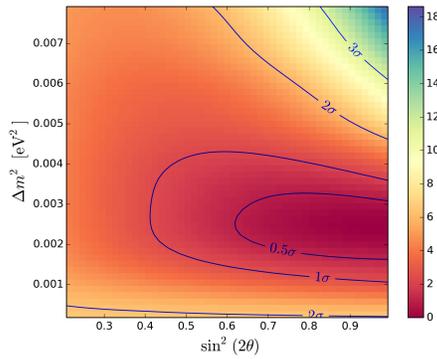


(b) Energy binning variation: 11 bins between 1.3 and 3 in $\log_{10}(E[GeV])$

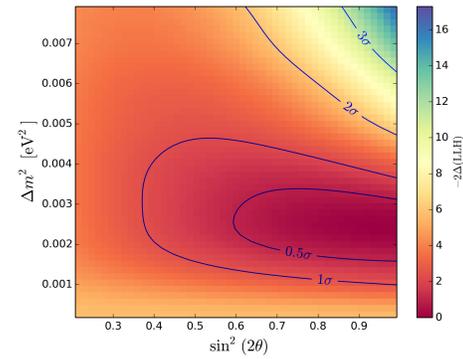
Figure 9.6: Results of round-trip trials comparing three different binning configuration for energy. See previous figure caption for more information.



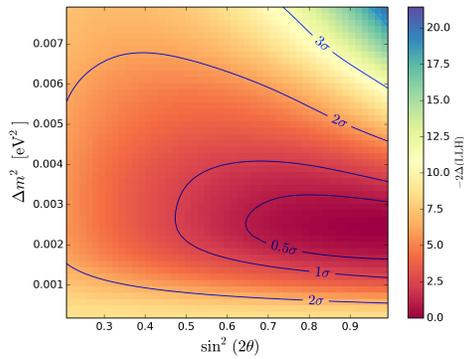
(a) Nominal simulated cosmic ray muon rate: no changes



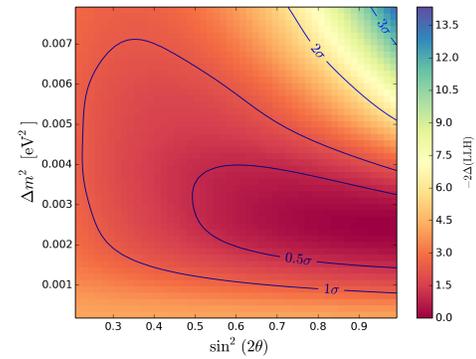
(b) Muon rate 0.5 times nominal



(c) Muon rate 2.0 times nominal

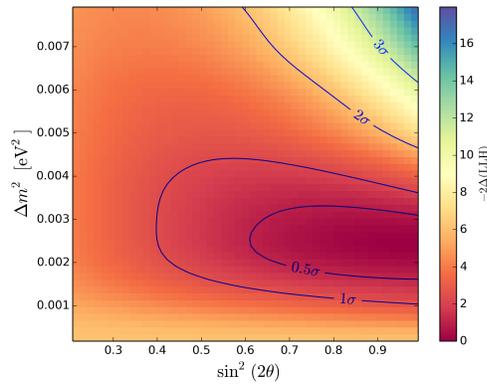


(d) Muon rate 0.1 times nominal

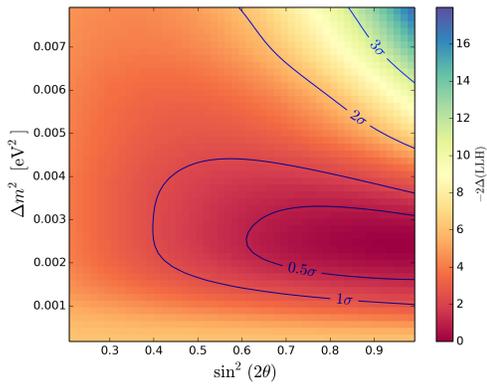


(e) Muon rate 10.0 times nominal

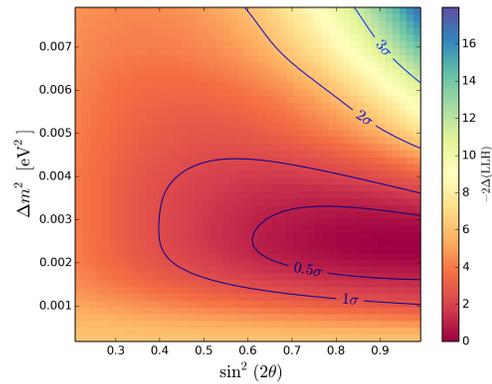
Figure 9.7: Round trip trials varying the rate of cosmic ray muon contamination. In all trials, the rate is varied both in the pseudodata and the expectation in the fitter, so agreement is preserved between “data” and simulation. Note that with this agreement, the impact of cosmic ray muon contamination is minimal. The rate for each sample is scaled to simulate the full IC79 detector year.



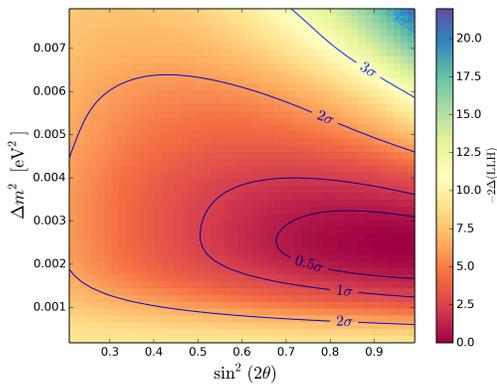
(a) Nominal simulated cosmic ray muon rate: no changes



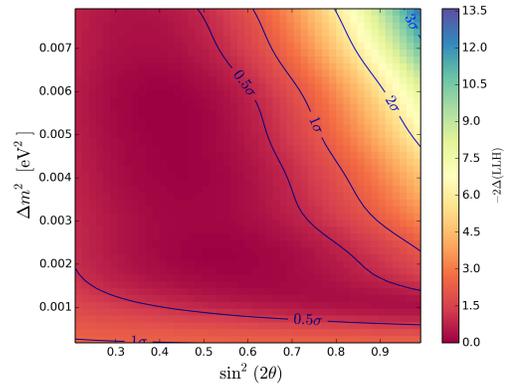
(b) Muon rate 2.0 times higher in fit than data



(c) Muon rate 2.0 times lower in fit than data



(d) Muon rate 10.0 times higher in fit than data



(e) Muon rate 10.0 times lower in fit than data

Figure 9.8: Round trip trials varying the rate of cosmic ray muon contamination. The data rate for all trials is kept constant while the fitter central value varies. The fitter is allowed to vary the cosmic ray muon normalization between 0.1 and 2.0 times its central value. The rate for each sample is scaled to simulate the full IC79 detector year.

the fitter: it may only change the muon background rate between 0.1 and 2.0 times the nominal value. If we set the pseudodata to something outside this range, it distorts the likelihood space.

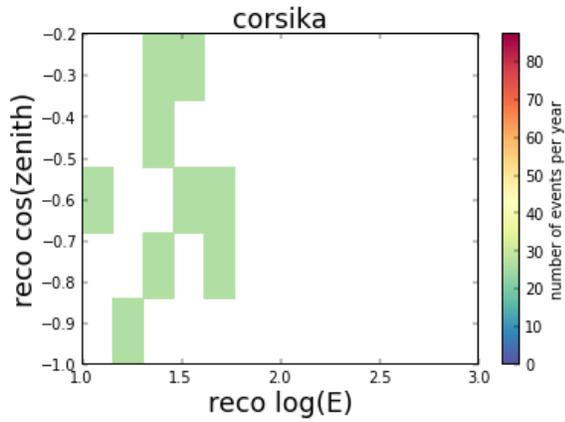
The results of a suite of tests in this style are shown in Fig. 9.8. For these tests, the cosmic ray muon rate in the pseudodata was kept constant, and the rate in the fitter expectation was varied¹. Note that in these results, distortion only arises when the rate difference is larger than the range allowed in the fitter.

A third test changes the shape of the cosmic ray muon background within its statistical errors. This is within the style of the second test (examining the effects of data-simulation). The reason for doing this test is that we know the CORSIKA sample is severely limited by statistics, so we should investigate how much that can hurt the sensitivity. Simulated background events are mostly (by design) cut out by the event selection. While avoiding muon contamination is a goal, it does limit the remaining sample statistically: from my selection, there are 9 CORSIKA simulated events remaining, which represent (with livetime weighting) 242 expected events in the final experimental sample. The CORSIKA files are generated with a natural spectrum, so each simulated event passing cuts gets the same weight. Considering \sqrt{N} errors on 9 events, the total rate can be considered to vary by ± 80 events, which is more than twice the entire ν_τ rate with standard oscillations.

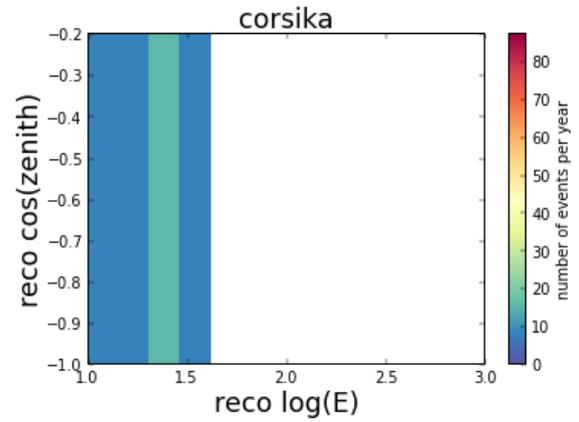
The test for changing cosmic ray muon background shape within statistical errors was performed in the following way. Two cosmic ray muon contribution histograms were constructed. The first had the 9 simulated events distributed as their reconstructions suggested (i.e., the way that all other component histograms were constructed). The second histogram was constructed by hand, by guessing a smooth distribution of events still consistent with the first histogram in the low-statistics case. The original histogram is shown in the upper-left panel of Fig. 9.9; the hand-smoothed histogram is shown in the upper-right panel of Fig. 9.9. The likelihood scan (lower panel of Fig. 9.9) uses the smooth map to construct the “data” and the bumpy map to do the fit; this simulates the current situation, where the fluctuations in the data are smaller than in the fitter.

This test shows that the shape of the cosmic ray muon background distribution has a huge effect on the sensitivity and accuracy of the analysis, but short of increasing the available simulation sample by at least a factor of 100 (which is not feasible on the timescale of this thesis), it does not suggest an easy way to fix this problem. This question will be raised again in Ch. 10. As a stop-gap measure, the smooth template is used for all further analysis, even though the shape was somewhat arbitrarily chosen.

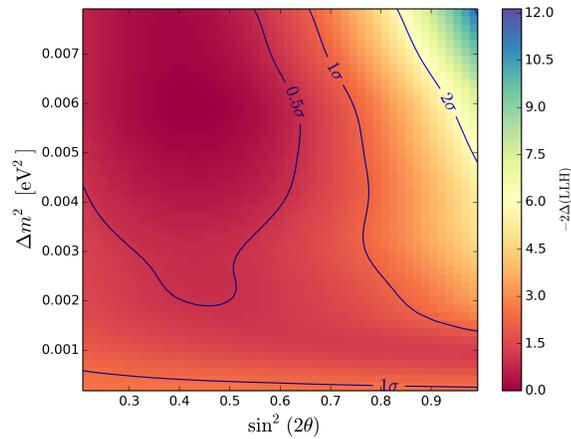
¹ More precisely, if someone uses my software to run this test again, make the following changes: add a factor x to the CORSIKA component in “DefineFitComponents.py” (which changes rates for all CORSIKA) then a factor of $1/x$ to cancel it out in the CORSIKA component of the pseudodata generation in “RunWilksRoundTrip.py”



(a) Distribution of the CORSIKA cosmic ray muon component predicted by 9 simulation events



(b) Distribution of the CORSIKA cosmic ray muon component smoothed by hand



(c) Result of using distribution (a) to try to fit (b)

Figure 9.9: Round trip trials comparing smoothed and unsmoothed CORSIKA distributions. The rate for each sample is scaled to simulate the full IC79 detector year.

9.5 Variations in ν_τ Rates

Contamination from ν_τ events is worth considering. Using methods similar to the cosmic ray rate variations, several trials were run varying the ν_τ rate (the test varying the shape of the cosmic ray variation was not repeated with ν_τ).

The results of varying the input data together with the fitter are shown in Fig. 9.10. These results (as with the cosmic ray muon results) show that in the case that experimental and simulation data agree, contamination is not a limiting factor compared to accurately describing the muon contamination, but lower contamination rates are still better.

The results of tests varying the fitter expectation separately from the input “data” are shown in Fig. 9.11. Again with tau as with cosmic ray muons, a disagreement between simulation and data has a bigger effect on the shape of the likelihood than changing the rate does. The effect is particularly noticeable in the case where the rate expected in the fitter is too high compared to the rate in the data.

9.6 Comparing Wilks Theorem and Feldman-Cousins Statistics

One of the major advances of this thesis is applying likelihood software to oscillations with enough speed optimizations that the full Feldman-Cousins statistical methods can be applied (see §8.4). The alternative method is to use Wilks theorem [132]. Both of these are methods for translating likelihood maps² into significances. To test the magnitude of the difference between these two cases, two round trip trials were run with the same simulation configuration but different statistical methods for determining the significance.

Both trials were set up in the usual way: the oscillation parameters used were the best-fit expected values [21]: maximal mixing at $\Delta m^2 = 0.00235$. Simulation data was used for both the input data and the fit variation; the same data was used in both cases, assuring a perfect match and limiting the scope of this test to only the differences in statistical significance method (i.e. this is not a “data challenge” test where every effort is made to create an independent simulation sample indistinguishable from experimental data). The simulation datasets used did not include the final Genie datasets with updated vuvuzela noise (see §4.3.2), because this test was run before those datasets were processed; the datasets used were instead those that the cuts were optimized for, which use the older noise-generator noise model.

The Wilks’ theorem results are shown in Fig. 9.12a. The results are as expected: the likelihood space is centered on the input oscillation parameter point, the contours are rounded with no unexplained local minima.

The results of the Feldman-Cousins trail are shown in Fig. 9.12b. The points included in this figure are the points that finished all their pseudo-data trials within the 12 hour cluster processing time limit; because of the structure of the significance software, higher significance points take longer to process. The color map in this figure shows p-value

²Recall: a likelihood map in this case is the result of a scan over the oscillation parameters, calculating the likelihood of each point, including a minimization of nuisance parameters at each point. In my Wilks theorem plots, likelihood is displayed as color, while significance is displayed as labeled contours.

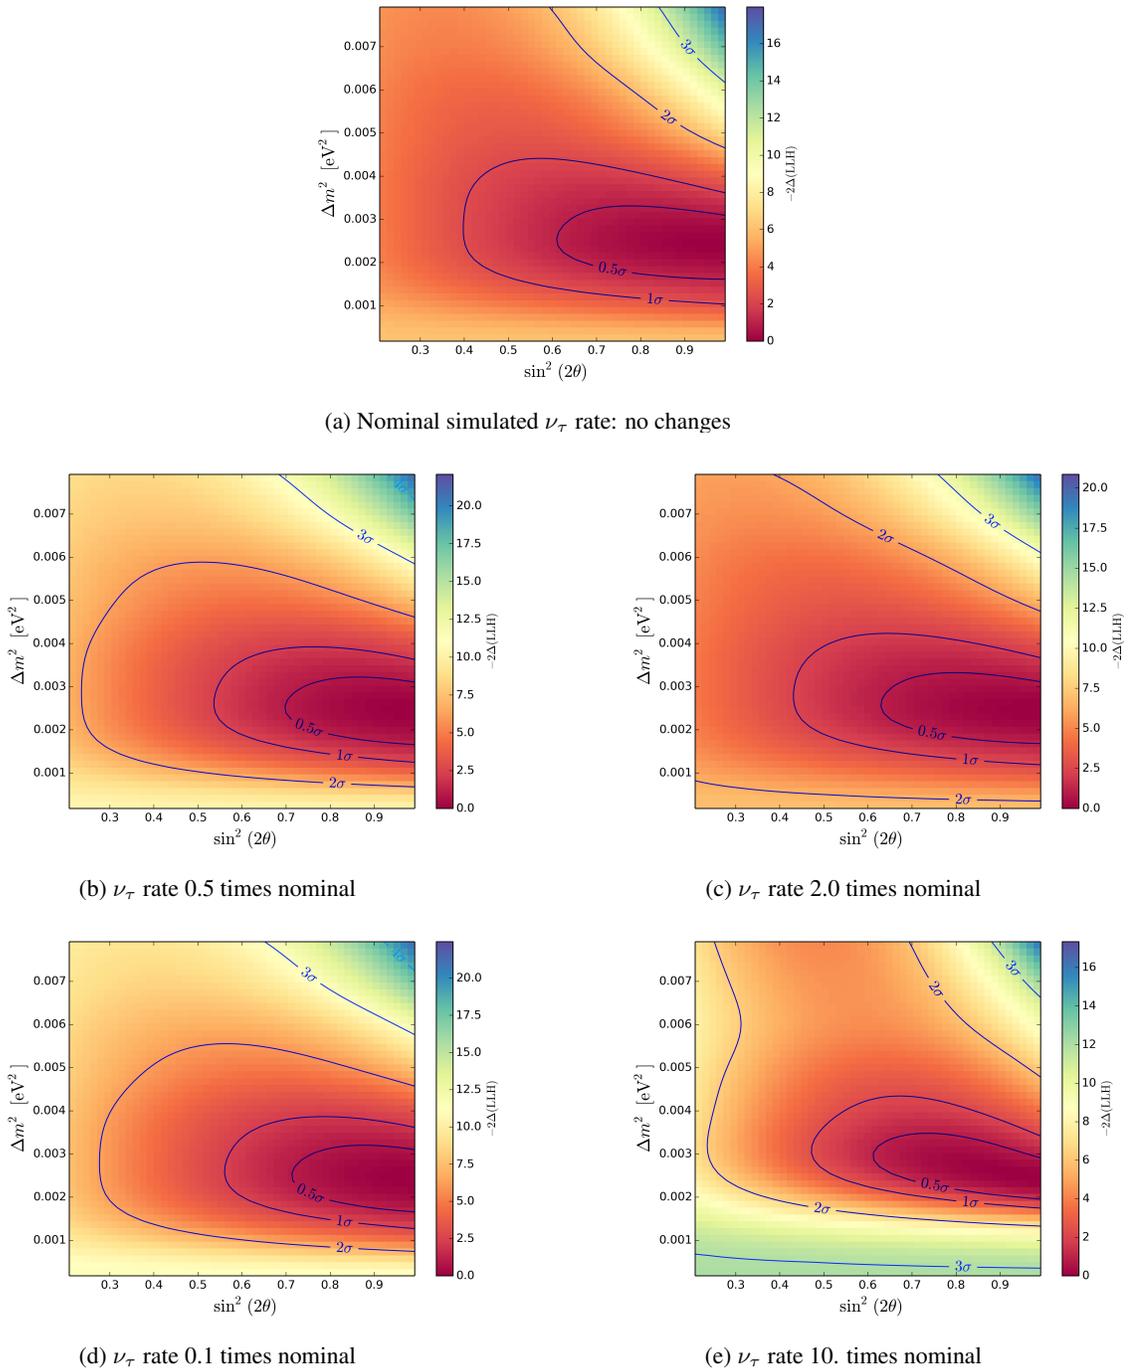


Figure 9.10: Round trip trials varying the rate of ν_τ contamination. In all trials, the rate is varied both in the pseudodata and the expectation in the fitter, so agreement is preserved between “data” and simulation. Lower contamination rates are better. The rate for each sample is scaled to simulate the full IC79 detector year.

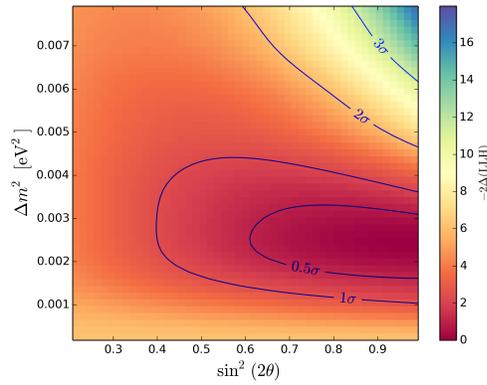
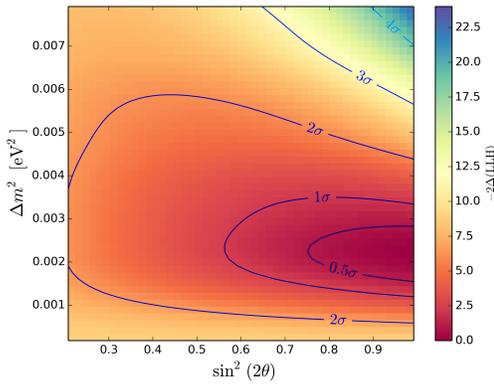
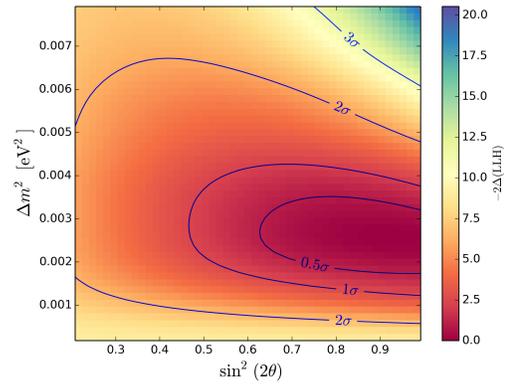
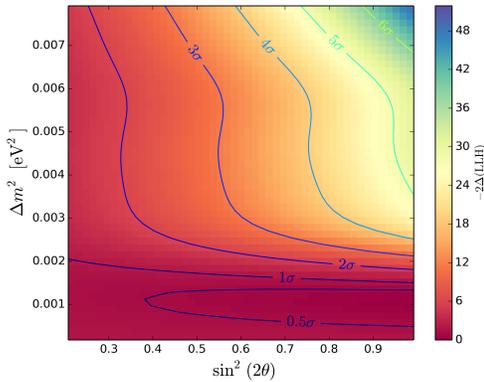
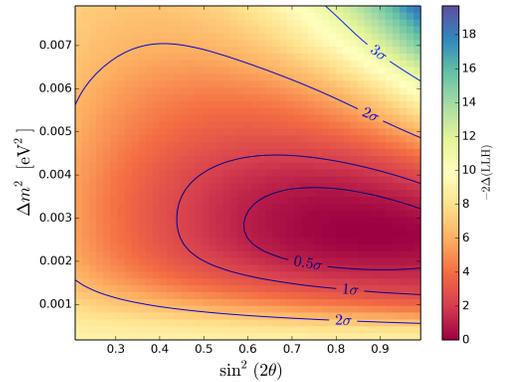
(a) Nominal simulated ν_τ rate: no changes(b) ν_τ rate 2.0 times higher in fit than data(c) ν_τ rate 2.0 times lower in fit than data(d) ν_τ rate 10.0 times higher in fit than data(e) ν_τ rate 10.0 times lower in fit than data

Figure 9.11: Round trip trials varying the rate of ν_τ contamination. The data rate for all trials is kept constant while the fitter central value varies. The fitter is allowed to vary the ν_τ normalization between 0.5 and 2.0 times its central value, but the ν_μ varies with the same factor, i.e. the rates are tied together. The rate for each sample is scaled to simulate the full IC79 detector year.

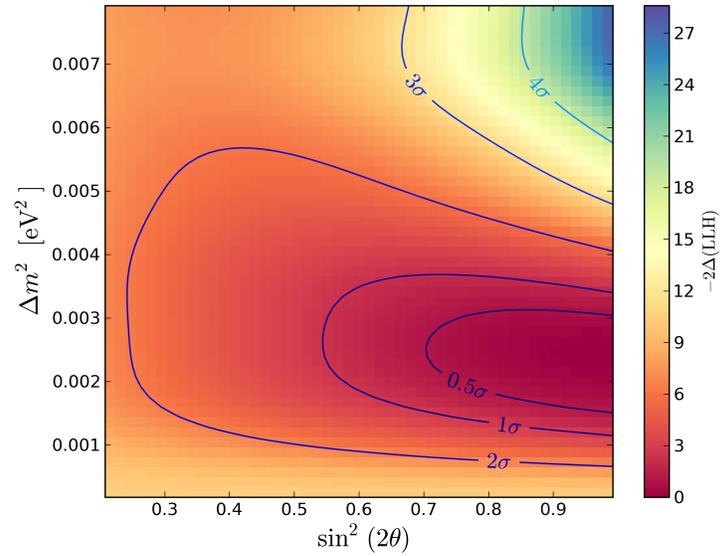
(instead of $\Delta\ln L$ -difference as the other figures do), and contours are only partly complete. As in the Wilks' figures, the contours included are 0.5σ , 1σ , and 2σ , although the 2σ contour is not complete.

Comparison between the two figures shows that while the basic shape is the same and centered on the same point, the contours fall in subtly different places. For example, the 1σ contour in the Wilks' case shows a Δm^2 range of $(1.2 - 3.4) \times 10^{-3}$ along the maximal mixing line, while the equivalent range in the Feldman-Cousins case is $(1.7 - 3.0) \times 10^{-3}$. This confirms the conventional wisdom that the two statistical methods give similar but not equivalent results.

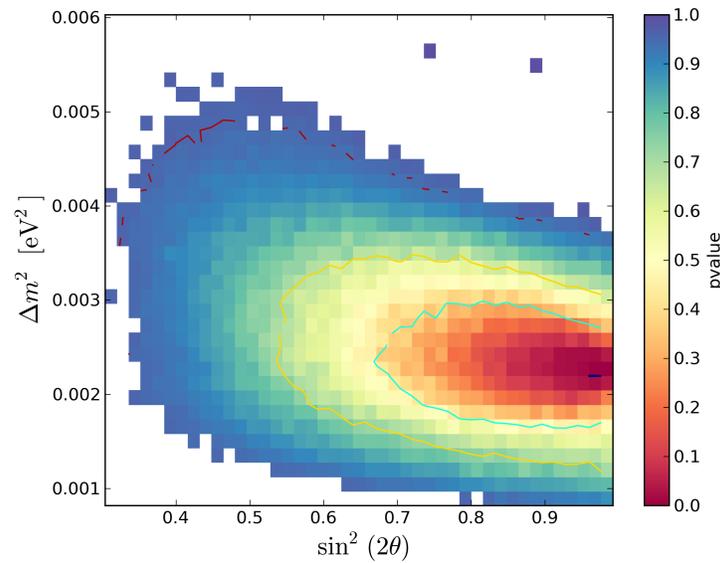
The original intention for this analysis was to use Wilks' theorem to establish a basic measurement on a small sample of the data, run tests to establish limits on the effects of various systematic effects, and estimate the difference between the expected significance of the smaller sample compared to the full sample. The intention was that Feldman-Cousins statistics would be used in the last step of the full central measurement, once a measurement was first established with Wilks' theorem contours. However, this step was deemed unnecessary in light of the large shape differences seen between the observed and expected results. The difference is large enough that it could not be created by the difference between Wilks' and Feldman-Cousins statistics.

The difference between these statistical methods remains important for any analysis reporting limits on oscillation parameters with the expectation of statistical formal correctness.

In conclusion, the round-trip trail is a powerful tool for exploring systematic effects and variations within a data sample. Round-trip trials were used to determine the placement of the final event selection cuts: the zenith angle and energy ranges for the 2-dimensional analysis histograms. Using this tool to explore various effects, including tau and cosmic-ray muon contamination and variations in statistical methods, we see a great variety of distortions to the likelihood space.



(a) Likelihood differences (in color) with Wilks' Theorem significance contours



(b) Significances (p-values, in color) from Feldman-Cousins statistics with contours for 0.5σ (green), 1.0σ (yellow), and 2.0σ (red). Each of the points in this plot had enough pseudo-data trials to produce 200 trials in the tail of the distribution.

Figure 9.12: Results of round-trip trials using the same setup to compare Wilks' theorem with Feldman-Cousins as methods for placing significance contours

Chapter 10

Results and Conclusions

This chapter provides an overview of the thesis as a whole, discussions of anomalies found, and several ideas for planning future work.

This thesis uses data from the IC79 configuration of the IceCube DeepCore detector to study standard oscillations in atmospheric neutrinos. Simulated data is used to design an event selection that isolates a sample of upgoing neutrinos in the energy range of 10-500 GeV, with tails on both ends. All events in this sample are reconstructed using likelihood-based reconstruction algorithms that integrate knowledge of the light scattering and absorption properties of the glacial ice. The energy and zenith arrival angle of each event is used to histogram the samples of both simulated and experimental data, and the histograms are then compared. The simulated data rates vary as a function of the assumed oscillation parameters. The normalization of each contribution (ν_μ , ν_e, ν_τ , and cosmic ray μ) to the total simulation is allowed to vary independently within a restricted range, usually 0.5 to 2.0 times the simulated central value, with the exception that the ν_μ and ν_τ rates vary by the same amount by construction. The likelihood of these fits is used to measure the likelihood of each tested value of the oscillation parameters. The likelihood software is tested in several ways, and these tests uncover limitations of the simulated data set, including limited background statistics.

One new contribution of this thesis is that it adds the new Monopod reconstruction algorithm (designed for higher energies) to a low-energy analysis. The Millipede reconstruction software suite has the first reconstructions with promise of working over the whole IceCube energy range. This work also adds fully correct statistical treatment of errors on the oscillation parameter contours using the Feldman-Cousins method. Using this reconstruction and statistical framework, simulation trials show that an accurate measurement of atmospheric oscillation parameters is possible. However, when this analysis technique is applied to experimental data from one year of IceCube running, the resulting likelihood scan does not match the expected effect or any of the simulated systematics effects.

An unexpected result is found and discussed below: the best fit values of the oscillation parameters are not the expected [21] best fit values. In fact, the likelihood fits better to the no-oscillation case. Considering the overwhelming body of evidence in favor of atmospheric neutrino oscillations, this result is considered to be an anomaly based on an as-yet not understood systematic error rather than a contradiction of oscillations. Thus, this chapter has a significant focus on ideas for further developing the currently limiting systematic errors, especially the estimated rate of the cosmic ray muon background.

10.1 Single Dimension Results

The first view of the results is shown here in the single-dimensional plots. Energy versus rate is shown in Fig. 10.1; zenith versus rate is shown in Fig. 10.2.

Consider first the energy distribution, in Fig. 10.1. The contribution from simulated cosmic ray muons is shown in green. The statistical error bars on this sample are remarkably large, which will be discussed further in later sections. As a reminder, the key oscillation region is 10-50 GeV ($1-1.7 \log_{10}(\text{GeV})$), using the units of the plot). The total rate predicted from simulation (assuming the expected [21] oscillation parameters) is shown as the red band in the upper panel. The experimental data agree well with this predicted rate. The lower panel shows the ratio to the no-oscillation simulated expected rate. As expected, the oscillation case simulation rate (in red) is lower, and the data follow this case in the key oscillation energy region. At higher energies, the total rates are lower and thus the statistical errors are higher; in this region, the increasing data/simulation disagreement is attributed to statistical fluctuation.

The plot of rate versus zenith angle (Fig. 10.2) uses the same color scheme: green denotes simulated cosmic ray muon contamination, and the red band shows the total predicted event rate. Once again, the simulated cosmic ray muons are distributed in a way that makes them difficult to cut out with simple cuts. The strongest difference between the predictions for oscillations and no-oscillations is seen on the left side of the plot, which shows events traveling directly upwards through the Earth. Likewise, no difference due to oscillations is expected at the right side of the plot, which shows events arriving from the angular region of the horizon. The data match the prediction at the horizon and most of the upgoing region. There is an unexpected discrepancy in the most upgoing bin (the left-most bin): the data here match the no-oscillations prediction nearly perfectly, while still falling within the statistical errors of the oscillation prediction. This demonstrates again the statistical limitations of the simulation.

10.2 Likelihood Scan Results with Variations in the Cosmic Ray Muon Rates

The anticipated result of this thesis is the likelihood scan of the oscillation parameter space, focusing on its best fit point and error range, which is presented in this section. All likelihood plots include significance contours based on Wilks' theorem [132]. The neutrino samples using the Genie Monte Carlo generator included the newest "vuvuzela" noise model; the other simulation samples were made using the older "noise-generator" sample.

The results of a scan over oscillation parameters vary greatly with the input simulation prediction, and specifically with the contribution of cosmic ray muons to the total simulated prediction. The cosmic ray muon contribution has the largest statistical error of all the components of the simulation. Three cases of cosmic ray muon distributions are considered, each of which lies within the statistical errors of the simulation prediction.

There are 9 events in the cosmic ray muon simulation sample passing analysis cuts. Their distribution in the reconstructed observables histogram is shown in the upper panel of Fig. 10.3. Each simulated cosmic ray muon event passing the analysis cuts represents 27 events predicted in the experimental sample. The statistical error on the rate

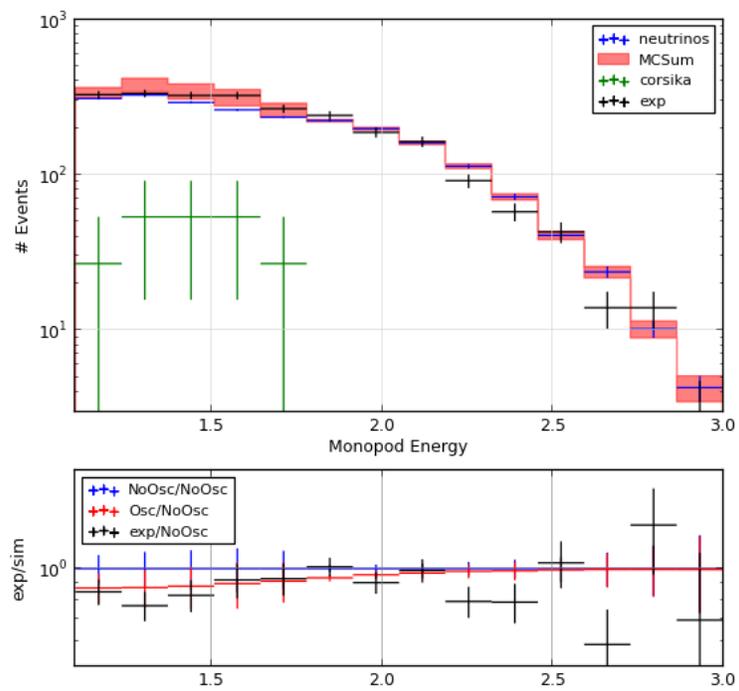


Figure 10.1: Rate versus reconstructed energy, comparing the full data to simulation predictions in both the expected oscillation and no oscillations cases.

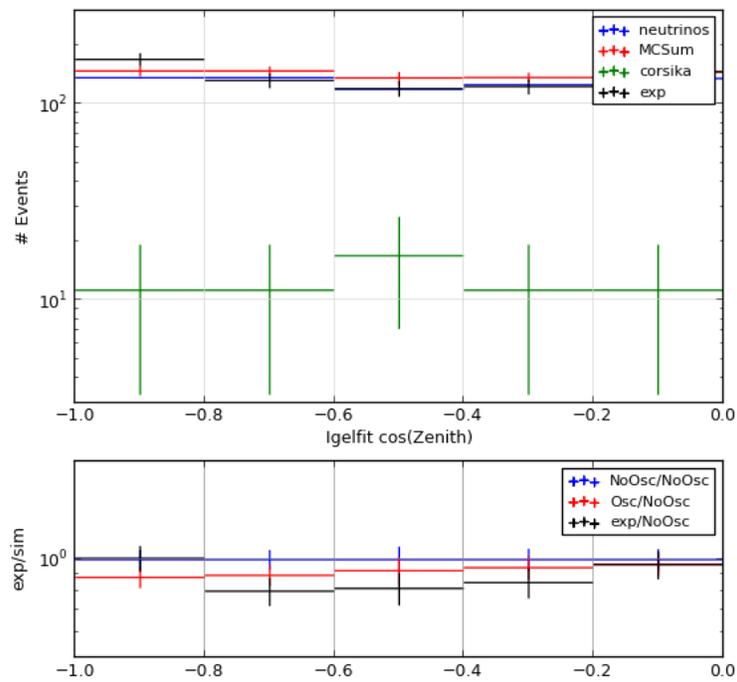
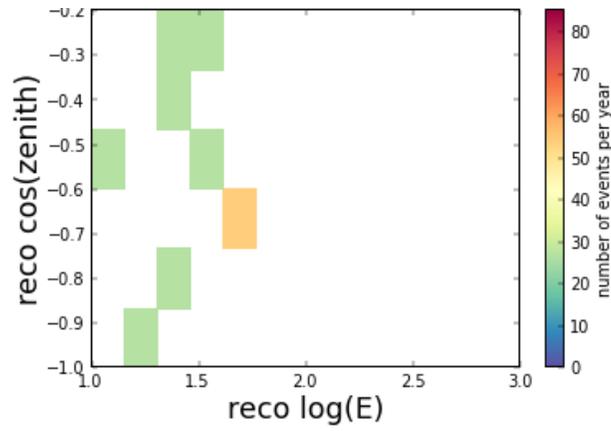
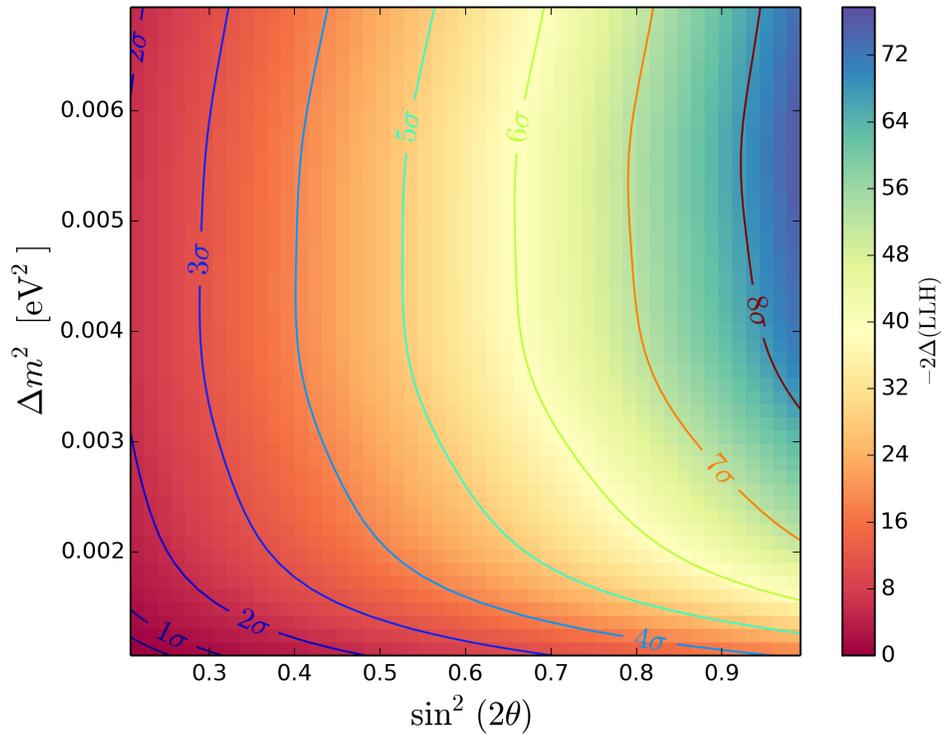


Figure 10.2: Rate versus reconstructed zenith angle, comparing the full data to simulation predictions in both the expected oscillation and no oscillations cases.

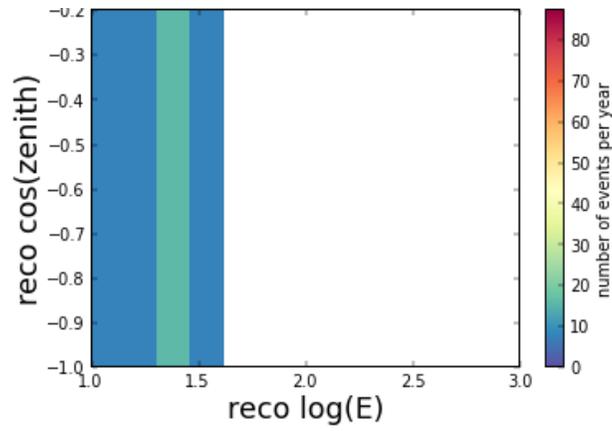


(a) Distribution of cosmic ray muon background assumed for the simulation estimated rate

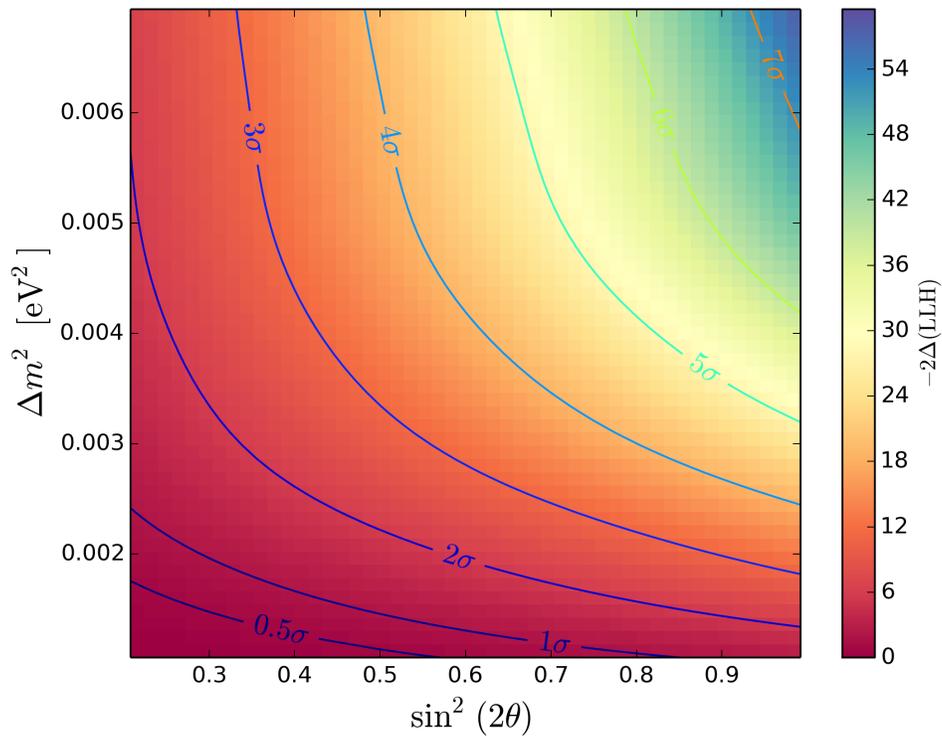


(b) Resulting scan over the oscillation parameter space, with significance contours assuming Wilks' theorem.

Figure 10.3: Result of a likelihood scan over the full IC79 year of data using a simulated expectation rate in which cosmic ray muon background rates are taken directly from simulation.

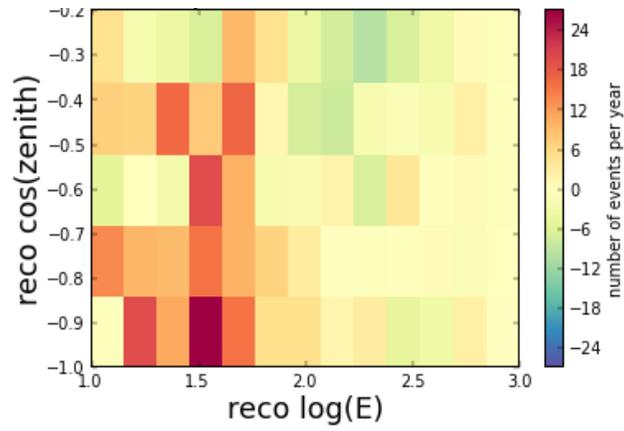


(a) Distribution of cosmic ray muon background assumed for the simulation estimated rate

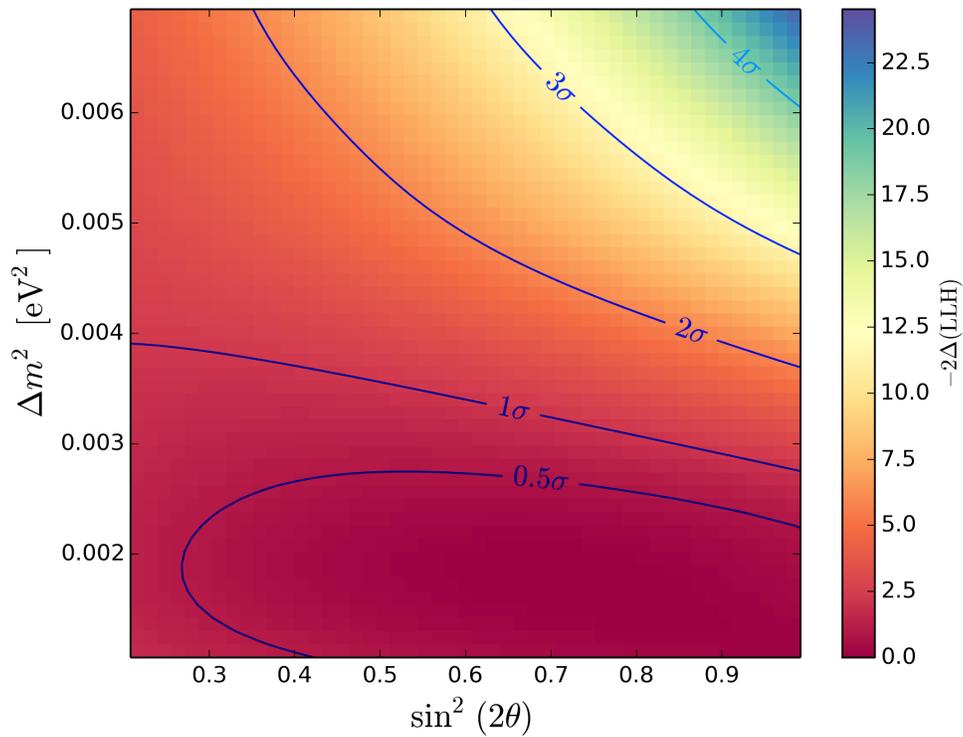


(b) Resulting scan over the oscillation parameter space, with significance contours assuming Wilks' theorem.

Figure 10.4: Result of a likelihood scan over the full IC79 year of data using a simulated expectation rate in which cosmic ray muon background rates are smoothed artificially.



(a) Distribution of cosmic ray muon background assumed for the simulation estimated rate



(b) Resulting scan over the oscillation parameter space, with significance contours assuming Wilks' theorem.

Figure 10.5: Result of a likelihood scan over the full IC79 year of data using a simulated expectation rate in which cosmic ray muon background rates are taken from the experimental data minus the simulated neutrino rate.

prediction of each bin of the observables histogram is thus ~ 1 simulated event, or ± 27 events in the total predicted rate.

The first case to consider is the cosmic ray muon prediction taken directly from the simulation, as distributed in the upper panel of Fig. 10.3. Adding to this the predicted neutrino rates and using analysis method described earlier, we compare simulation to experimental data, which results in the likelihood map shown in the lower panel of Fig. 10.3. The significance contours shown in this figure use the Wilks' theorem assumption. This result is unexpected in the sense that it excludes the expected [21] oscillation parameter point with more than 7σ statistical errors, and instead favors the no-oscillation case as the best fit. It is worth pointing out here that the analysis method does not incorporate the statistical errors of the simulated sample, only the inherent statistical errors from the predicted rate.

The second case to consider is a hand-smoothed histogram of cosmic ray muons. This histogram, shown in the upper panel of Fig. 10.4, was constructed (see §9.4) to examine the possible effects of comparing data to simulation with a different shape of the cosmic ray muon contribution. The rates for each bin were chosen with the following points in mind: the total rate of the observables histogram should be the same in the original and the smoothed histogram, and the smoothed histogram should be a shape that could plausibly generate the original histogram in a low-statistics case. The results of using this histogram for the cosmic ray muon contribution in a likelihood scan are shown in the lower panel of Fig. 10.4. Comparing this likelihood scan to that in Fig. 10.3, one can see that the smoothing represents a move towards the expected oscillation result; the expected best-fit point with these assumptions is excluded by only 4σ .

The third case to consider tests the logical limit of this method. For this case, we construct a cosmic ray muon histogram in a somewhat circular fashion by subtracting the statistically well-constrained neutrino predicted rate from the total experimental rate. The remaining events may be assumed to be cosmic ray muons. This is circular in the sense that we assume the expected values of the oscillation parameters and perfect agreement between the experimental and simulated data apart from the simulated cosmic ray muon rate. The rates for cosmic ray muons found in this way are plotted in the upper panel of Fig. 10.5. All fall within the statistical error range of the original simulated cosmic ray muon rates, i.e. the rates are all within 27 ± 27 for the lower energies and close to 0 for the higher energies. The likelihood scan of simulation to experimental data when assuming this distribution of the cosmic ray muons is shown in the lower panel of Fig. 10.5. This likelihood shape is the closest to the expected shape: it excludes the expected best-fit point by less than 1σ and favors oscillations over no oscillations.

Comparing these three cases shows extreme variation in the final result, which demonstrates that the accuracy of the cosmic ray muon shape dominates the accuracy of the result.

10.3 Specific Ideas for Future Improvements

This analysis is necessarily limited by the time constraints of the PhD process. While developing the analysis for a thesis, certain paths became clear for future improvements. These could make interesting topics for a new student's

summer project, or a service project for a beginning oscillation graduate student. Each project has the potential to improve the oscillation analysis.

10.3.1 Additional Cosmic Ray Muon Simulation

The most limiting factor of this analysis is the limited sample of cosmic ray muon background. There were 9 events from the muon background simulation surviving the event selection; the statistical errors on these 9 events were the largest single known error in the analysis. A straightforward (but difficult) way to improve this is to simulate more background. Another possibility is to allow more background to pass the event selection, reducing the statistical errors while increasing the errors from contamination. The contamination tests of §9.4 suggest that this may have been a wise approach. Yet another approach would be to include the statistical errors of each simulation component more explicitly in the analysis likelihood calculation. This may be the most feasible option, but would require a nearly complete rewrite of the analysis software. A middle ground could be to include the errors for the cosmic ray muon sample only.

The sample of simulated muons used for this analysis represents 11 days of simulated livetime compared to approximately one year of experimental data. Future results will need to take into account this dearth of simulated background relative to experimental signal, and several methods may be considered.

One method for further constraining the predicted cosmic ray muon rate is to simulate more data. This method is the most straightforward to configure, but would also require prohibitive amounts of computing time.

The additional simulation can be achieved either by running the standard CORSIKA simulation or by running a more targeted single-muon simulation. Running CORSIKA and associated detector simulations requires an approximate factor of 5 to 1 in time: 5 days of real time are required to simulate one day of simulated background with CORSIKA, assuming that all or most of the computing power of the Madison IceCube group is available for the task (or an equivalent amount of distributed computing power). This takes a prohibitively long amount of time to reach the required factor of 10:1 (or at least 1:1) simulated background to experimental data.

The more targeted method for generating cosmic ray muons uses the “muon gun” software, which includes more configurable detector volume and primary energies. By making three assumptions, the background simulation production speed can be greatly increased. Those three assumptions are that the background is only single muons and not muon bundles, that all the background muons passing analysis cuts transverse the DeepCore central volume, and that the muons passing analysis cuts deposit less than 1 TeV of visible energy within IceCube. With these assumptions, generating a larger sample of cosmic ray muon background would take about 5 GPU-years. Considering current collaboration GPU resources, a single student could produce this sample within a one- to two-year timeframe. As that timeline still falls outside the scope of this thesis, both of these methods are left as future work.

10.3.2 Spline Tables

The reconstruction algorithms described in Ch.7 rely on descriptions of how light travels in the ice, called “photonics tables,” or after a fitting process is applied, “spline tables”. These tables are described in §5.8. While the tables are adequate, improvements could be made which would lead immediately to improvements in the reconstruction resolution.

One potential improvement would be to improve the spline fit at small distances from the DOMs. This would be done by scanning through fits by hand and adding knots as needed. This is currently being addressed by Moriah Tobin.

Another potential improvement is to change the simulation that generates the photonics data. The simulation currently assumes that neutrino interactions in the ice produce only three output particles: a charged lepton, a neutral lepton (neutrino), and a single particle encompassing the average of everything else, which we call “hadrons.” The photonics tables could be made more accurate by including more accurate descriptions of the many particles involved in low-energy showers. An easy way to do this is to use the Geant simulation software [137], which already includes data about such interactions. Currently, Sarah Nowicki is working on these tables. When they become available and are used for reconstructions, the reconstruction resolutions will improve.

The spline tables were found to be a limiting factor during early development of the Biped reconstruction (a millipede-likelihood based fit using the hypothesis of two particles: a cascade and a track). Early development showed no improvement when a more accurate hypothesis was used. Rather than abandon the model, a problem was found and solved at a lower level: the spline tables. These descriptions of how light travels in the ice were found to be inaccurate at the short distances that were most relevant to low-energy fits. Fixing the tables greatly increased the accuracy of Biped and any reconstruction using the new tables, although it required the investment of a full student-year.

10.3.3 Noise Modeling

Measurements of noise rates and probability distributions exist and could be better integrated into reconstructions and simulations. A great deal of improvement has already been made towards a better noise model, and while it is not the dominant error in this analysis, it still cannot be ruled out as a major systematic effect.

It is well known that the dark noise in IceCube DOMs does not strictly follow a Poisson distribution. The distribution that it does follow has been measured and described [83], but including this noise model in the various aspects of simulation and reconstruction is a continuing process. At the moment of this work being finalized, the new noise simulation project vuvuzela (see §4.3.2) has been included as a standard part of new simulation datasets. The additional step of including noise-only events with each simulated data set has not been completed. As an intermediate step, simulated data is available with noise-only events separately from physics events. Consequently, it should be possible to simulate a fully correct spectrum by including all types of relevant simulation, where noise events are an additional entry on that list. This spectrum still needs calibration at the lowest-energy end.

A notable difference between millipede and some other reconstructions currently under consideration is that millipede includes noise hits by design in the likelihood. A straw-man alternative method is to cut out the noise hits before giving the pulse information to the fitter. This method runs into the insurmountable problem that some noise hits may occur close to physics hits, and it is impossible to separate them reliably and accurately. The better alternative that millipede uses is to give some noise probability to every observed hit.

Adding new noise models to reconstruction algorithms is a more involved process. Noise in reconstruction is often included as a constant term (as a function of the time length of the event) in the likelihood. Correcting the shape of the noise distribution in reconstruction is equivalent to changing that addition to likelihood function. While this is a non-trivial change, it is possible to first make an incremental improvement. One can use the measurements taken during the creation of the vuvuzela noise model to update the values of the constant noise term. While this is not the most accurate possible model for the noise, it does make the noise assumptions in reconstruction as correct as possible within the constraints of a Poisson model. This correction is $\sim 10\%$ in the noise rate.

10.3.4 Updating the Genie Simulation Generator

The simulations used in this analysis were generated with two different Monte Carlo simulation generators: Genie and NuGen. Each one has advantages. IceCube would benefit from combining the positive qualities of each into a single generator, or else more explicitly separating them. For example, NuGen includes the effects of propagation through the Earth but makes bad assumptions about the details of relatively low energy interactions in the ice. If these two functions were separated into different software steps, each analyzer (or simulation coordinator) could decide what assumptions to combine. This effort is currently underway, led by various groups working on sterile neutrino searches.

10.3.5 Particle Reconstructions

The choice of energy and zenith angle reconstructions has turned out to be one of the largest differences between different students' approaches within the IceCube oscillations working group, and as such this choice merits a higher level of scrutiny than the other cuts.

Monopod and Igelfit were chosen for several reasons. Generally, during the time of the development of this thesis analysis, the Millipede reconstruction library was the most well-developed and cross-checked set of reconstructions. Millipede was the first software to use all pulses directly in a reconstruction (without dropping after-first hits, or approximating N hits as 1 at a carefully chosen position). Millipede was the first reconstruction to successfully include a noise model within the likelihood function, obviating the need for error-prone hit cleaning beyond a time-window cut. Millipede reconstructions were shown to be accurate with higher energy events ($\mathcal{O}(\text{TeV})$ to $\mathcal{O}(\text{PeV})$), and the analysis of this thesis represented an opportunity to test them at lower IceCube energies. Initial tests of the energy

resolution, as well as the final measured energy resolution of 35%, support the idea of expanding Millipede-based reconstructions to lower energies.

While maintaining all of those positive reasons to try an analysis with Millipede-based reconstructions, there are also some reasons against it.

The first and easiest to dismiss is that any Millipede-based reconstruction is slow. However, this slowness can be circumvented particularly easily at low energies by setting up the reconstruction carefully. We can assume that the tracks are shorter than the full detector, and that light will not travel the full width of the detector, and thus there are many simple speed optimizations for running Millipede at low energies. These principles guide the several efforts to extend existing reconstructions to lower energy ranges.

A second way in which all Millipede reconstructions could be improved is by improving the accuracy of the light propagation tables for short distances and low energies. This should improve the quality of all reconstruction algorithms, but it will especially benefit Millipede-based algorithms because they contain so little other than the light propagation likelihood calculation.

A final known shortcoming of Millipede and Monopod in their current state concerns the noise model, and this argument has not been resolved by the work of this thesis. While it is an excellent advance in reconstruction algorithms that we now include the noise within the reconstruction likelihood, we also know that the accuracy of that noise model is critical for any task concerning low energy events. The noise model within Millipede is currently a constant probability, that is, any hit has a small probability of originating from noise. Equivalently stated, the noise model is inherently Poissonian within Millipede even though we know the real detector noise has a substantial non-Poissonian component. Much of the noise at low energies is dominated by correlated bursts that cannot be described by a single rate. As such, it would be difficult to integrate more accurate noise modeling into Millipede. It is still possible to improve the average noise rate, which is about 10% low in the old (and Millipede) model. While this should have some effect, it does not fully account for the inaccuracies observed.

There are currently many efforts within the collaboration to improve event reconstruction. The most promising of these incorporate the existing gulliver software framework, which facilitates debugging efforts and requires documented and modularized code. The most promising likelihood function is the millipede likelihood, which includes information from all pulses on all DOMs and considers no-hit DOMs in the likelihood.

10.4 Conclusion

In conclusion, this analysis has demonstrated the feasibility of measuring atmospheric neutrino oscillations with IceCube, but has not yet measured oscillations. Some ideas are presented for further analysis, focusing on the cosmic ray muon background simulation.

LIST OF REFERENCES

- [1] Sir Arthur Conan Doyle. The sign of the four, 1890. Page 111 in *The Penguin Complete Sherlock Holmes* (1981) ISBN 0713914440.
- [2] Katrin spectrum. Zykure - Own work. Licensed under Creative Commons Attribution-Share Alike 3.0-2.5-2.0-1.0 via Wikimedia Commons.
- [3] Wolfgang Pauli. Open letter to radioactive ladies and gentlemen (1930). *Physics Today*, 31:27, 1978.
- [4] Ruth Lewin Sime. *Lise Meitner: A Life in Physics*. University of California Press, 1996.
- [5] Enrico Fermi. Tentativo di una teoria dei raggi β . *Il Nuovo Cimento*, 11(1):1–19, 1934.
- [6] E. Fermi. Versuch einer theorie der β -strahlen. i. *Zeitschrift für Physik*, 88(3-4):161–177, 1934.
- [7] J. Chadwick. Possible Existence of a Neutron. *Nature*, 129:312, 1932.
- [8] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [9] Press release: The 1995 nobel prize in physics. Nobelprize.org, retrieved 14 June 2013.
- [10] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys.Rev.Lett.*, 20:1205–1209, 1968.
- [11] P. Anselmann et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Phys.Lett.*, B285:376–389, 1992.
- [12] Dzh.N. Abdurashitov, V.N. Gavrin, S.V. Girin, V.V. Gorbachev, Tatiana V. Ibragimova, et al. The Russian-American gallium experiment (SAGE) Cr neutrino source measurement. *Phys.Rev.Lett.*, 77:4708–4711, 1996.
- [13] V.N. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Phys.Lett.*, B28:493, 1969.
- [14] K.G. Libbrecht and Woodard M.F. Advances in helioseismology. *Science*, 253:152, 1992.
- [15] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys.Rev.Lett.*, 81:1562–1567, 1998.
- [16] Q.R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory. *Phys.Rev.Lett.*, 87:071301, 2001.
- [17] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti-neutrino disappearance. *Phys.Rev.Lett.*, 90:021802, 2003.
- [18] T. Araki et al. Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion. *Phys.Rev.Lett.*, 94:081801, 2005.

- [19] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog.Theor.Phys.*, 28:870–880, 1962.
- [20] B. Pontecorvo. Mesonium and anti-mesonium. *Sov.Phys.JETP*, 6:429, 1957.
- [21] J. Beringer et al. Review of Particle Physics (RPP). *Phys.Rev.*, D86:010001, 2012. and 2013 partial update for the 2014 edition.
- [22] F. P. An, J. Z. Bai, A. B. Balantekin, H. R. Band, D. Beavis, W. Beriguete, M. Bishai, S. Blyth, K. Boddy, R. L. Brown, and et al. Observation of Electron-Antineutrino Disappearance at Daya Bay. *Physical Review Letters*, 108(17):171803, April 2012.
- [23] Y. Abe et al. Indication for the disappearance of reactor electron antineutrinos in the Double Chooz experiment. *Phys.Rev.Lett.*, 108:131801, 2012.
- [24] Darrel Smith. Calculating the probability for neutrino oscillations. Student lecture series for miniboone, Embry-Riddle University, July 18 2001. Retrieved August 10, 2013 <http://physicsx.pr.erau.edu/Office/oscillations.pdf>.
- [25] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.
- [26] S.P. Mikheev and A. Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov.J.Nucl.Phys.*, 42:913–917, 1985.
- [27] Vernon D. Barger, K. Whisnant, S. Pakvasa, and R.J.N. Phillips. Matter Effects on Three-Neutrino Oscillations. *Phys.Rev.*, D22:2718, 1980.
- [28] M.C. Gonzalez-Garcia and Carlos Pena-Garay. Status of the MSW solutions to the solar neutrino problem. *Nucl.Phys.Proc.Suppl.*, 87:204–207, 2000.
- [29] M.G. Aartsen et al. Letter of Intent: The Precision IceCube Next Generation Upgrade (PINGU). 2014.
- [30] John Updike. Cosmic gall. *The New Yorker*. December 17, 1960, page 36.
- [31] A.A. Aguilar-Arevalo et al. A Search for electron neutrino appearance at the $\Delta m^2 \sim 1\text{eV}^2$ scale. *Phys.Rev.Lett.*, 98:231801, 2007.
- [32] J.A. Formaggio and G.P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev.Mod.Phys.*, 84:1307, 2012.
- [33] Sheldon L. Glashow. Resonant scattering of antineutrinos. *Phys. Rev.*, 118:316–317, Apr 1960.
- [34] A. de Gouvêa, K. Pitts, K. Scholberg, G.P. Zeller, et al. Working Group Report: Neutrinos. 2013.
- [35] Mike Perricone. Hitting the broad side of a (classified) barn. *Symmetry Magazine*, 3:4, February 2006. Signal to Background.
- [36] M.G. Aartsen et al. Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector. *Science*, 342(6161):1242856, 2013.
- [37] Francis Halzen and Alan Martin. *Quarks & Leptons: An introductory course in modern particle physics*. John Wiley & Sons, New York, USA, 1984.

- [38] M.G. Aartsen et al. Searches for Extended and Point-like Neutrino Sources with Four Years of IceCube Data. 2014.
- [39] Victor f. hess - biographical. In *Nobel Lectures, Physics 1922-1941*. Elsevier Publishing Company, Amsterdam, 1965.
- [40] M.G. Aartsen et al. Observation of Cosmic Ray Anisotropy with the IceTop Air Shower Array. *Astrophys.J.*, 765:55, 2013.
- [41] T.K. Gaisser and M. Honda. Flux of atmospheric neutrinos. *Ann.Rev.Nucl.Part.Sci.*, 52:153–199, 2002.
- [42] R. Abbasi et al. Cosmic Ray Composition and Energy Spectrum from 1-30 PeV Using the 40-String Configuration of IceTop and IceCube. *Astropart.Phys.*, 42:15–32, 2013.
- [43] James J. Beatty and Stefan Westerhoff. The Highest-Energy Cosmic Rays. *Ann.Rev.Nucl.Part.Sci.*, 59:319–345, 2009.
- [44] Y.S. Yoon, H.S. Ahn, P.S. Allison, M.G. Bagliesi, J.J. Beatty, et al. Cosmic-Ray Proton and Helium Spectra from the First CREAM Flight. *Astrophys.J.*, 728:122, 2011.
- [45] T.K. Gaisser. Cosmic rays and particle physics. 1990.
- [46] Sandip Pakvasa, Werner Rodejohann, and Thomas J. Weiler. Flavor Ratios of Astrophysical Neutrinos: Implications for Precision Measurements. *JHEP*, 0802:005, 2008.
- [47] Kurt Woschnagg, personal communication. Also available from the IceCube internal wiki page “Atmospheric neutrino fluxes”.
- [48] Paolo Lipari. Lepton spectra in the earth’s atmosphere. *Astroparticle Physics*, 1(2):195 – 227, 1993.
- [49] Morihiro Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki. Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data. *Phys.Rev.*, D75:043006, 2007.
- [50] G.D. Barr, T.K. Gaisser, P. Lipari, Simon Robbins, and T. Stanev. A Three - dimensional calculation of atmospheric neutrinos. *Phys.Rev.*, D70:023006, 2004.
- [51] A. Achterberg et al. First Year Performance of The IceCube Neutrino Telescope. *Astropart.Phys.*, 26:155–173, 2006.
- [52] M. G. Aartsen et al. Measurement of atmospheric neutrino oscillations with icecube. *Phys. Rev. Lett.*, 111:081801, Aug 2013.
- [53] Michael R. Dziomba. *A Study of Neutrino Oscillation Models with Super-Kamiokande Atmospheric Neutrino Data*. PhD thesis, University of Washington, 2012.
- [54] Roger Wendell. Latest super-k atmospheric neutrino oscillation analysis. Now2012, Conca Specchiulla.
- [55] K. Abe et al. Precise Measurement of the Neutrino Mixing Parameter θ_{23} from Muon Neutrino Disappearance in an Off-axis Beam. *Phys.Rev.Lett.*, 112:181801, 2014.
- [56] P. Adamson et al. Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys.Rev.Lett.*, 110(25):251801, 2013.
- [57] accessed 20 Sep 2014.

- [58] R. Abbasi et al. The Design and Performance of IceCube DeepCore. *Astropart.Phys.*, 35:615–624, 2012.
- [59] R. Abbasi et al. IceTop: The surface component of IceCube. *Nucl.Instrum.Meth.*, A700:188–220, 2013.
- [60] Stefan Schonert, Thomas K. Gaisser, Elisa Resconi, and Olaf Schulz. Vetoing atmospheric neutrinos in a high energy neutrino telescope. *Phys.Rev.*, D79:043009, 2009.
- [61] J. Cherwinka, R. Co, D.F. Cowen, D. Grant, F. Halzen, et al. A Search for the Dark Matter Annual Modulation in South Pole Ice. *Astropart.Phys.*, 35:749–754, 2012.
- [62] H.V. Klapdor-Kleingrothaus, A. Dietz, H.L. Harney, and I.V. Krivosheina. Evidence for neutrinoless double beta decay. *Mod.Phys.Lett.*, A16:2409–2420, 2001.
- [63] E. Andres, P. Askebjerg, X. Bai, G. Barouch, S.W. Barwick, et al. Observation of high-energy neutrinos using Cherenkov detectors embedded deep in Antarctic ice. *Nature*, 410:441–443, 2001.
- [64] J. Ahrens et al. Muon track reconstruction and data selection techniques in AMANDA. *Nucl.Instrum.Meth.*, A524:169–194, 2004.
- [65] P. Allison, J. Auffenberg, R. Bard, J.J. Beatty, D.Z. Besson, et al. Design and Initial Performance of the Askaryan Radio Array Prototype EeV Neutrino Detector at the South Pole. *Astropart.Phys.*, 35:457–477, 2012.
- [66] United States Geological Survey. Geographic names information system feature detail report for: Mount mitchell. geonames.usgs.gov.
- [67] R. Abbasi et al. Calibration and Characterization of the IceCube Photomultiplier Tube. *Nucl.Instrum.Meth.*, A618:139–152, 2010.
- [68] Kristan Hutchison. Icecube - one hole done, 79 more to go. SpaceRef.com, press release from The Antarctic Sun, 24 October 2005.
- [69] Icecube filesystem, filtered data.
- [70] Oladipo Fadiran. wiki/goodrunlist. IceCube internal Good Run List, available upon request.
- [71] M.G. Aartsen et al. Measurement of South Pole ice transparency with the IceCube LED calibration system. *Nucl.Instrum.Meth.*, A711:73–89, 2013.
- [72] Dmitry Chirkin. Study of ice transparency with icecube flashers. Technical report, 2009.
- [73] Per Askebjerg, Steven W. Barwick, Lars Bergström, Adam Bouchta, Staffan Carius, et al. Optical properties of deep ice at the South Pole: Absorption. *Appl.Opt.*, 36:4168–4180, 1997.
- [74] P. Buford Price and Lars Bergström. Optical properties of deep ice at the South Pole: scattering. *Appl.Opt.*, 36:4181–4194, 1997.
- [75] R. C. Bay, R. A. Rohde, Price P. B., and N. E. Bramall. South Pole paleowind from automated synthesis of ice core records. *J.Geophys.Res.*, 115:D14126, 2010.
- [76] U. Ruth et al. Edml1: a chronology for the epica deep ice core from dronning maud land, antarctica, over the last 150 000 years. *Clim. Past.*, pages 475–484, 2007.
- [77] M. Ackermann et al. Optical properties of deep glacial ice at the south pole. *J. Geophys. Res.*, 111:D13203, 2006.

- [78] P. Liu. A new phase function approximating to Mie scattering for radiative transport equations. *Phys. Med. Biol.*, 39(6):1025–36, 1994.
- [79] R. Abbasi et al. The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping. *Nucl.Instrum.Meth.*, A601:294–316, 2009.
- [80] Synchrotron radiation center. The Synchrotron Radiation Center is primarily funded by the University of Wisconsin-Madison with supplemental support from facility Users and the University of Wisconsin-Milwaukee.
- [81] The LBNE Collaboration. The Long Baseline Neutrino Experiment (LBNE) Water Cherenkov Detector (WCD) Conceptual Design Report (CDR), 2012.
- [82] H.O. Meyer. Dark Rate of a Photomultiplier at Cryogenic Temperatures. 2008.
- [83] Michael James Larson. Simulation and identification of non-poissonian noise triggers in the icecube neutrino detector. Master’s thesis, University of Alabama, Tuscaloosa, Alabama, 2013.
- [84] Troy D. Straszheim. Icefishing for neutrinos (with boost). Boostcon 2009, Aspen Colorado.
- [85] M.G. Aartsen, R. Abbasi, M. Ackermann, J. Adams, J.A. Aguilar, et al. The IceProd Framework: Distributed Data Processing for the IceCube Neutrino Observatory. 2013.
- [86] How npulses hurts everyone everywhere. IceCube internal Wiki, 22 November 2011, retrieved 31 July 2014.
- [87] K. Hagiwara et al. Monte carlo particle numbering scheme. *Physical Review*, D66:010001–1, 2002. available on the PDG WWW pages (URL: <http://pdg.lbl.gov/>).
- [88] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, et al. The GENIE Neutrino Monte Carlo Generator. *Nucl.Instrum.Meth.*, A614:87–104, 2010.
- [89] M.G. Aartsen, R. Abbasi, Y. Abdou, M. Ackermann, J. Adams, et al. Improvement in Fast Particle Track Reconstruction with Robust Statistics. *Nucl.Instrum.Meth.*, A736:143–149, 2014.
- [90] *Measurement of the atmospheric neutrino energy spectrum with IceCube*, 2009. Proceedings of the 31st International Cosmic Ray Conference, Łódź 2009.
- [91] Dmitry Chirkin. Neutrino search with icecube. Technical report, University of Wisconsin-Madison, 2008.
- [92] Douglas R. Hofstadter. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, 1979.
- [93] Jakob van Santen and Nathan Whitehorn. Photospline: smooth, semi-analytic interpolation for photonics tables. Technical report, University of Wisconsin-Madison, 2011.
- [94] Nathan Whitehorn, Jakob van Santen, and Sven Lafebre. Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets. *Comput.Phys.Commun.*, 184:2214–2220, 2013.
- [95] M.G. Aartsen et al. Energy Reconstruction Methods in the IceCube Neutrino Telescope. *JINST*, 9:P03009, 2014.
- [96] F. James and M. Roos. Minuit - a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343 – 367, 1975.
- [97] Dirk Pandel. Determination of water and detector parameters and reconstruction of muons up to 100 TeV with the baikal neutrino telescope NT72. Diplom Physiker, Humboldt University, Berlin, 1996.

- [98] V.A. Balkanov, I.A. Belolaptikov, Leonid B. Bezrukov, N.M. Budnev, A.G. Chensky, et al. In-situ measurements of optical parameters in Lake Baikal with the help of a neutrino telescope. *Appl.Opt.*, 33:6818, 1999.
- [99] George Japaridze, Oladipo Fadiran, and Nick van Eijndhoven. Implementation of a gauss convoluted pandel pdf for track reconstruction in neutrino telescopes. Technical report, 2007.
- [100] Dmitry Chirkin. New implementation of convoluted mpe pandel function. Technical report, University of Wisconsin-Madison, 2011.
- [101] GI Angelis, AJ Reader, FA Kotasidis, WR Lionheart, and JC Matthews. The performance of monotonic and new non-monotonic gradient ascent reconstruction algorithms for high-resolution neuroreceptor pet imaging. *Phys Med Biol*, 56(13):3895–917, Jul 7 2011.
- [102] Kelen Tuttle. Logbook: Ultra-high-energy neutrinos. *Symmetry Magazine*, November 2013.
- [103] Lewis Carroll. *Alice’s Adventures in Wonderland*. Macmillan, 1865.
- [104] George W. Clark. Arrival Directions of Cosmic-Ray Air Showers from the Northern Sky. *Phys.Rev.*, 108:450–457, 1957.
- [105] A.U. Abeysekara et al. The HAWC Gamma-Ray Observatory: Observations of Cosmic Rays. 2013.
- [106] B. Bartoli et al. Measurement of the cosmic ray antiproton/proton flux ratio at TeV energies with the ARGO-YBJ detector. *Phys.Rev.*, D85:022002, 2012.
- [107] B. Bartoli, P. Bernardini, X.J. Bi, C. Bleve, I. Bolognino, et al. Observation of the Cosmic Ray Moon shadowing effect with ARGO-YBJ. *Phys.Rev.*, D84:022003, 2011.
- [108] Pierre Colin et al. Probing the CR positron/electron ratio at few hundreds GeV through Moon shadow observation with the MAGIC telescopes. 2011.
- [109] Morgan O’Hagan Wascko. Study of the shadow of the moon in very high-energy cosmic rays with the Milagrito water Cherenkov detector. 2001.
- [110] M. Amenomori. Multi-Tev Gamma-Ray Observation from the Crab Nebula Using the Tibet-III Air Shower Array Finely Tuned by the Cosmic-Ray Moon’s Shadow. *Astrophys.J.*, 692:61–72, 2009.
- [111] P. Adamson et al. Observation in the MINOS far detector of the shadowing of cosmic rays by the sun and moon. *Astropart.Phys.*, 34:457–466, 2011.
- [112] Carla Distefano and Colas Riviere. Observation of the Moon shadow with the ANTARES neutrino telescope. *PoS*, EPS-HEP2011:385, 2011.
- [113] N. Giglietto. Performance of the MACRO detector at Gran Sasso: Moon shadow and seasonal variations. *Nuclear Physics B Proceedings Supplements*, 61:180–184, February 1997.
- [114] George W. Clark. Arrival Directions of Cosmic-Ray Air Showers from the Northern Sky. *Phys.Rev.*, 108:450–457, 1957.
- [115] M.G. Aartsen et al. Observation of the cosmic-ray shadow of the Moon with IceCube. *Phys.Rev.*, D89:102004, 2014.
- [116] *Moon Shadow Observation by IceCube*, 2009. Proceedings of the 31st International Cosmic Ray Conference, Łódź 2009.

- [117] T.-P. Li and Y.-Q. Ma. Analysis methods for results in gamma-ray astronomy. *Astrophys.J.*, 272:317–324, 1983.
- [118] R. Abbasi et al. Time-Integrated Searches for Point-like Sources of Neutrinos with the 40-String IceCube Detector. *Astrophys.J.*, 732:18, 2011.
- [119] *The Shadow of the Moon in IceCube*, 2010. Rencontres de Moriond Electroweak, La Thuile, Italy.
- [120] Hugo Stiebel. Study of the angular resolution of the IceCube 59 string detector using the shadowing of cosmic rays by the Moon. Master’s thesis, Stockholm University, 2011. Available upon request from IceCube internal documents, collection 8819.
- [121] Jan Blumenthal. Measurements of the shadowing of cosmic rays by the moon with the icecube neutrino observatory. Diplomarbeit in Physik, Rheinisch-Westfälischen Technischen Hochschule Aachen, 2011. Available upon request from IceCube internal documents, collection 8819.
- [122] J.R. Klein and A. Roodman. Blind analysis in nuclear and particle physics. *Ann.Rev.Nucl.Part.Sci.*, 55:141–163, 2005.
- [123] Sebastian Euler. *Observations of oscillations of atmospheric neutrinos with the IceCube Neutrino Observatory*. PhD thesis, Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, 2013.
- [124] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J.N. Capdevielle. CORSIKA: A Monte Carlo code to simulate extensive air showers, 1998.
- [125] Askhat Gazizov and Marek P. Kowalski. ANIS: High energy neutrino generator for neutrino telescopes. *Comput.Phys.Commun.*, 172:203–213, 2005.
- [126] M. Dhanamjaya, V.G. Talawar, K.R.Mulla, and N. Chowdappa. Visibility of reference patterns in the doctoral theses of engineering and technology: A citation analysis study. *PEARL - A Journal of Library and Information Science*, 5:2:59–70, 2011.
- [127] Adam M. Dziewonski and Don L. Anderson. Preliminary reference earth model. *Physics of the Earth and Planetary Interiors*, 25(4):297 – 356, 1981.
- [128] Thomas K. Gaisser, Kyle Jero, Albrecht Karle, and Jakob van Santen. A generalized self-veto probability for atmospheric neutrinos. *Phys.Rev.*, D90:023009, 2014.
- [129] Dmitry Chirkin. Likelihood description for comparing data with simulation of limited statistics. 2013.
- [130] Roger J. Barlow and Christine Beeston. Fitting using finite Monte Carlo samples. *Comput.Phys.Commun.*, 77:219–228, 1993.
- [131] Louis Lyons. Bayes and Frequentism: a Particle Physicist’s perspective. 2013.
- [132] S.S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Annals Math.Statist.*, 9(1):60–62, 1938.
- [133] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys.Rev.*, D57:3873–3889, 1998.
- [134] Mayly C. Sanchez et al. Measurement of the L/E distributions of atmospheric neutrinos in Soudan 2 and their interpretation as neutrino oscillations. *Phys.Rev.*, D68:113004, 2003.

- [135] Sean Grullon. *A Search for a Diffuse Flux of Astrophysical Muon Neutrinos With the IceCube Neutrino Observatory in the 40-String Configuration*. PhD thesis, University of Wisconsin Madison, 2010.
- [136] Jose R. Alonso for the DAE δ ALUS Collaboration. The dae δ alus project: Rationale and beam requirements. 2010.
- [137] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl.Instrum.Meth.*, A506:250–303, 2003.

Appendix A: Supplemental Plots for Level 8 Cuts

Many plots went into designing the event selection presented in this analysis, especially Ch. 7. Some of the distributions are especially relevant to the analysis; the others are presented in this chapter for future reference and comparison. The purpose of many of these plots is to investigate whether the variable in question would be a good cut variable. As such, the relevant questions are:

- Do the experimental and simulated data agree reasonably well across a wide range, or if not, can we cut out the region of disagreement without damaging the oscillation signal?
- Does this variable separate neutrinos (especially muon neutrinos) from cosmic ray background generated by the CORSIKA simulation generator?
- Does this variable correlate with energy reconstruction resolution, and if so, is it in a specific region?
- Does this variable correlate with zenith angle reconstruction resolution, and if so, is it in a specific region?

The plots on the following pages are designed to address these questions. Each of these questions is best addressed with a specific format of plot. The purpose and conclusion from each plot is addressed in its caption.

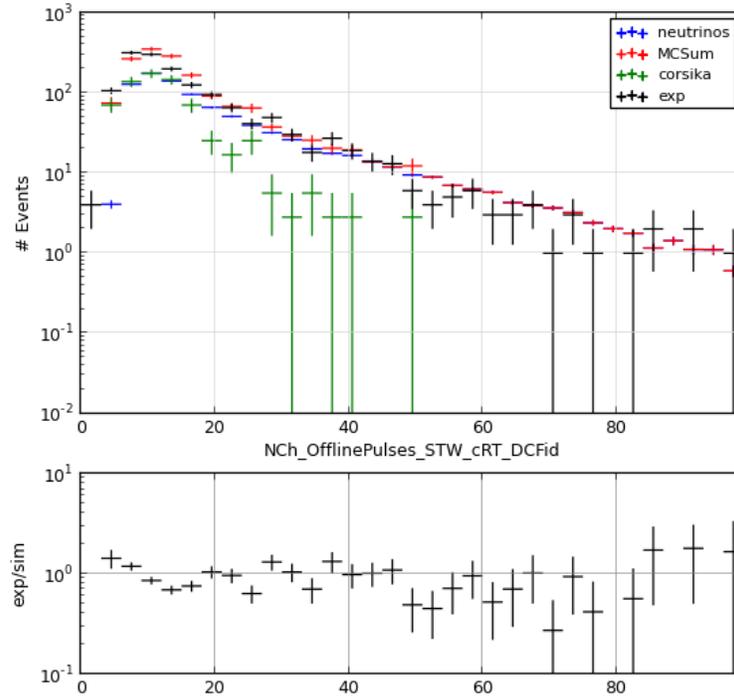
A.1 Designing Level 8 Cuts

The plots shown in this section reflect the statistics of the Level 7 sample, but with the reconstruction algorithms of Level 8 processing.

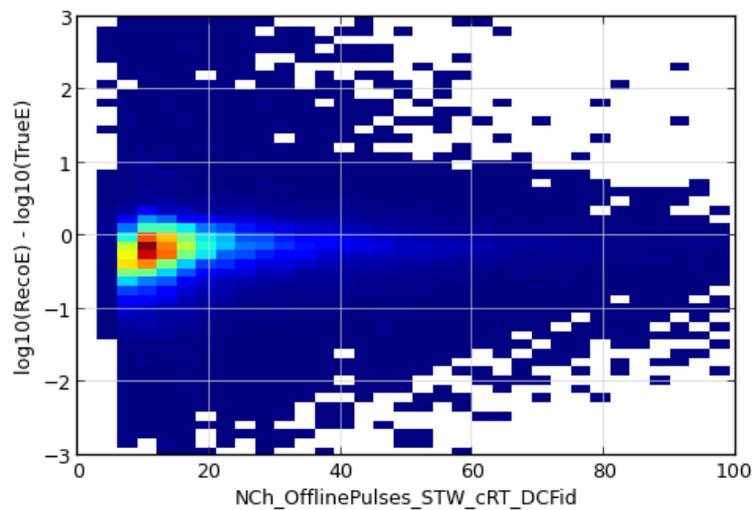
At cut Level 7, a reasonable pure neutrino sample has already been isolated. The Level 8 processing and cuts are based around the energy reconstruction Monopod, which was run on all events in the Level 7 sample (after the reconstruction and cuts are both performed, the sample is called “Level 8”).

The main reconstruction in the Level 8 processing is Monopod, which is used as the energy reconstruction in the final analysis. The purpose of the Level 8 cuts is to identify events with good energy resolution, and to establish agreement between experimental data and simulation. Zenith angle reconstruction and quality is not considered until the next cut level, where a zenith angle reconstruction is performed.

Figure A.1: NCh, after hit cleaning, at cut level 7. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

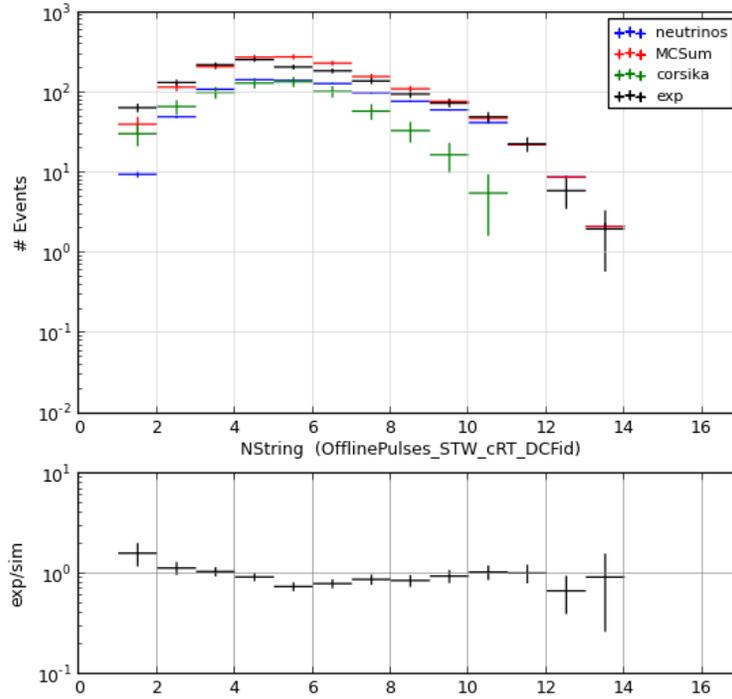


(a) NCh after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

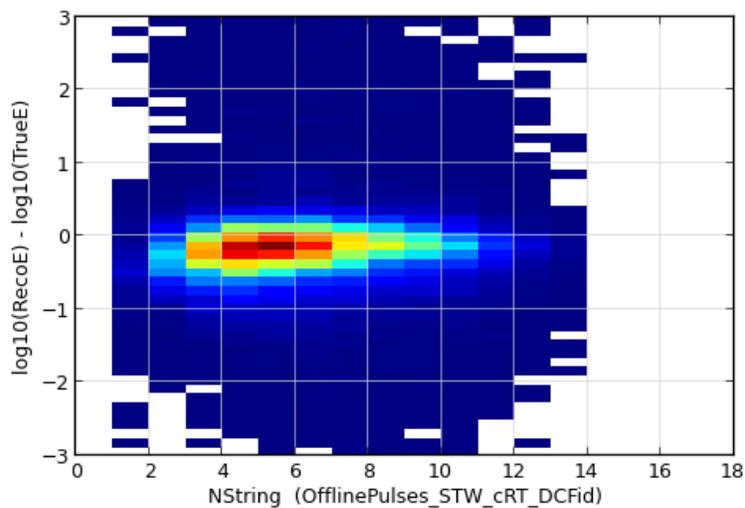


(b) NCh after hit cleaning versus energy reconstruction error

Figure A.2: NString, after hit cleaning, at cut level 7. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

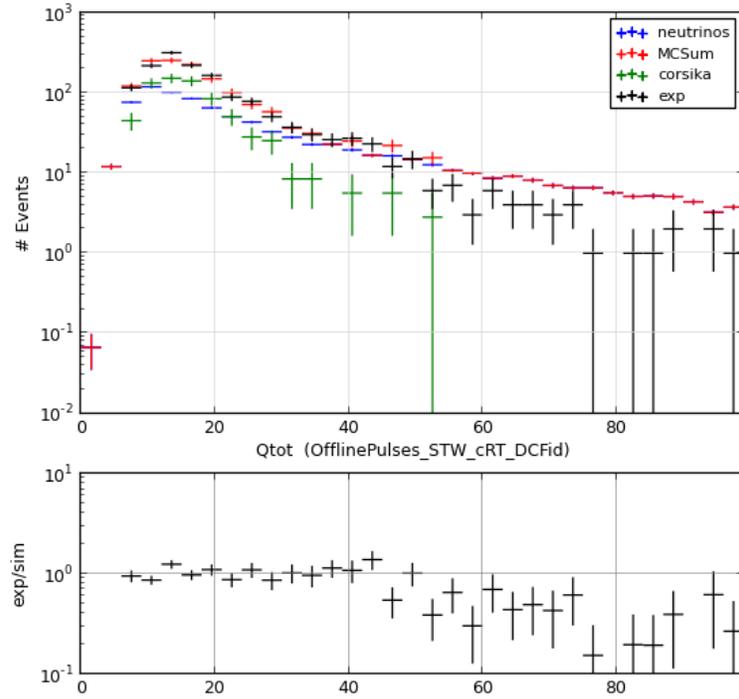


(a) NString after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

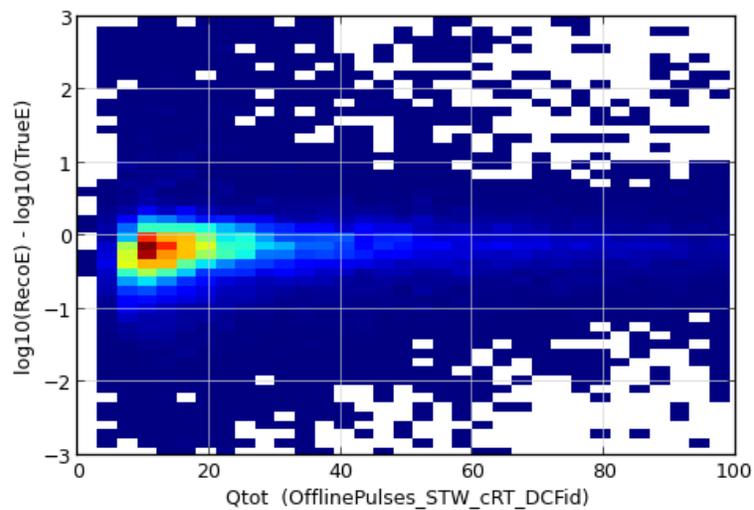


(b) NString after hit cleaning versus energy reconstruction error

Figure A.3: Total charge, after hit cleaning, at cut level 7. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

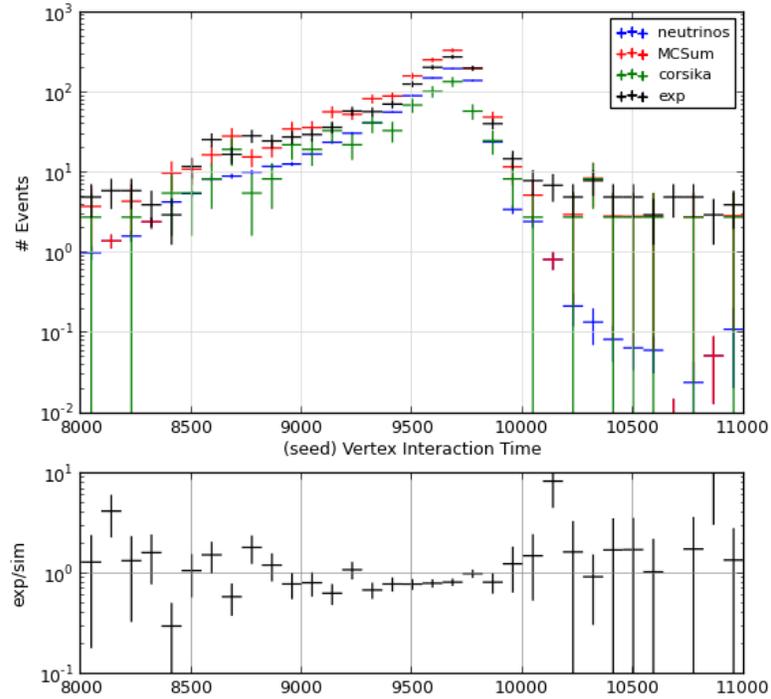


(a) Total charge after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

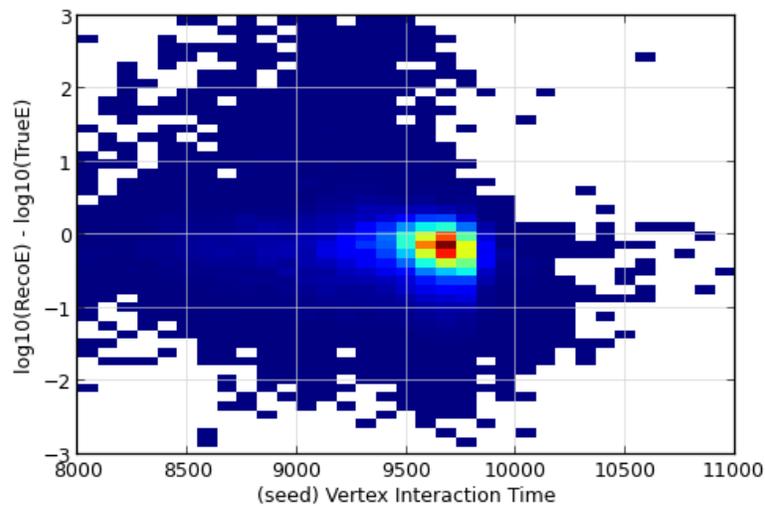


(b) Total charge after hit cleaning versus energy reconstruction error

Figure A.4: CLast reconstructed vertex time at cut level 7. For an explanation of how the CLast vertex time is calculated, see §5.7.

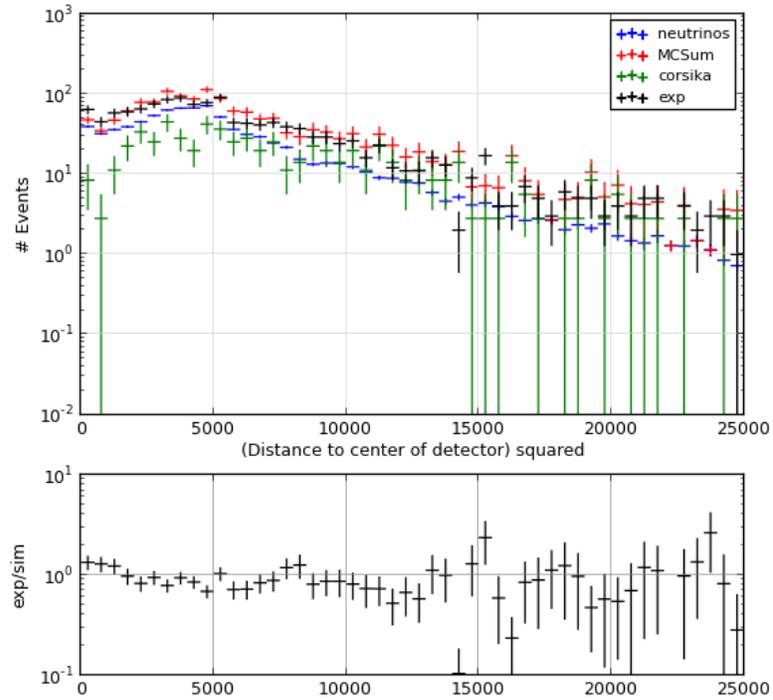


(a) CLast reconstructed vertex time versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

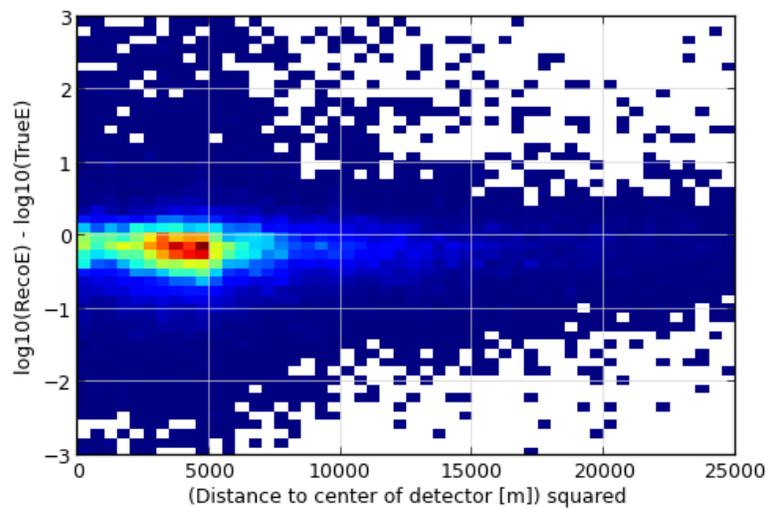


(b) CLast reconstructed vertex time versus energy reconstruction error

Figure A.5: Distance between the central string and center-of-gravity reconstructed vertex, squared, at cut level 7

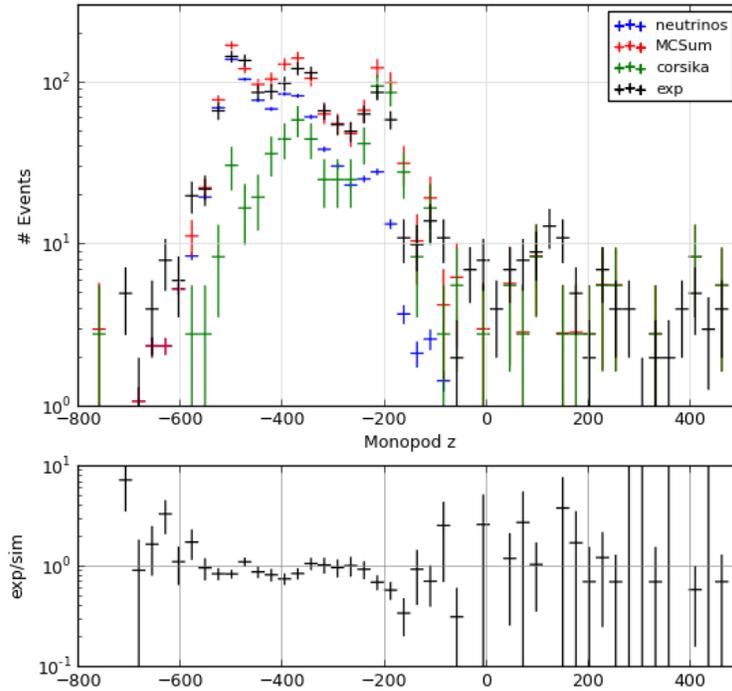


(a) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

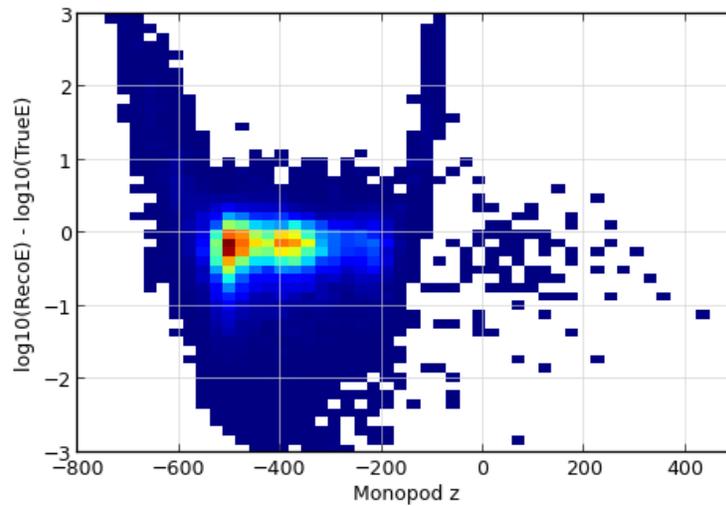


(b) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus energy reconstruction error

Figure A.6: Monopod reconstructed depth at cut level 7

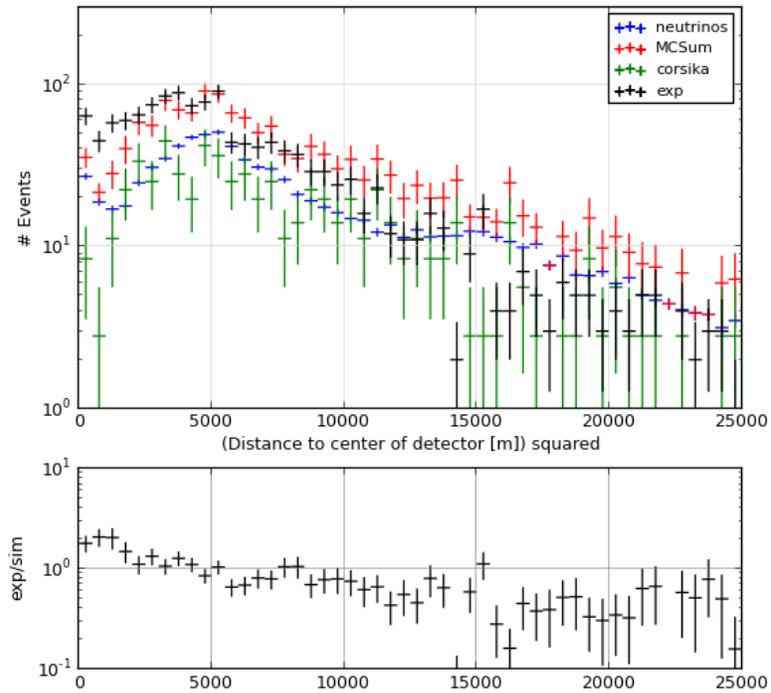


(a) Monopod reconstructed depth versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

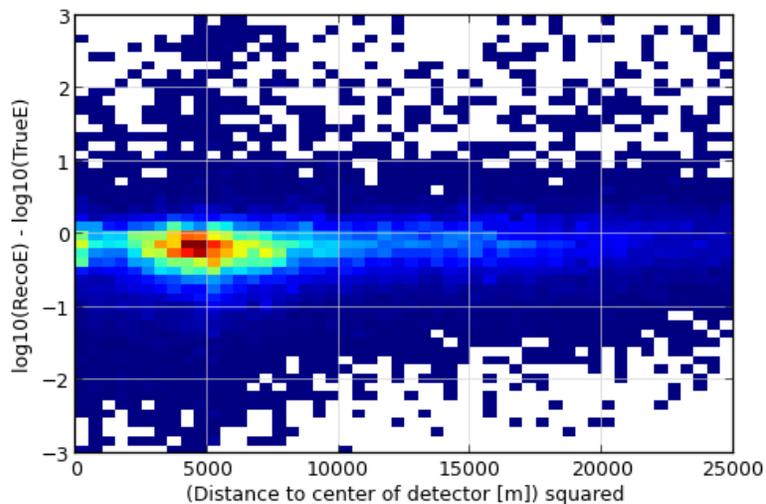


(b) Monopod reconstructed depth versus energy reconstruction error

Figure A.7: Distance between the central string and the Monopod reconstructed vertex, squared, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1.

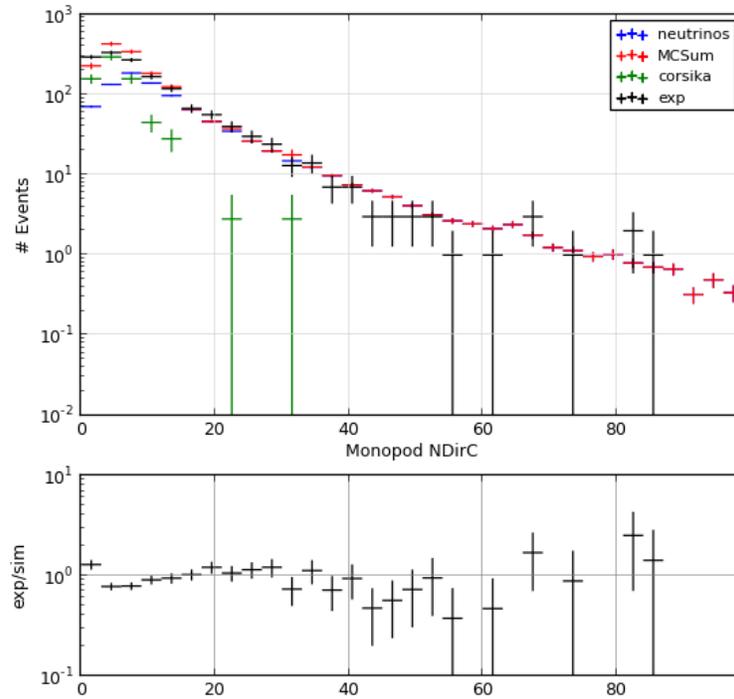


(a) Distance between the central string and the Monopod reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

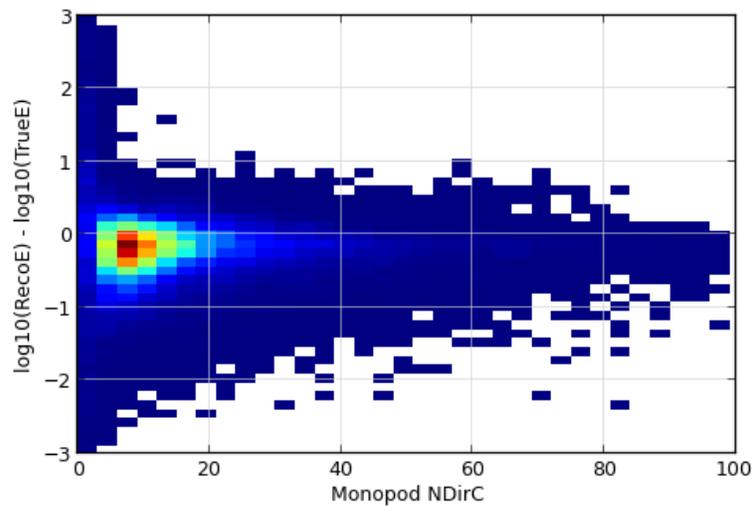


(b) Distance between the central string and the Monopod reconstructed vertex, squared, versus energy reconstruction error

Figure A.8: Number of direct hits relative to the Monopod reconstructed interaction vertex, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1. Direct hits are those that arrive between 15 nsec before the expected light arrival time and 75 nsec after the expected time.

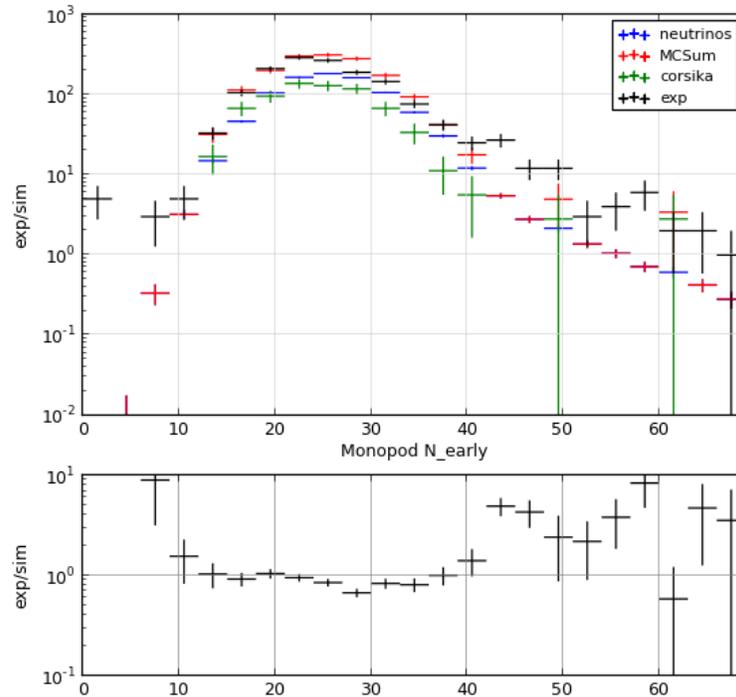


(a) Number of Direct hits relative to Monopod versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

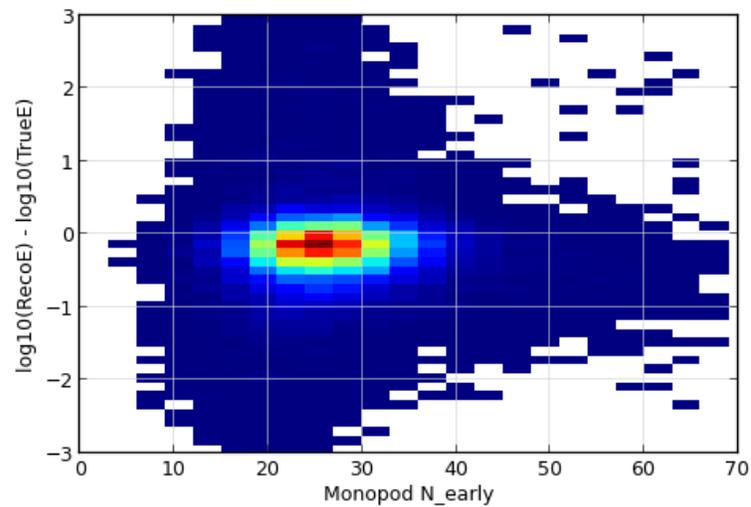


(b) Number of direct hits relative to Monopod versus energy reconstruction error

Figure A.9: Number of early pulses relative to the Monopod reconstructed interaction vertex, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1. Early pulses are those that arrive 15 nsec or more before the expected light arrival time.

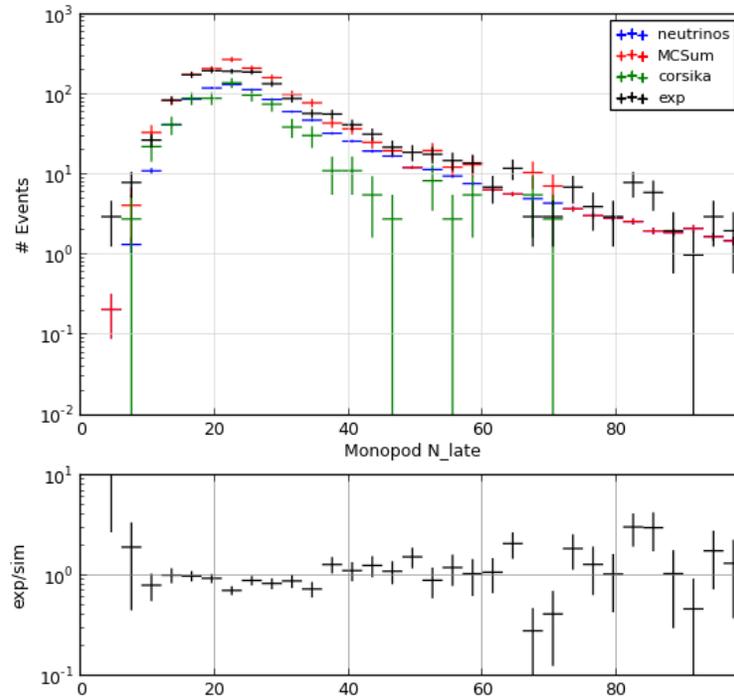


(a) Number of early pulses relative to the Monopod reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

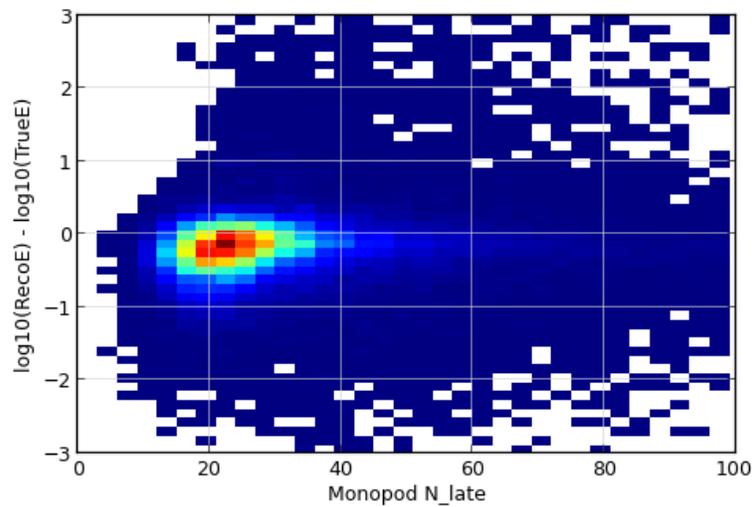


(b) Number of early pulses relative to the Monopod reconstructed interaction vertex versus energy reconstruction error

Figure A.10: Number of late pulses relative to the Monopod reconstructed interaction vertex, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1. Late pulses are those that arrive 250 nsec or more after the expected light arrival time.

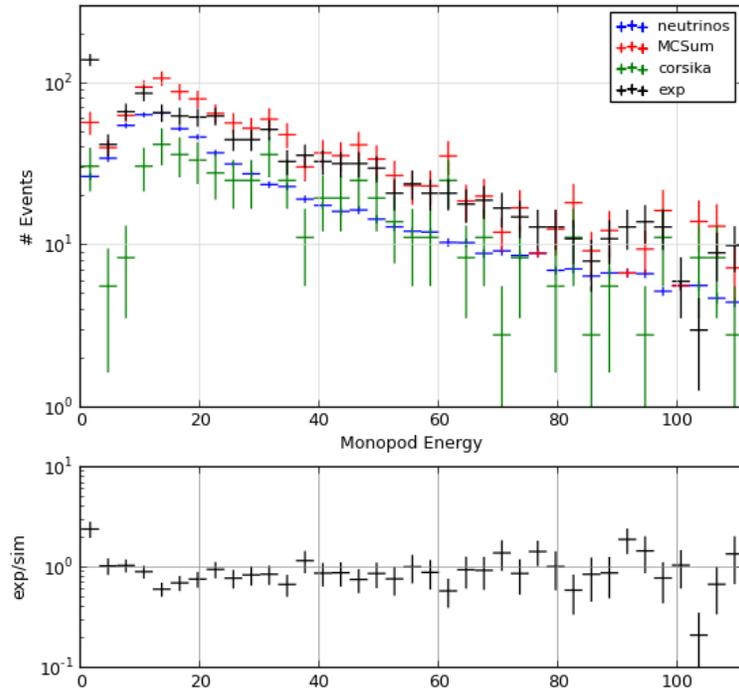


(a) Number of late pulses relative to the Monopod reconstructed interaction vertex versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

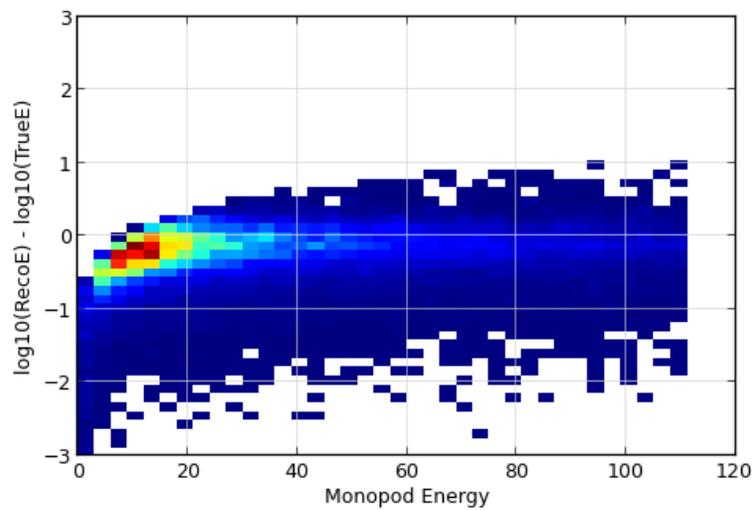


(b) Number of late pulses relative to the Monopod reconstructed interaction vertex versus energy reconstruction error

Figure A.11: Reconstructed energy, using the Monopod reconstruction algorithm, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1.

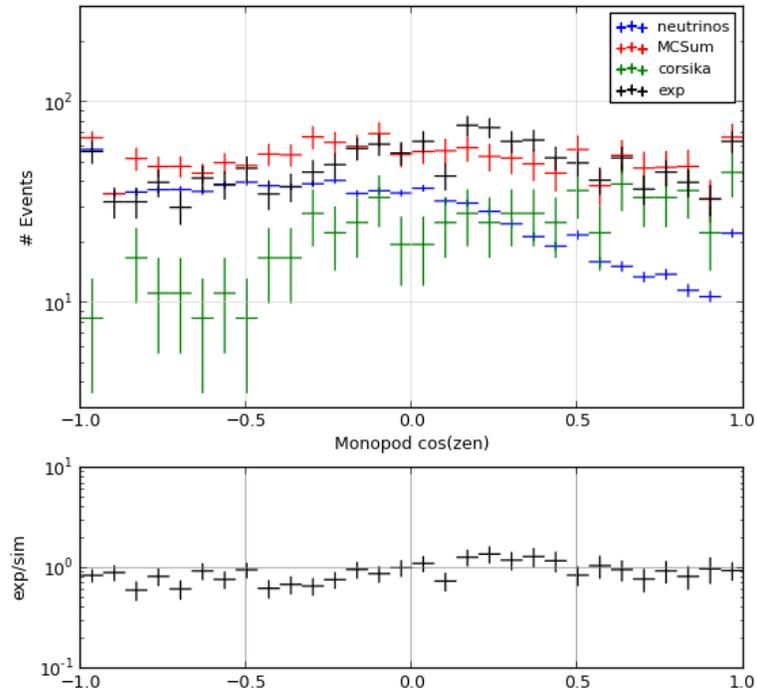


(a) Reconstructed Energy versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

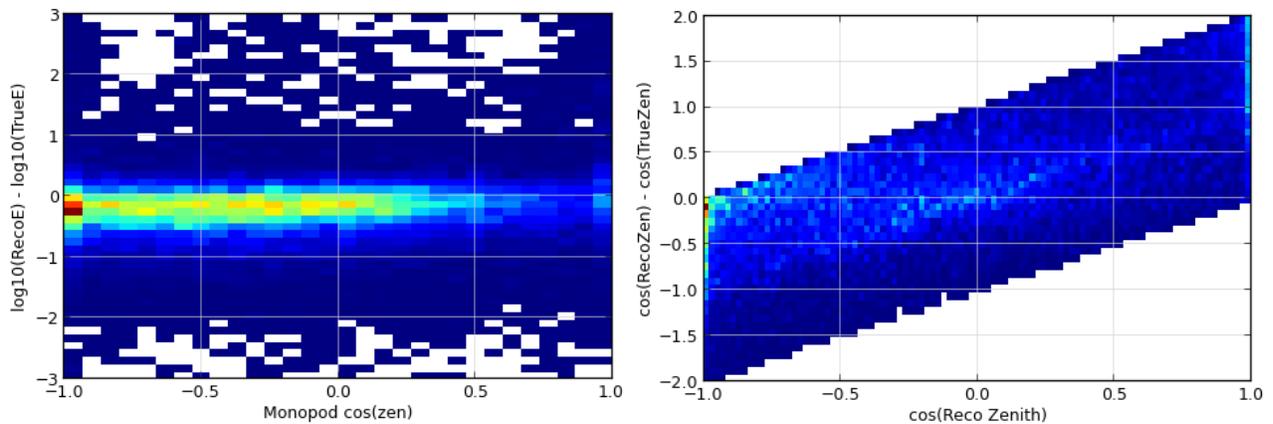


(b) Reconstructed Energy versus energy reconstruction error

Figure A.12: Reconstructed Zenith, using the Monopod reconstruction algorithm, at cut level 7. For more information about the Monopod reconstruction algorithm, see §5.11.1. Monopod is not expected to be an accurate zenith angle reconstruction, but is included here to allow comparison.



(a) Reconstructed Energy versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



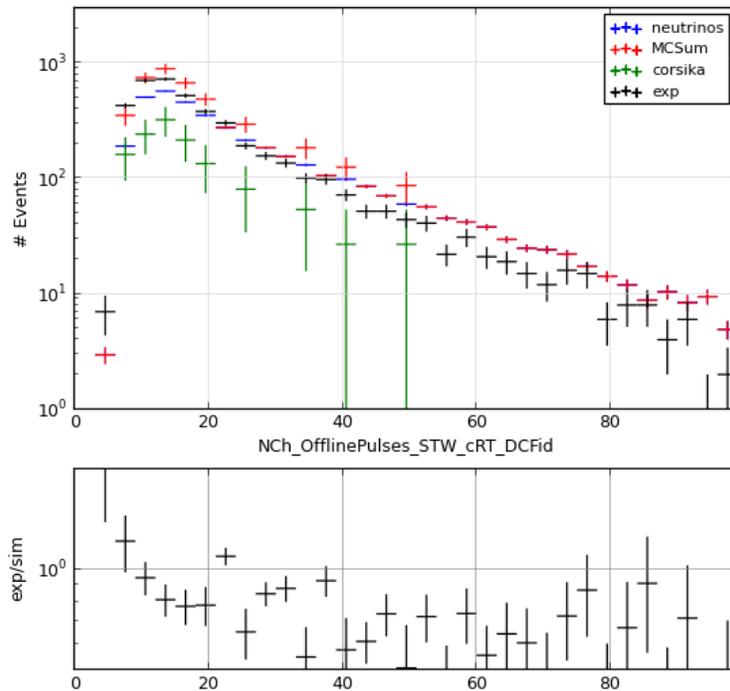
(b) Reconstructed Energy versus energy reconstruction error (c) Reconstructed Energy versus zenith angle reconstruction error.

Appendix B: Supplemental Plots for Level 9 Cuts

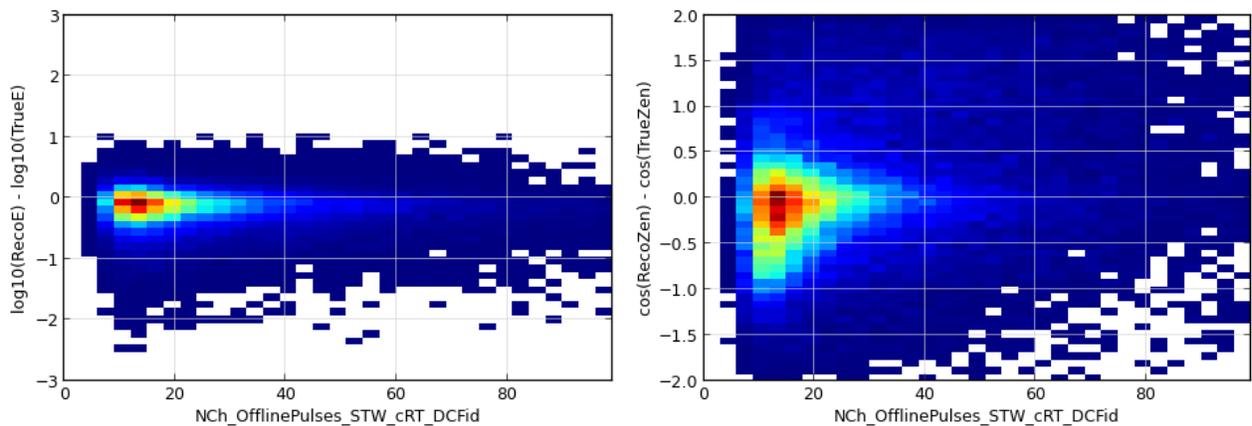
The plots shown in this section reflect the statistics of the Level 8 sample while considering variables calculated during Level 9 processing. For continuity, variables from the Level 8 processing are included also.

The main reconstruction at Level 9 is the Igelfit directional reconstruction; the zenith angle component of the output is used as an observable in the final analysis. The cuts at level 9 are deigned to identify an event sample with good zenith angle resolution while preserving: agreement between experimental and simulated data, purity of the neutrino sample against the cosmic ray muon background, and a large enough statistical sample to analyze.

Figure B.1: NCh, after hit cleaning, at cut level 8. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

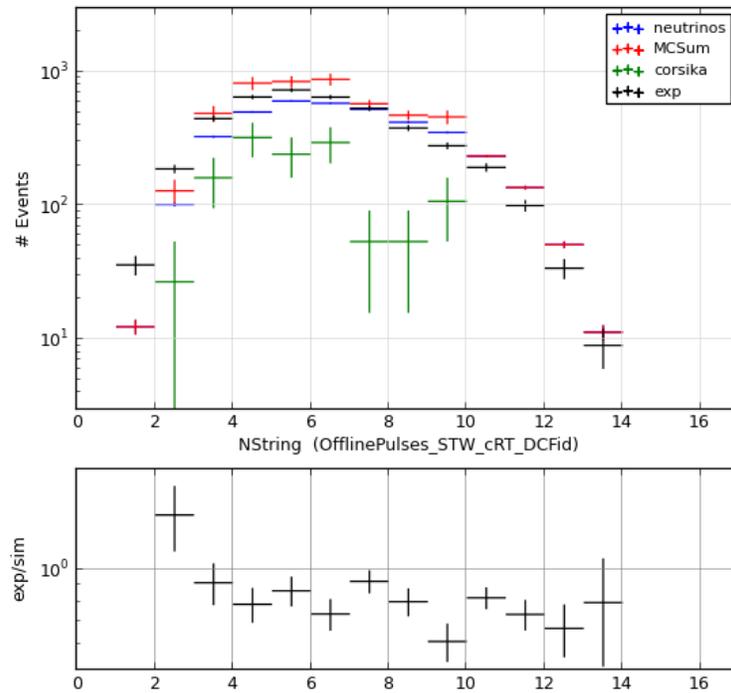


(a) NCh after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

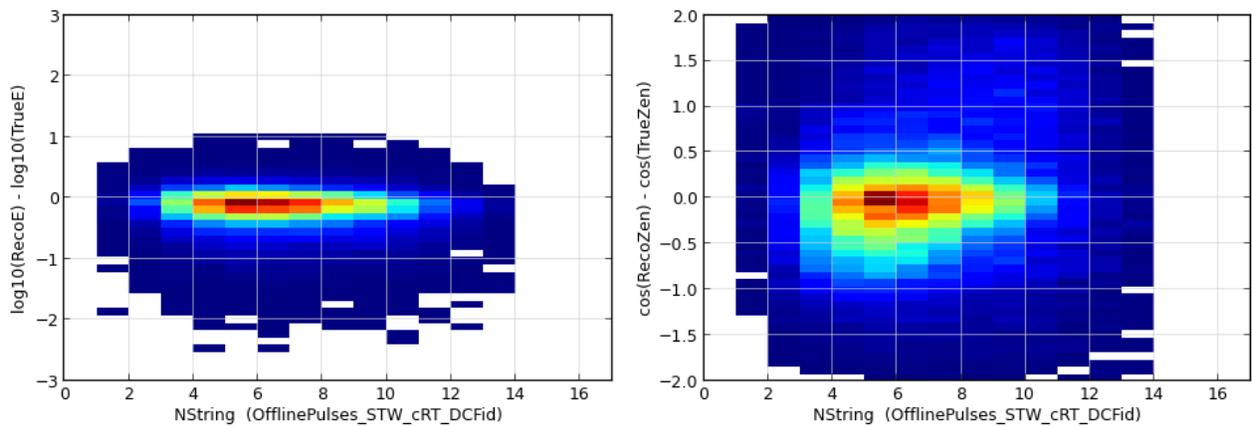


(b) NCh after hit cleaning versus Monopod energy reconstruction error (c) NCh after hit cleaning versus Igefit energy reconstruction error

Figure B.2: NString, after hit cleaning, at cut level 8. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

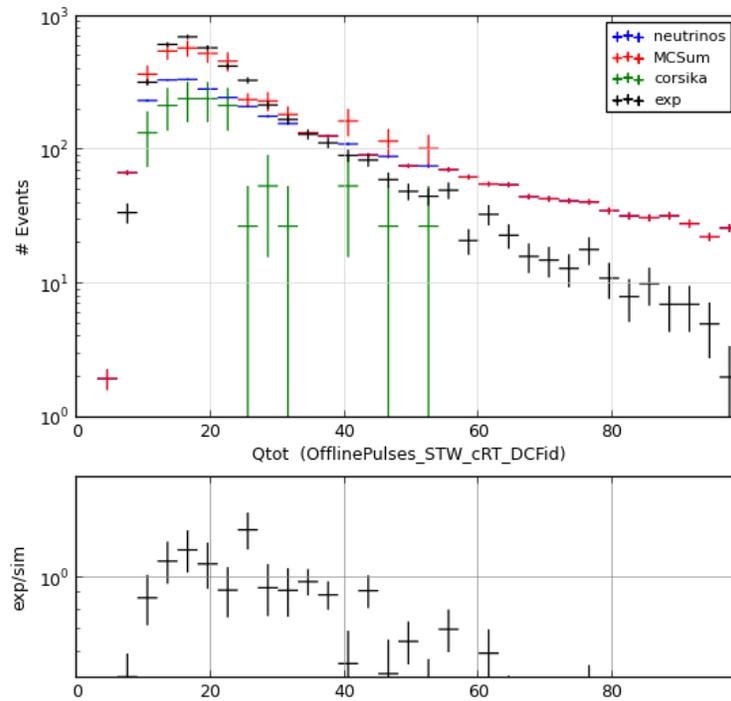


(a) NString after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

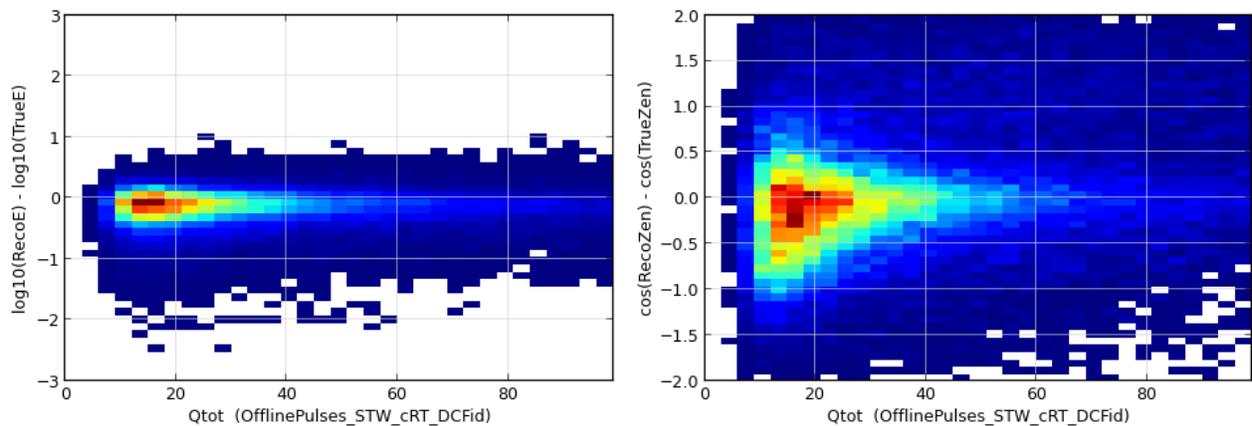


(b) NString after hit cleaning versus Monopod energy reconstruction error (c) NString after hit cleaning versus Igelfit zenith angle reconstruction error

Figure B.3: Total charge, after hit cleaning, at cut level 8. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

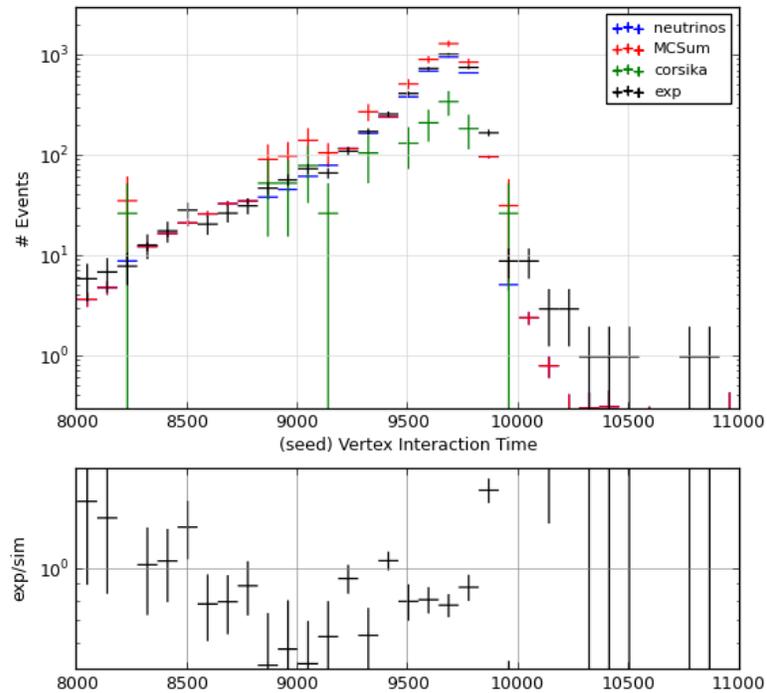


(a) Total charge after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

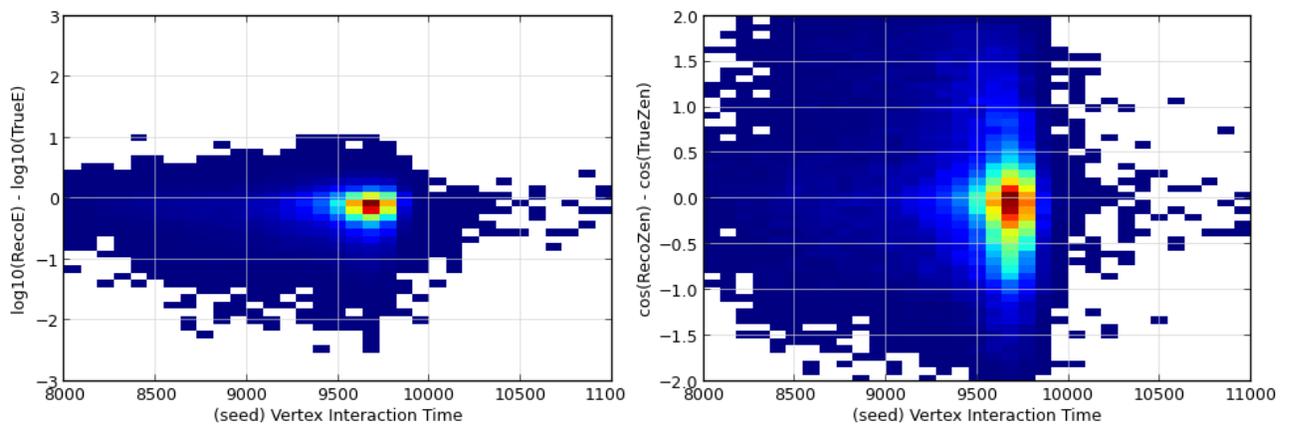


(b) Total charge after hit cleaning versus Monopod energy reconstruction error (c) Total charge after hit cleaning versus Igelfit zenith angle reconstruction error

Figure B.4: CLast reconstructed vertex time at cut level 8. For an explanation of how the CLast vertex time is calculated, see §5.7.



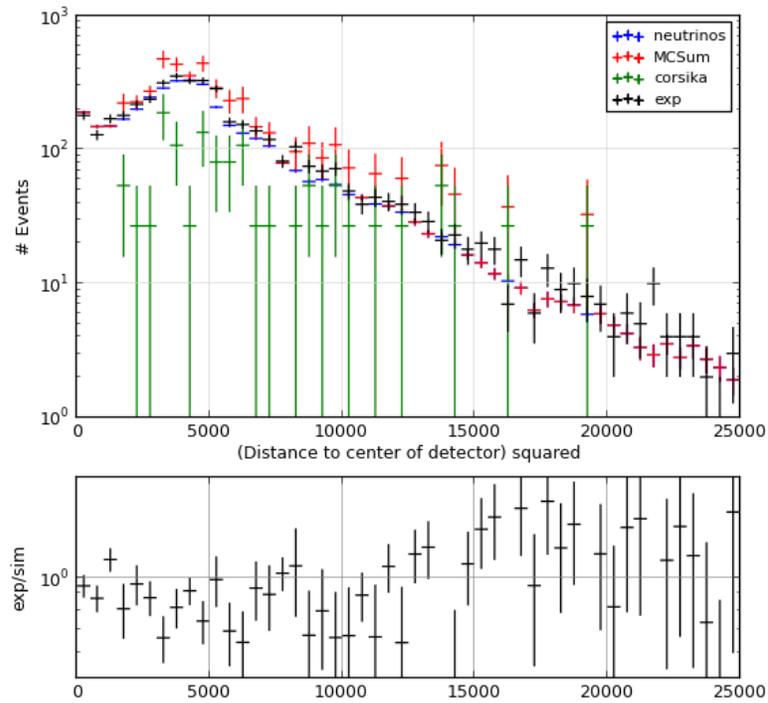
(a) CLast reconstructed vertex time versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



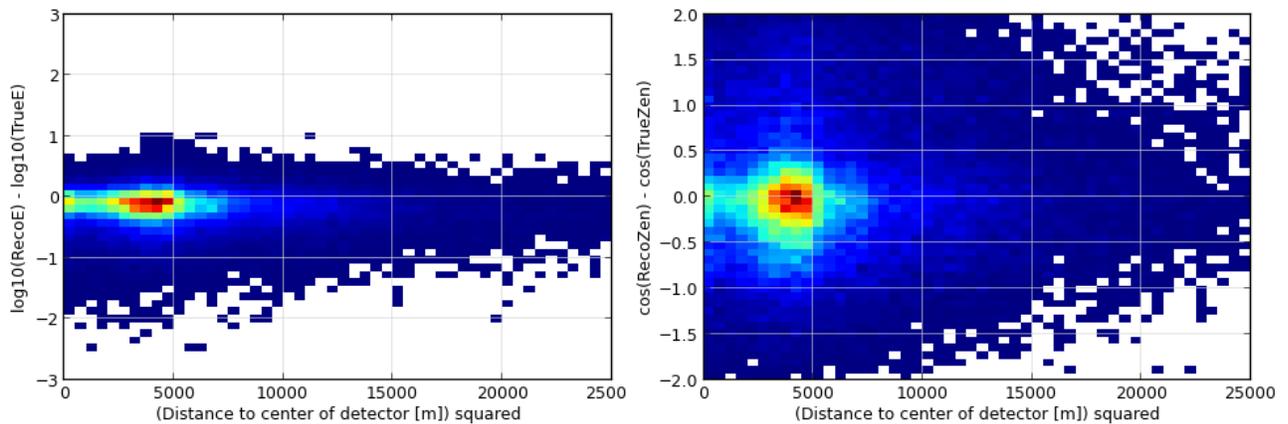
(b) CLast reconstructed vertex time versus Monopod energy reconstruction error

(c) CLast reconstructed vertex time versus Igelift zenith angle reconstruction error

Figure B.5: Distance between the central string and center-of-gravity reconstructed vertex, squared, at cut level 8

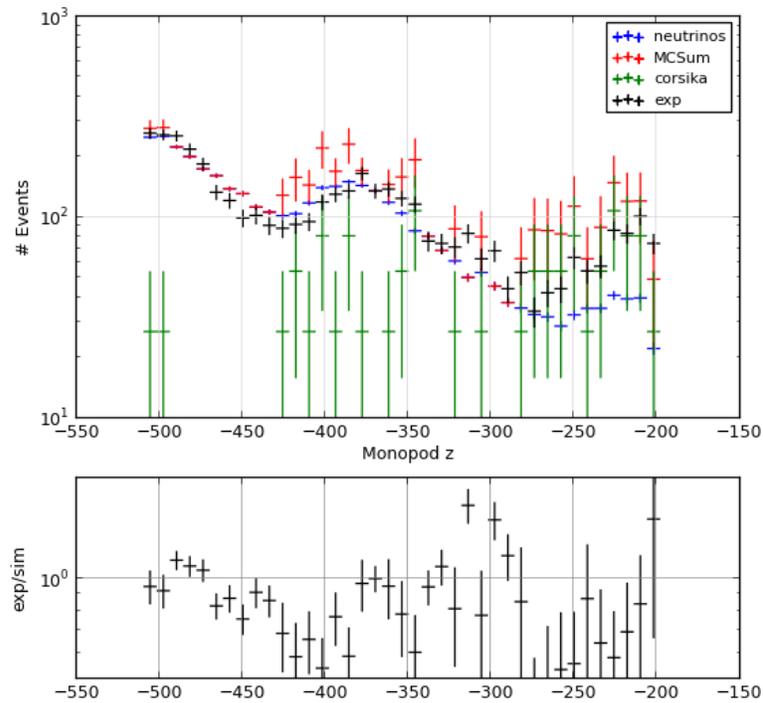


(a) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

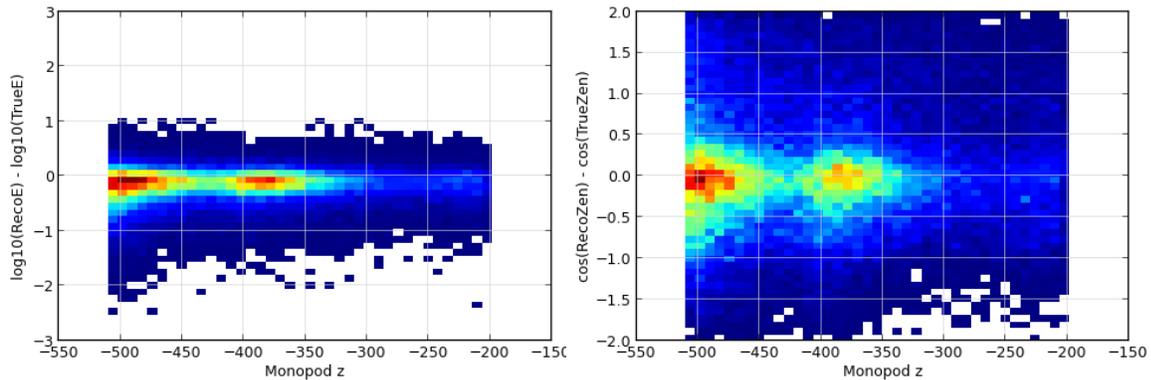


(b) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus Monopod energy reconstruction error
 (c) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus Igelift zenith angle reconstruction error

Figure B.6: Monopod reconstructed depth at cut level 8

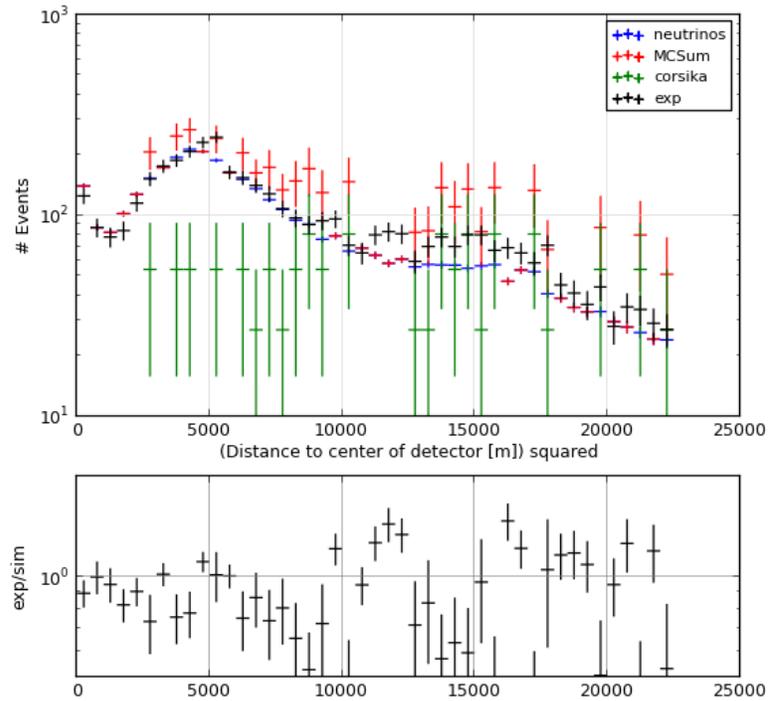


(a) Monopod reconstructed depth versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

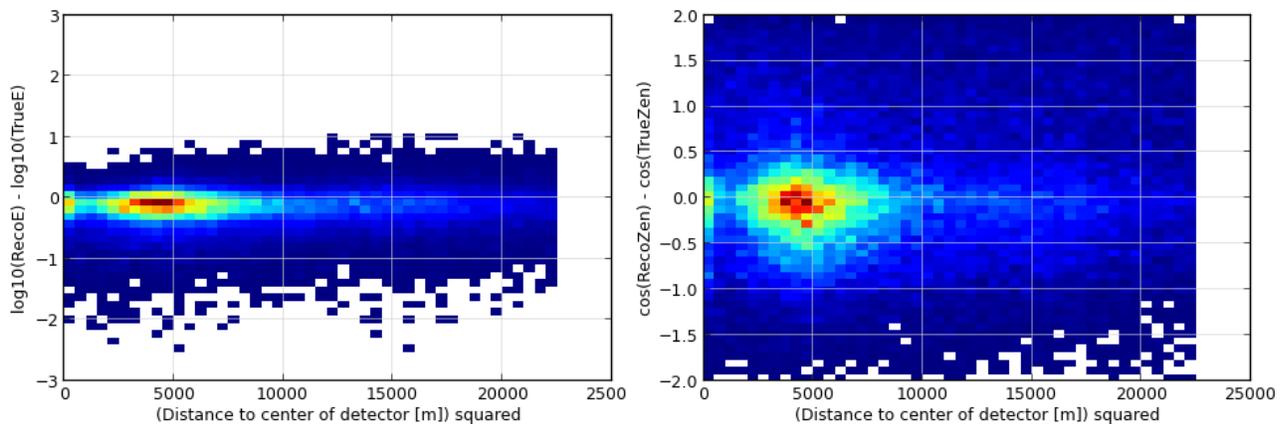


(b) Monopod reconstructed depth versus Monopod energy (c) Monopod reconstructed depth versus Igefit zenith angle reconstruction error

Figure B.7: Distance between the central string and the Monopod reconstructed vertex, squared, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1.

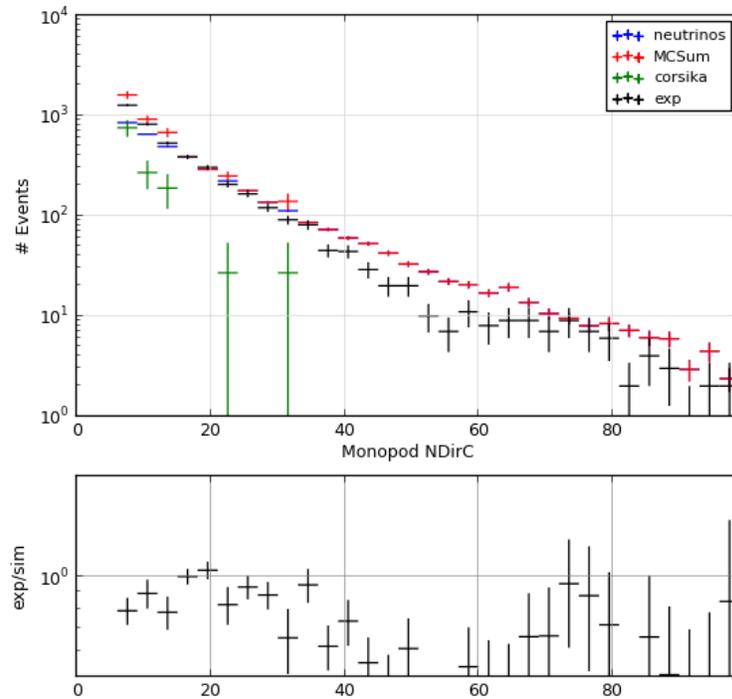


(a) Distance between the central string and the Monopod reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

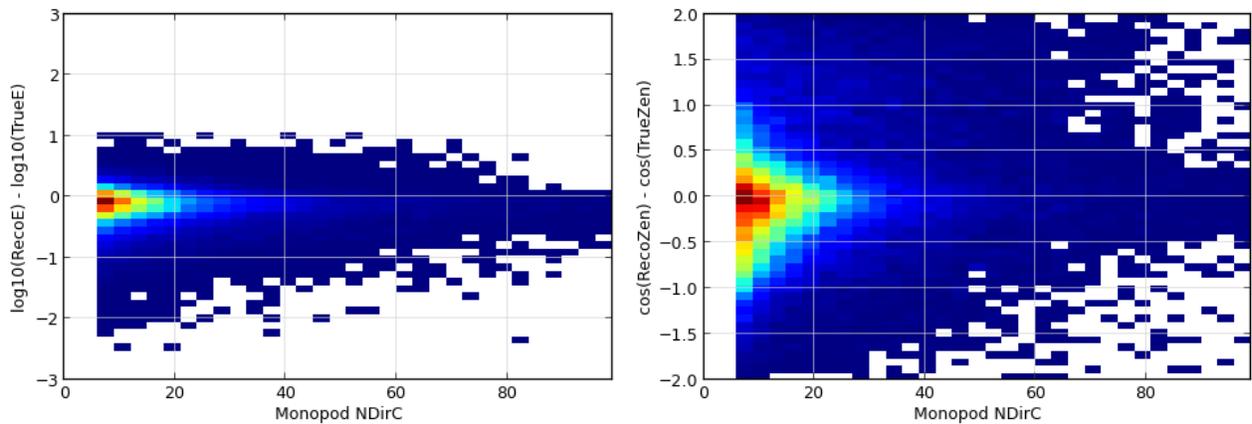


(b) Distance between the central string and the Monopod reconstructed vertex, squared, versus Monopod energy reconstruction error
 (c) Distance between the central string and the Monopod reconstructed vertex, squared, versus Igelift zenith angle reconstruction error

Figure B.8: Number of direct hits relative to the Monopod reconstructed interaction vertex, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1. Direct hits are those that arrive between 15 nsec before the expected light arrival time and 75 nsec after the expected time.

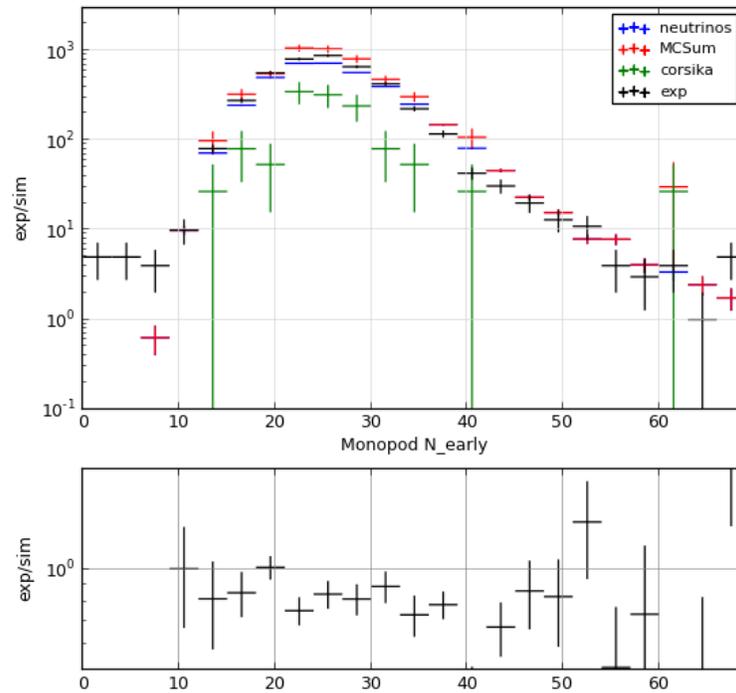


(a) Number of direct hits relative to Monopod versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

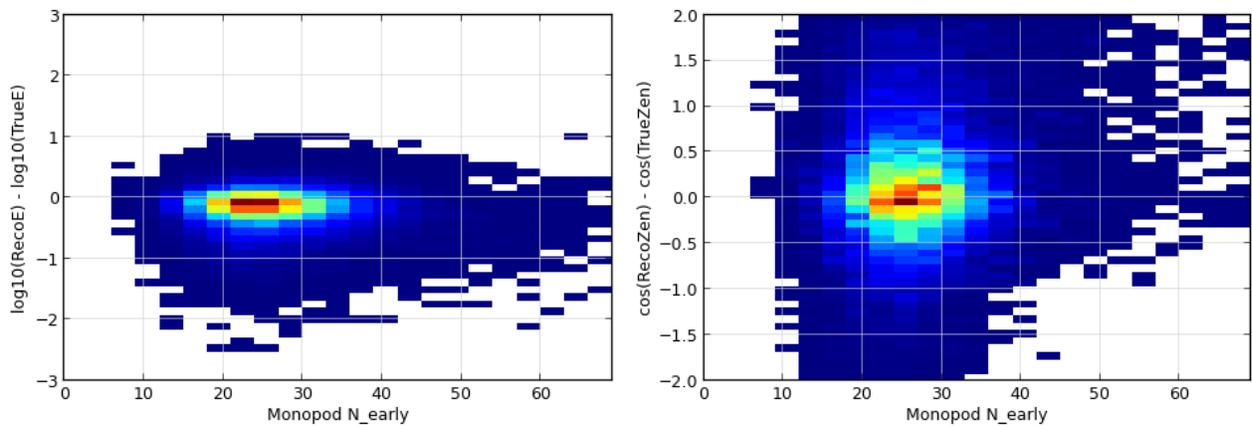


(b) Number of direct hits relative to Monopod versus Monopod energy reconstruction error (c) Number of direct hits relative to Monopod versus Igefitt zenith angle reconstruction error

Figure B.9: Number of early pulses relative to the Monopod reconstructed interaction vertex, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1. Early pulses are those that arrive 15 nsec or more before the expected light arrival time.

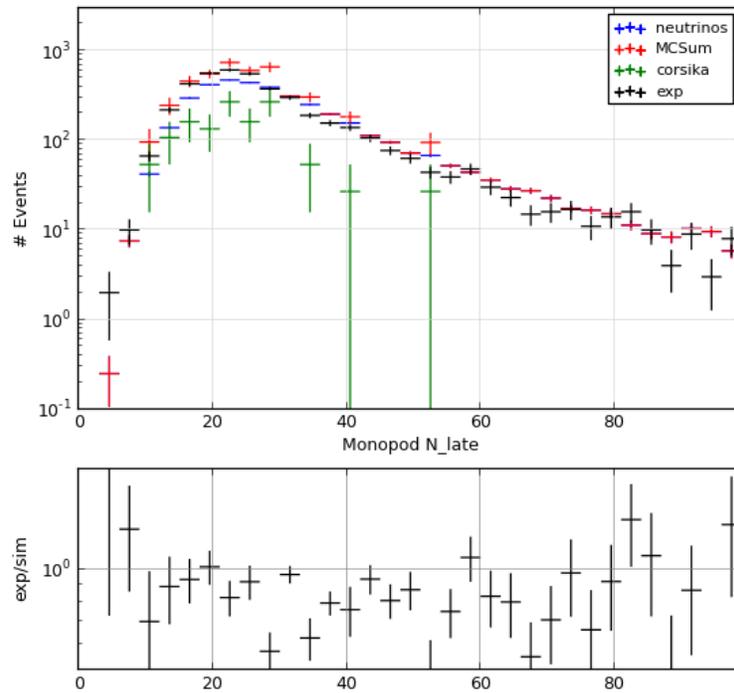


(a) Number of early pulses relative to the Monopod reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

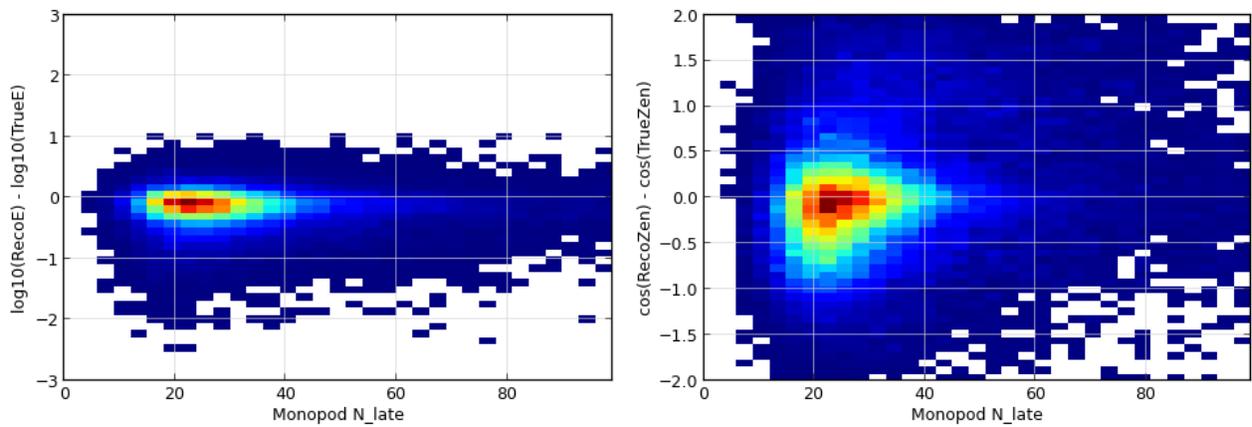


(b) Number of early pulses relative to the Monopod reconstructed interaction vertex versus Monopod energy reconstruction error (c) Number of early pulses relative to the Monopod reconstructed interaction vertex versus Igefitt zenith angle reconstruction error

Figure B.10: Number of late pulses relative to the Monopod reconstructed interaction vertex, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1. Late pulses are those that arrive 250 nsec or more after the expected light arrival time.

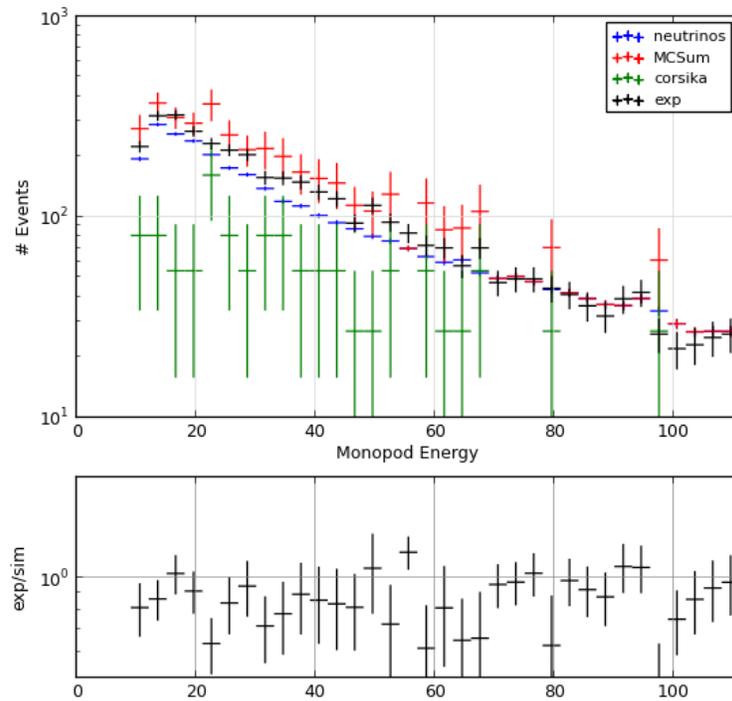


(a) Number of late pulses relative to the Monopod reconstructed interaction vertex versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

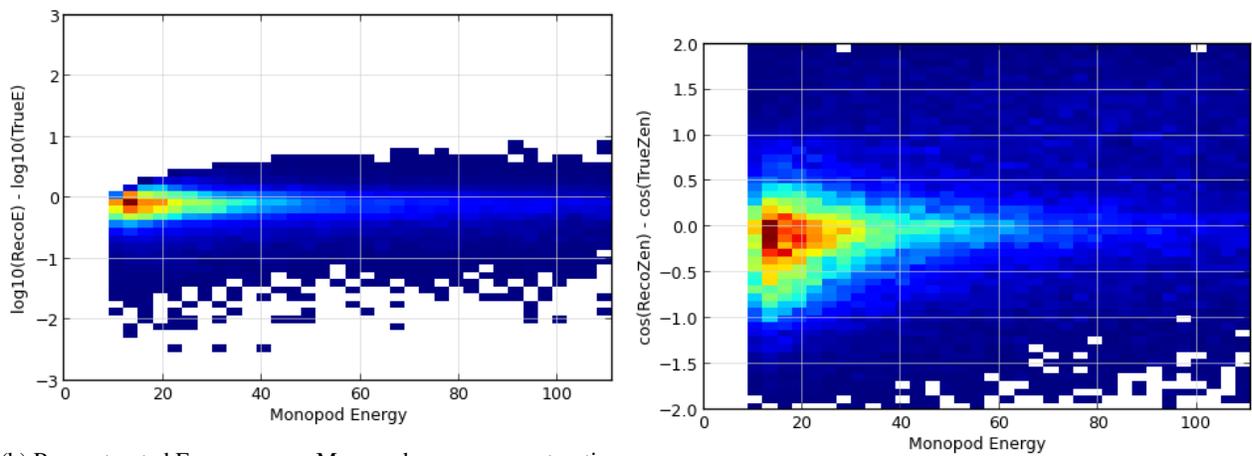


(b) Number of late pulses relative to the Monopod reconstructed interaction vertex versus Monopod energy reconstruction error (c) Number of late pulses relative to the Monopod reconstructed interaction vertex versus Igefit zenith angle reconstruction error

Figure B.11: Reconstructed energy, using the Monopod reconstruction algorithm, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1.



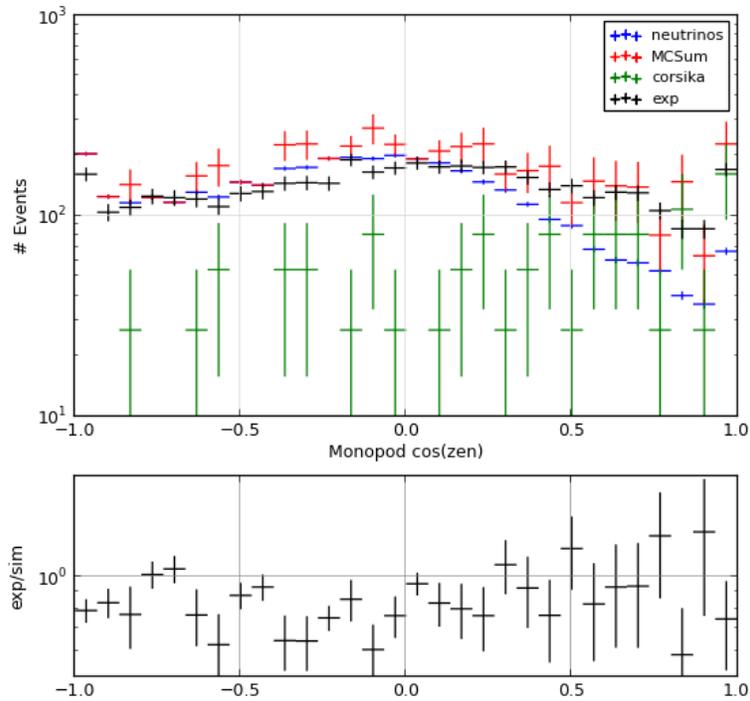
(a) Reconstructed Energy versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



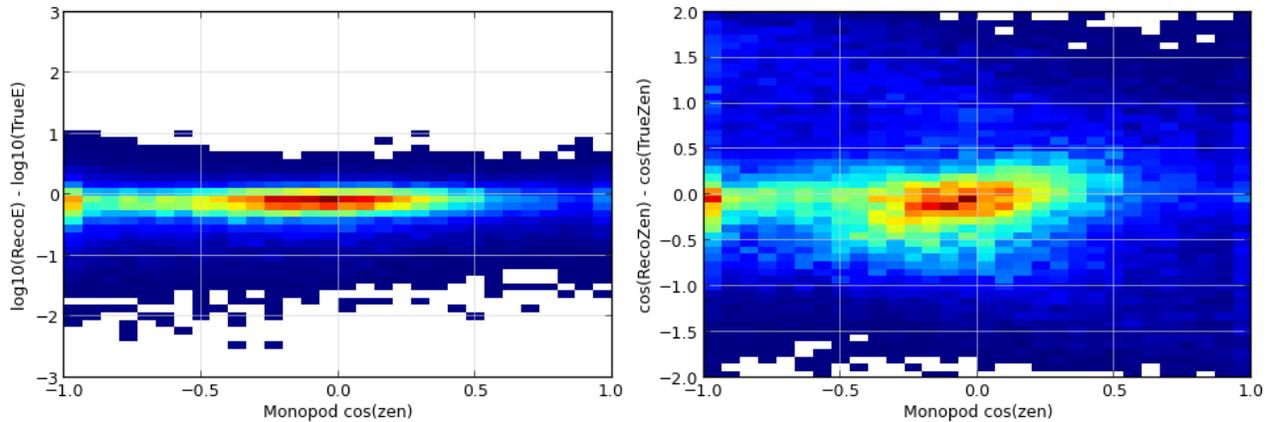
(b) Reconstructed Energy versus Monopod energy reconstruction error

(c) Reconstructed Energy versus Igelift reconstruction error

Figure B.12: Reconstructed Zenith, using the Monopod reconstruction algorithm, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1. Monopod is not expected to be an accurate reconstruction for zenith angle, but the plot is included here for comparison to the Igelfit reconstructed zenith.



(a) Reconstructed zenith angle versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



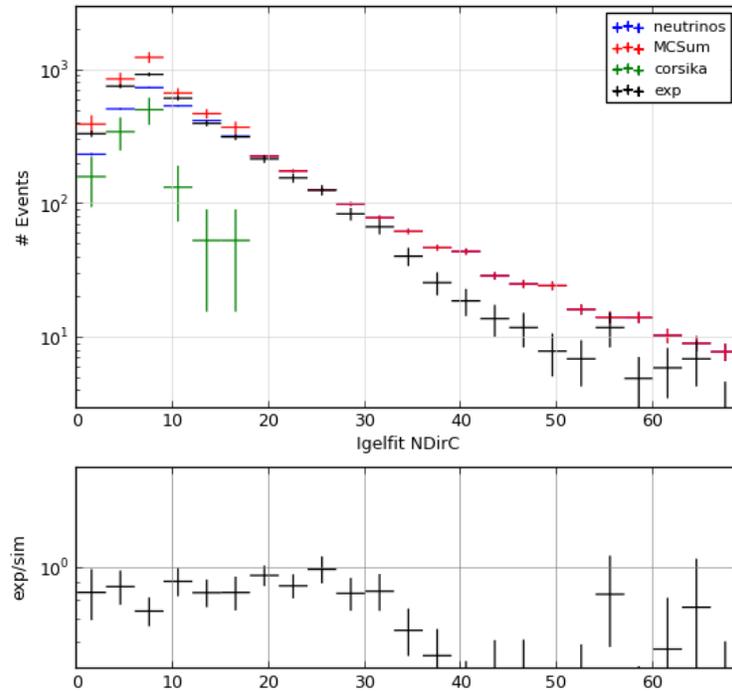
(b) Monopod reconstructed zenith angle versus Monopod energy reconstruction error (c) Monopod reconstructed zenith angle versus Igelfit zenith angle reconstruction error

B.1 Results of Running Igelfit

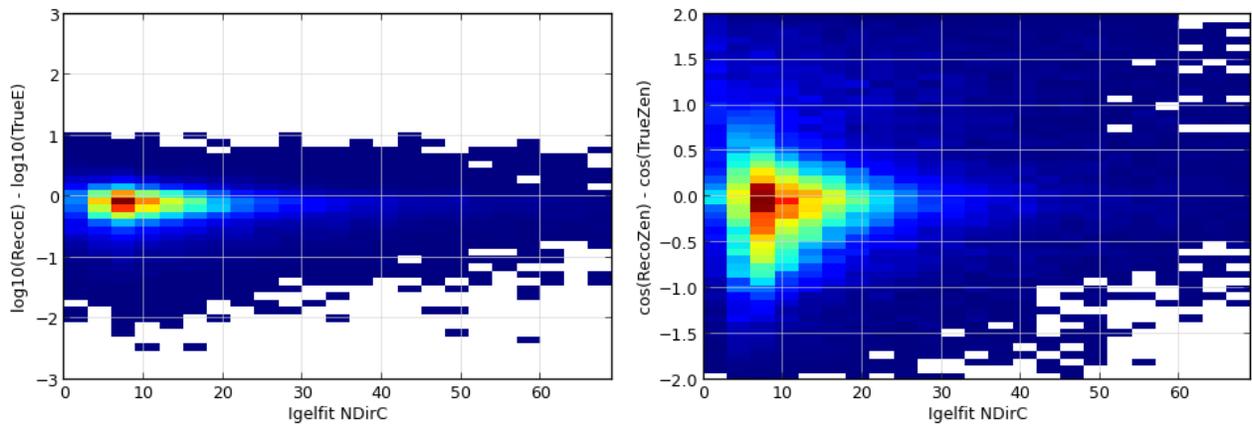
Because no new cuts are applied at L9 (i.e., there are no cuts on Igelfit quality), the plots from the previous section also apply to this cut level. The only plots included in this section show the new variables calculated with Igelfit.

Note that reconstructed depth is left out of these plots because Igelfit does not vary the vertex position away from the seed it is given, thus the depth distribution is the same for both Monopod and Igelfit.

Figure B.13: Number of direct hits relative to the Igelfit reconstruction, at cut level 9. For more information about the Igelfit reconstruction algorithm, see §5.11.3. Direct hits are those that arrive between 15 nsec before the expected light arrival time and 75 nsec after the expected time.



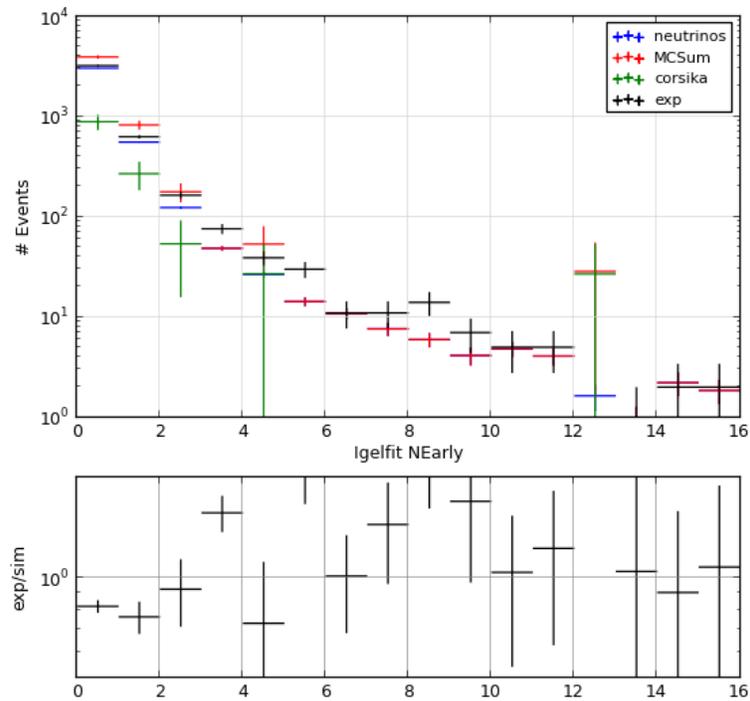
(a) Number of direct hits relative to Igelfit versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



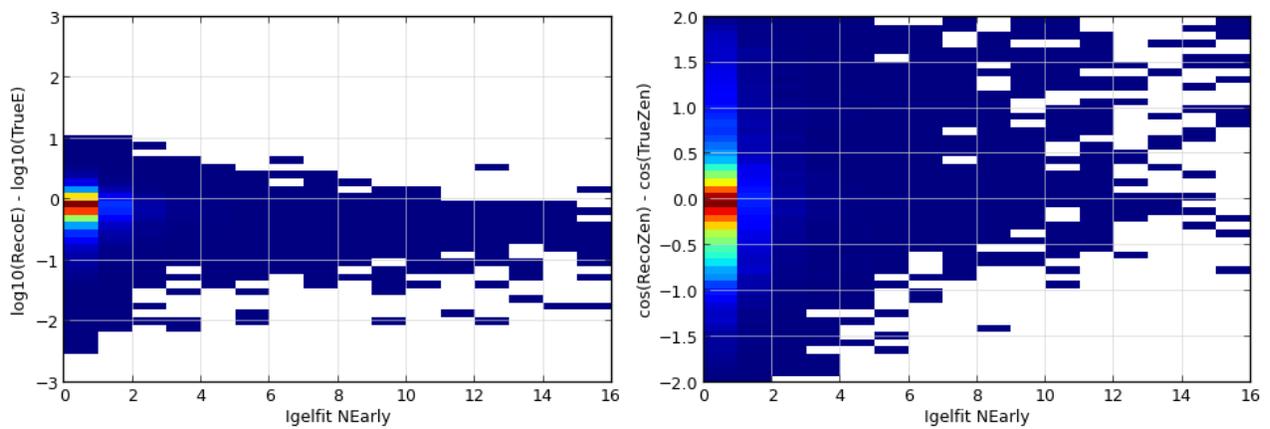
(b) Number of direct hits relative to Igelfit versus Monopod energy reconstruction error

(c) Number of direct hits relative to Igelfit versus Igelfit zenith angle reconstruction error

Figure B.14: Number of early pulses relative to the Igelfit reconstruction, at cut level 8. For more information about the Monopod reconstruction algorithm, see §5.11.1. Early pulses are those that arrive 15 nsec or more before the expected light arrival time.

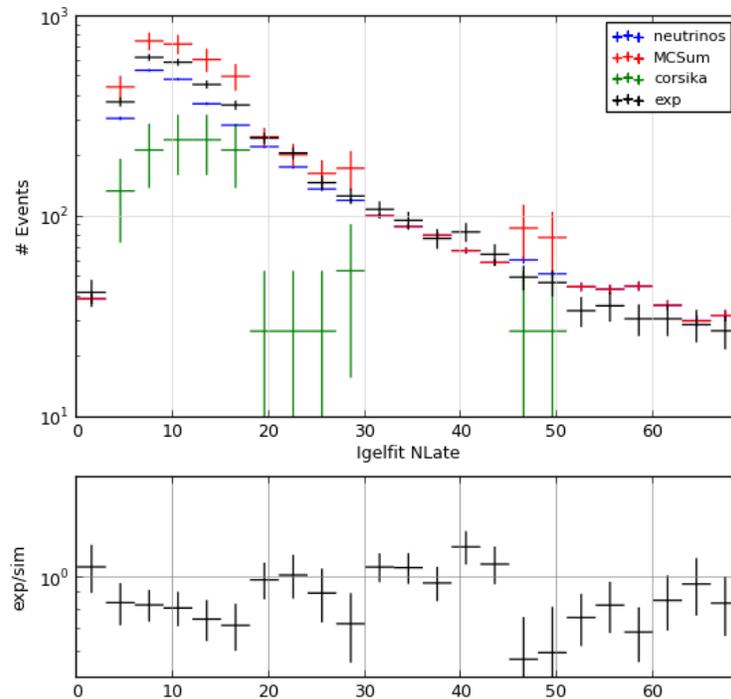


(a) Number of early pulses relative to the Igelfit reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

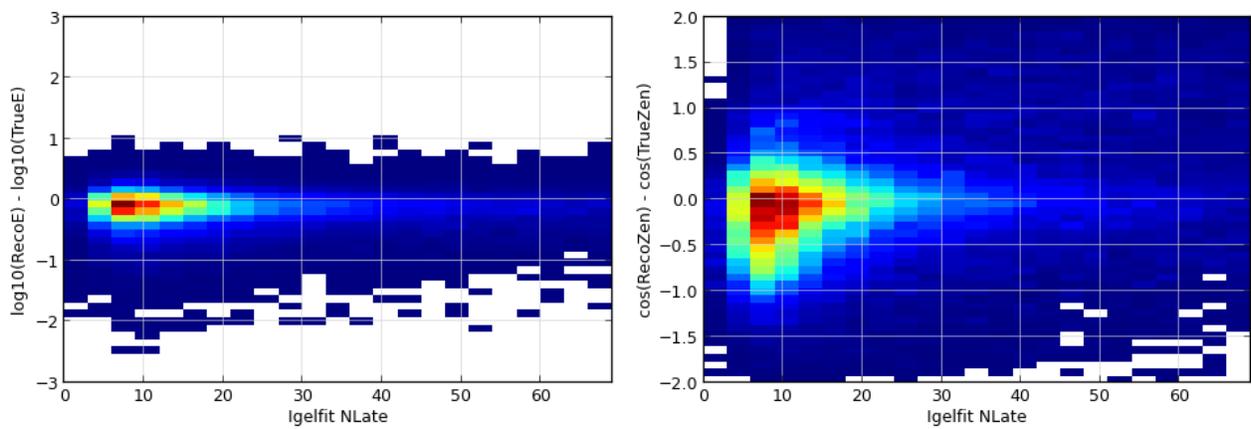


(b) Number of early pulses relative to the Igelfit reconstruction versus Monopod energy reconstruction error (c) Number of early pulses relative to the Igelfit reconstruction versus Igelfit zenith angle reconstruction error

Figure B.15: Number of late pulses relative to the Igelfit reconstruction, at cut level 8. For more information about the Igelfit reconstruction algorithm, see §5.11.3. Late pulses are those that arrive 250 nsec or more after the expected light arrival time.

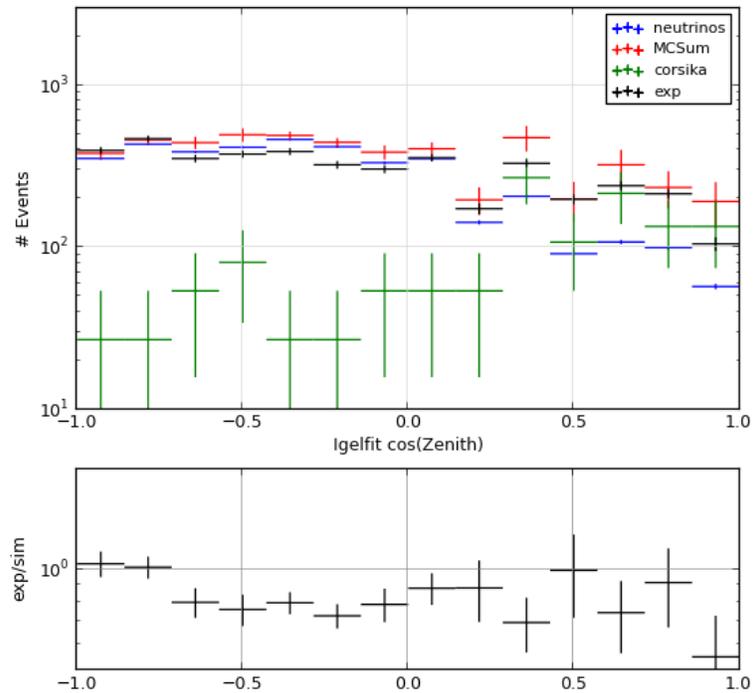


(a) Number of late pulses relative to the Igelfit reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

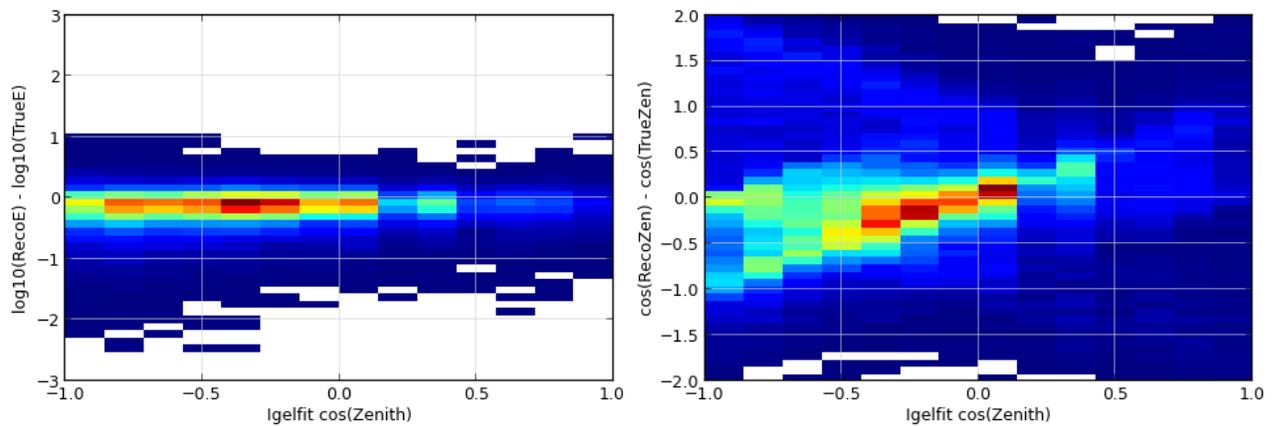


(b) Number of late pulses relative to the Igelfit reconstruction versus Monopod energy reconstruction error (c) Number of late pulses relative to the Igelfit reconstruction versus Igelfit zenith angle reconstruction error

Figure B.16: Reconstructed Zenith, using the Igelfit reconstruction algorithm, at cut level 8. For more information about the Igelfit reconstruction algorithm, see §5.11.3.



(a) Igelfit reconstructed zenith angle versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

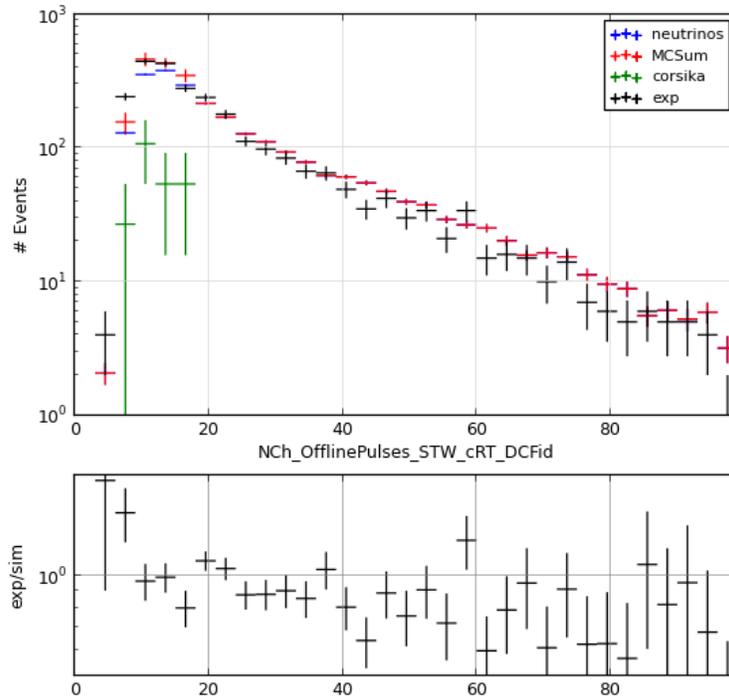


(b) Igelfit reconstructed zenith angle versus Monopod energy reconstruction error (c) Igelfit reconstructed zenith angle versus Igelfit zenith angle reconstruction error

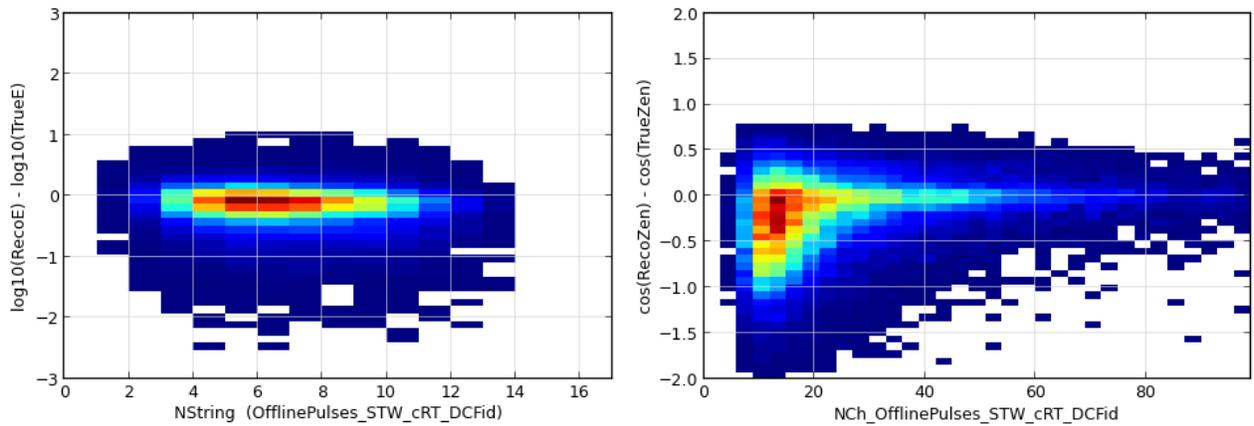
Appendix C: Supplemental Plots for Level 10 Cuts

The final set of cuts is a simple cut on Monopod reconstructed energy and Igelfit reconstructed zenith. The events passing these cuts are used for the final analysis. The full IC79 dataset is used.

Figure C.1: NCh, after hit cleaning, at cut level 10. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

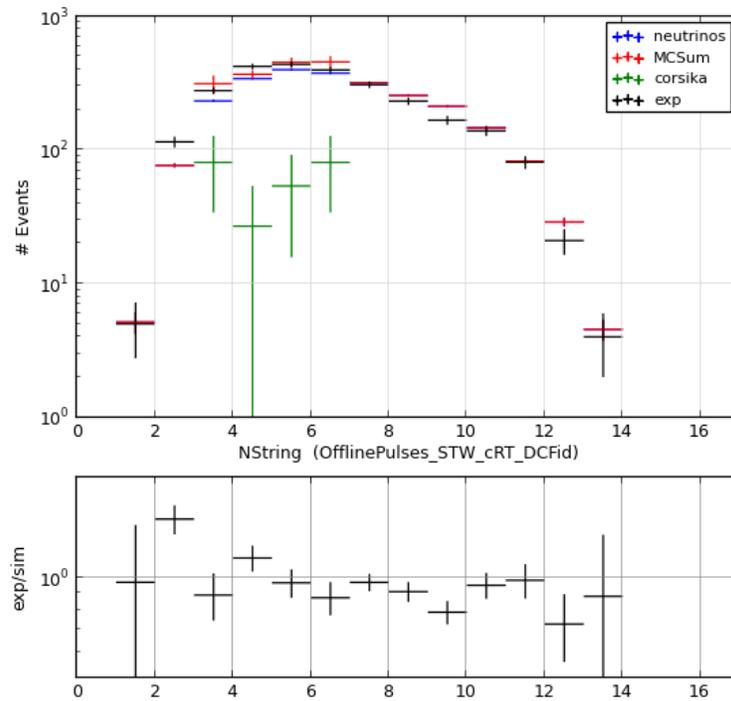


(a) NCh after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

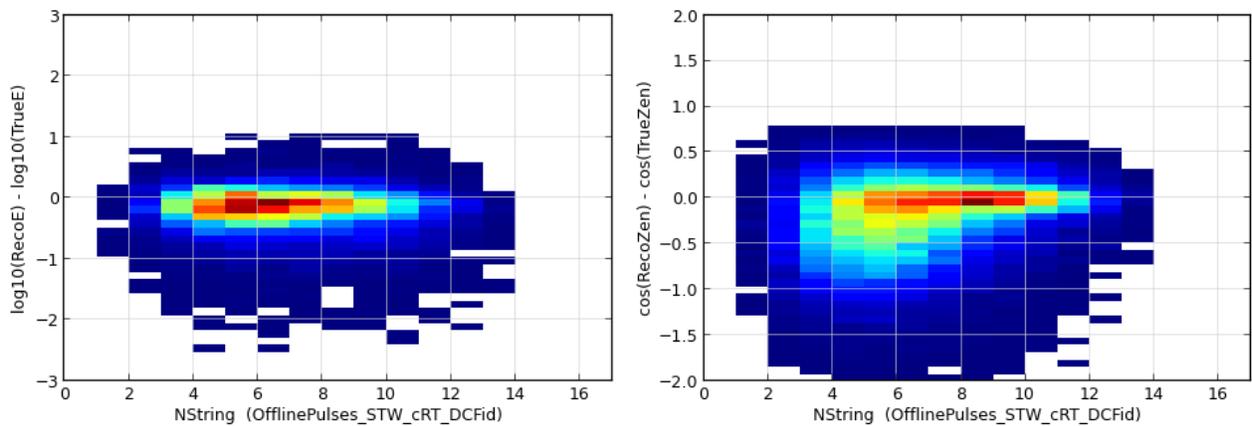


(b) NCh after hit cleaning versus Monopod energy reconstruction error (c) NCh after hit cleaning versus Igelfit energy reconstruction error

Figure C.2: NString, after hit cleaning, at cut level 10. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

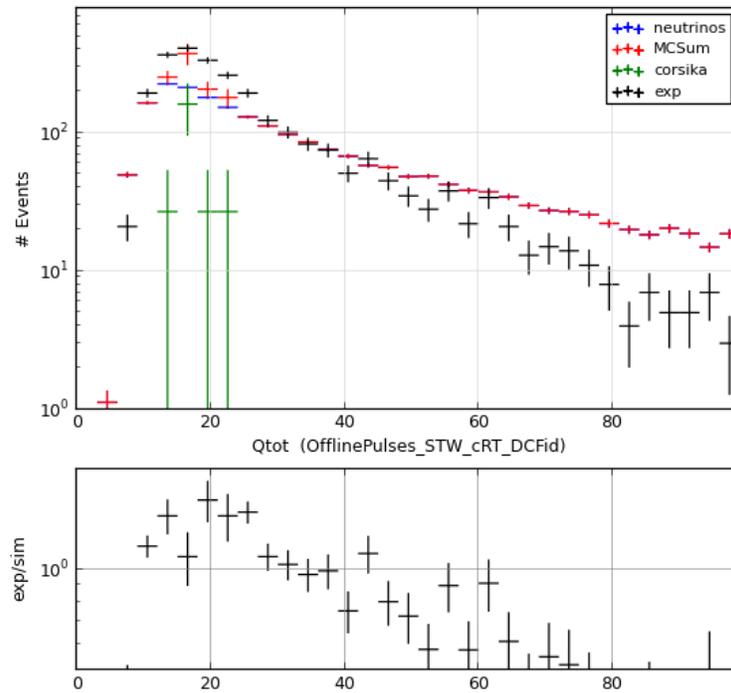


(a) NString after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

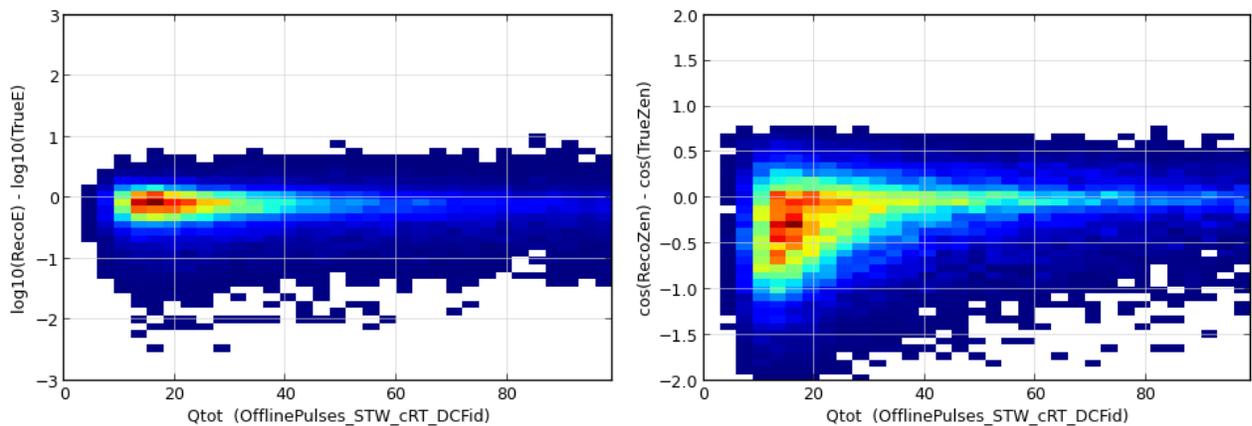


(b) NString after hit cleaning versus Monopod energy reconstruction error (c) NString after hit cleaning versus Igelfit zenith angle reconstruction error

Figure C.3: Total charge, after hit cleaning, at cut level 10. The hit cleaning used here is a time window cleaning, followed by a seeded RT cleaning, then only counting the DOMs in the DeepCore fiducial region.

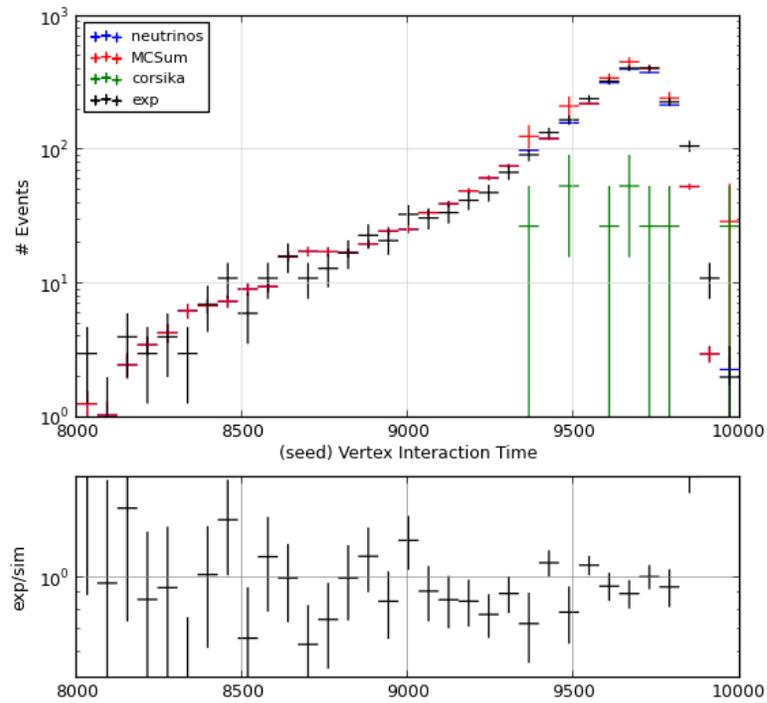


(a) Total charge after hit cleaning versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

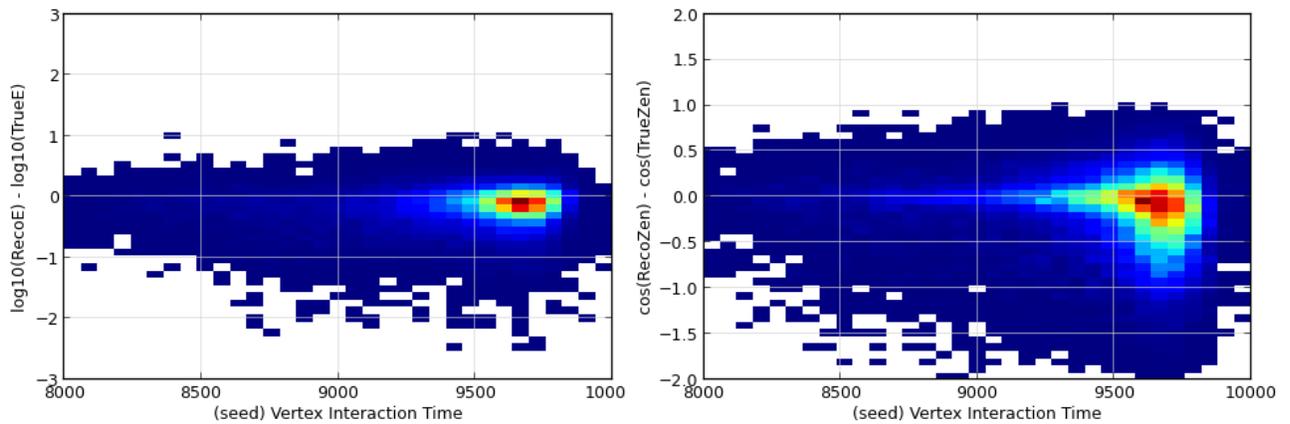


(b) Total charge after hit cleaning versus Monopod energy reconstruction error (c) Total charge after hit cleaning versus Igefit zenith angle reconstruction error

Figure C.4: CLast reconstructed vertex time at cut level 10. For an explanation of how the CLast vertex time is calculated, see §5.7.

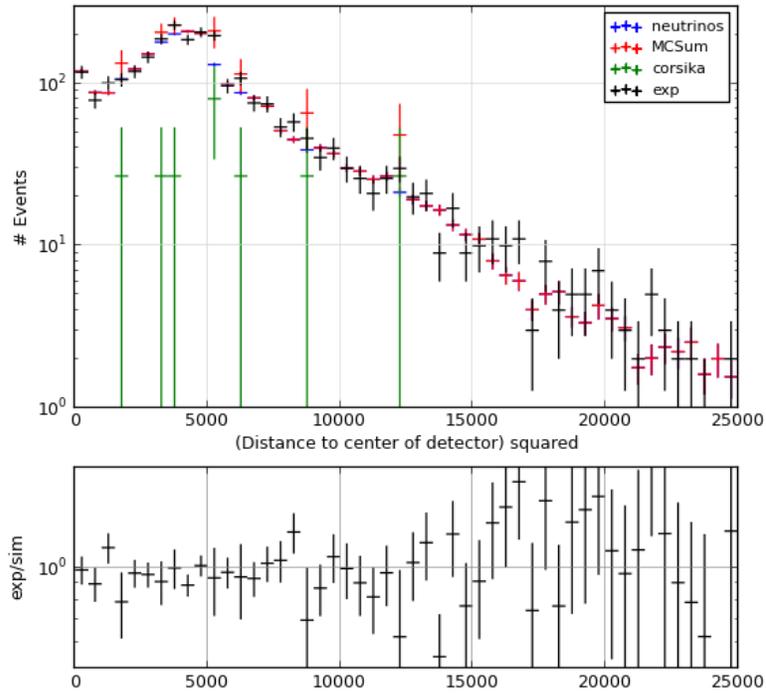


(a) CLast reconstructed vertex time versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

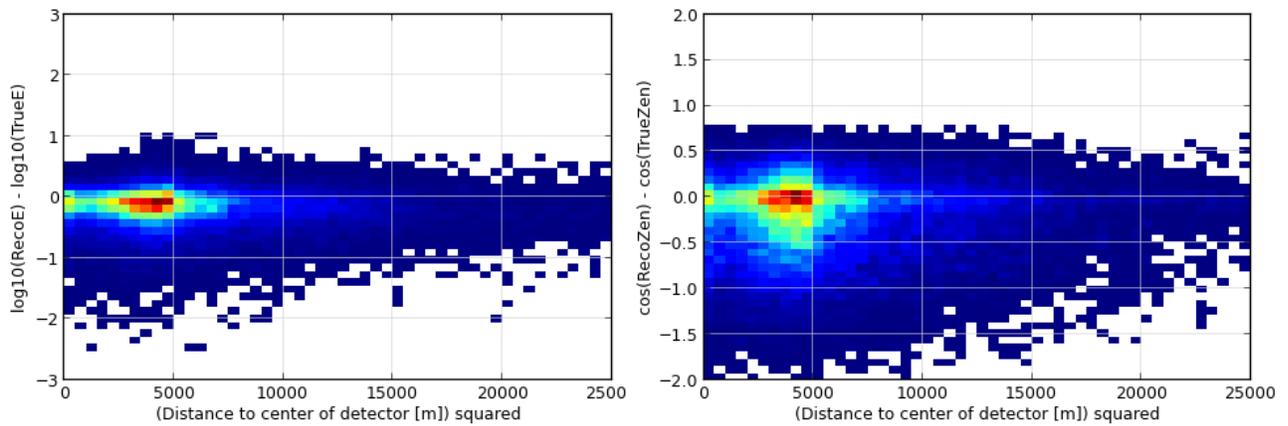


(b) CLast reconstructed vertex time versus Monopod energy reconstruction error (c) CLast reconstructed vertex time versus Igelfit zenith angle reconstruction error

Figure C.5: Distance between the central string and center-of-gravity reconstructed vertex, squared, at cut level 10

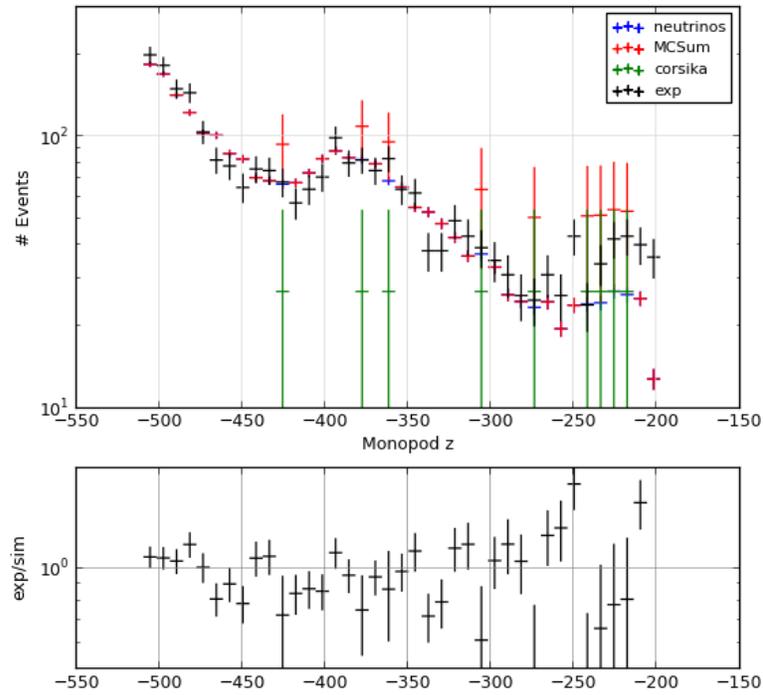


(a) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

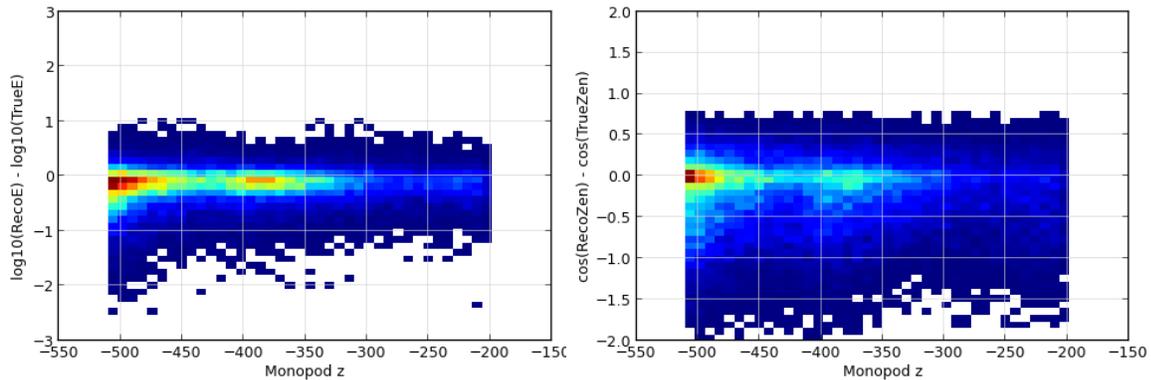


(b) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus Monopod energy reconstruction error
 (c) Distance between the central string and center-of-gravity reconstructed vertex, squared, versus IgefIt zenith angle reconstruction error

Figure C.6: Monopod reconstructed depth at cut level 10

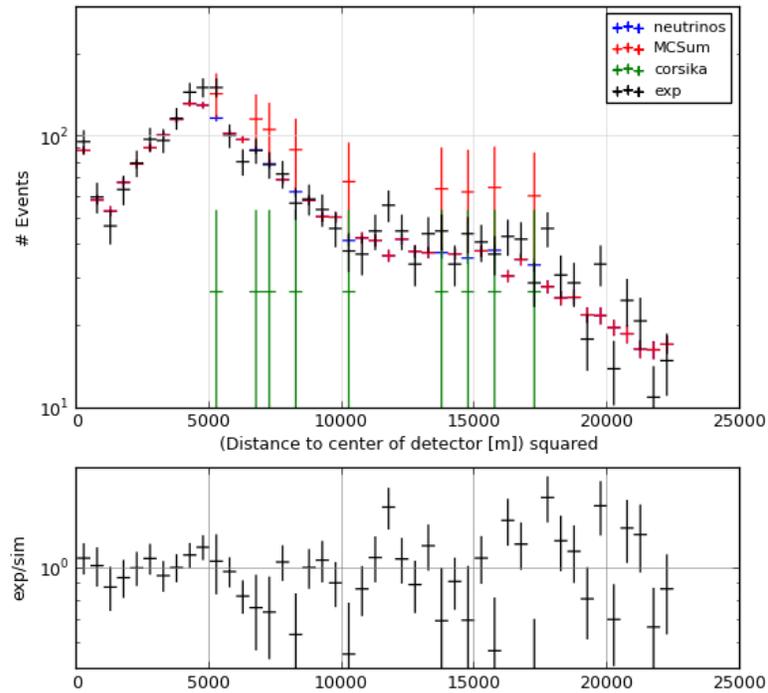


(a) Monopod reconstructed depth versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

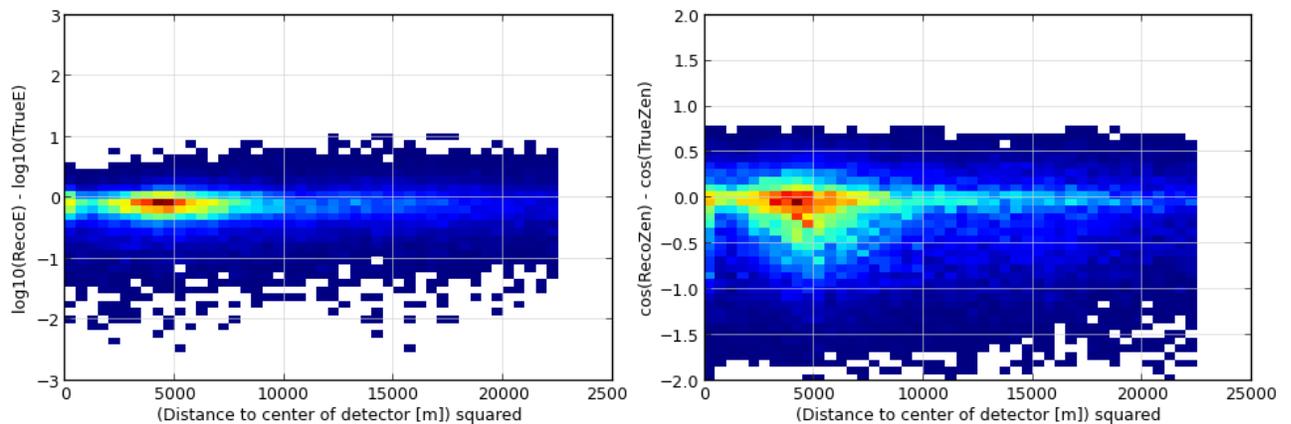


(b) Monopod reconstructed depth versus Monopod energy (c) Monopod reconstructed depth versus Igefit zenith angle reconstruction error

Figure C.7: Distance between the central string and the Monopod reconstructed vertex, squared, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1.

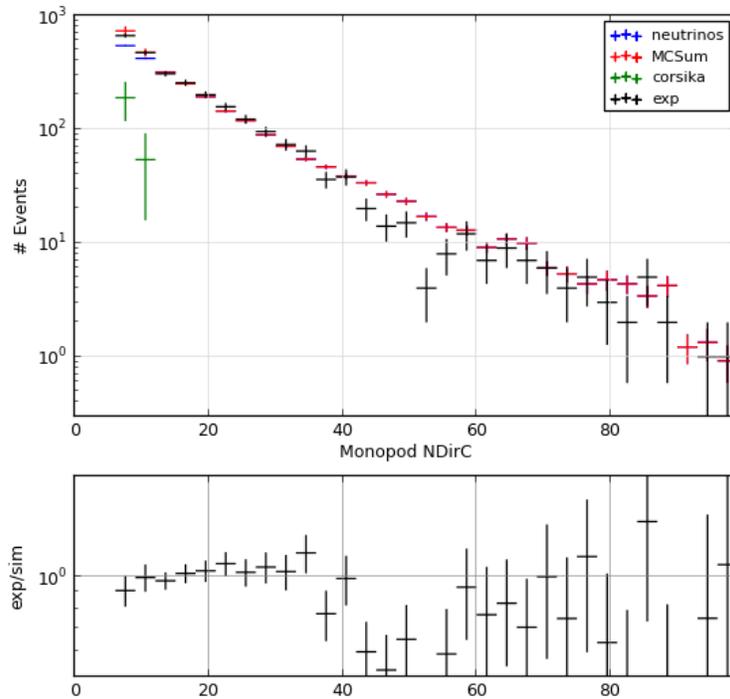


(a) Distance between the central string and the Monopod reconstructed vertex, squared, versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

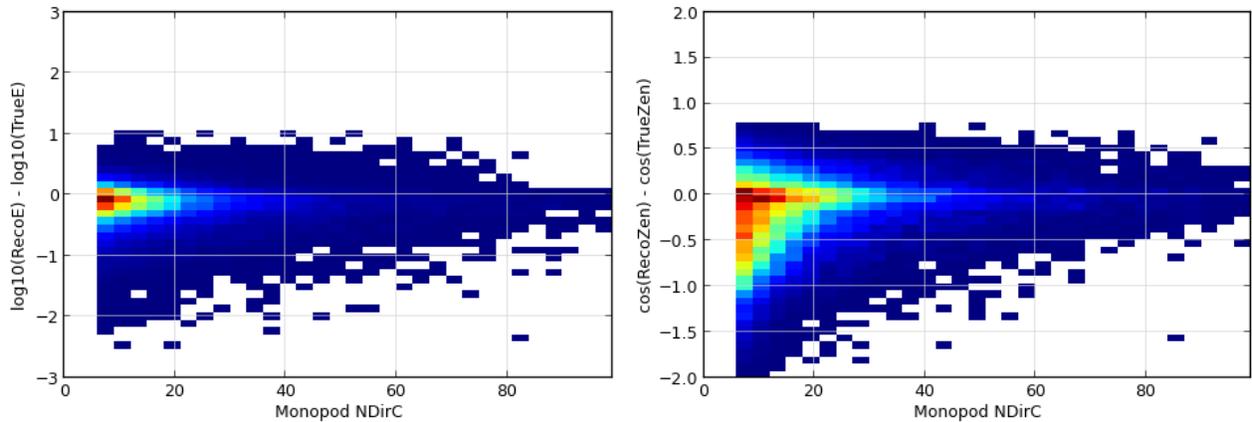


(b) Distance between the central string and the Monopod reconstructed vertex, squared, versus Monopod energy reconstruction error (c) Distance between the central string and the Monopod reconstructed vertex, squared, versus Igefit zenith angle reconstruction error

Figure C.8: Number of direct hits relative to the Monopod reconstructed interaction vertex, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1. Direct hits are those that arrive between 15 nsec before the expected light arrival time and 75 nsec after the expected time.

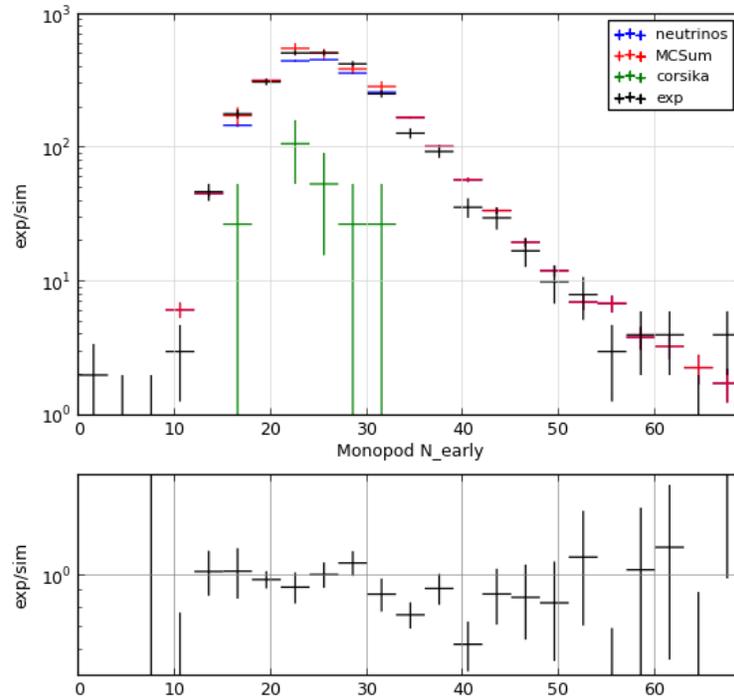


(a) Number of direct hits relative to Monopod versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

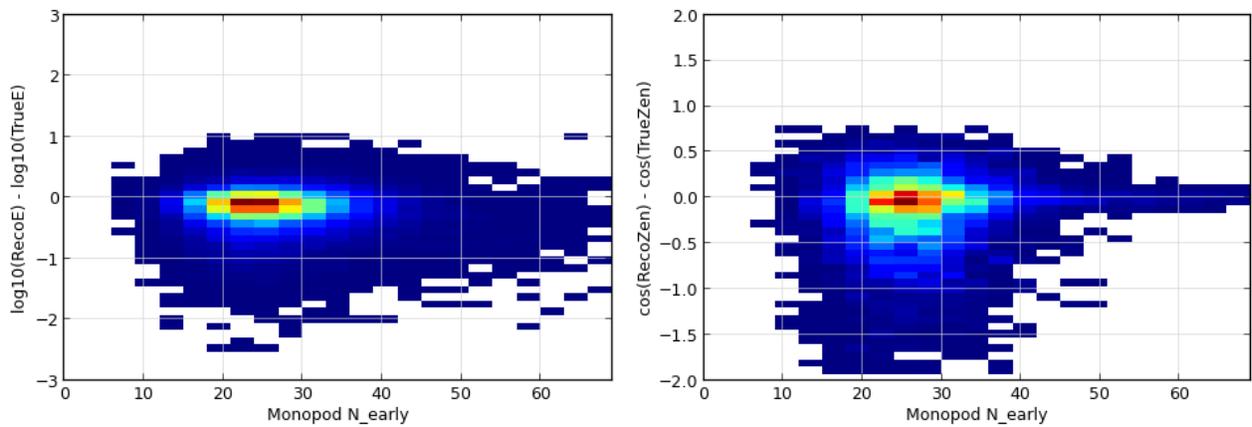


(b) Number of direct hits relative to Monopod versus Monopod energy reconstruction error (c) Number of direct hits relative to Monopod versus Igefilt zenith angle reconstruction error

Figure C.9: Number of early pulses relative to the Monopod reconstructed interaction vertex, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1. Early pulses are those that arrive 15 nsec or more before the expected light arrival time.

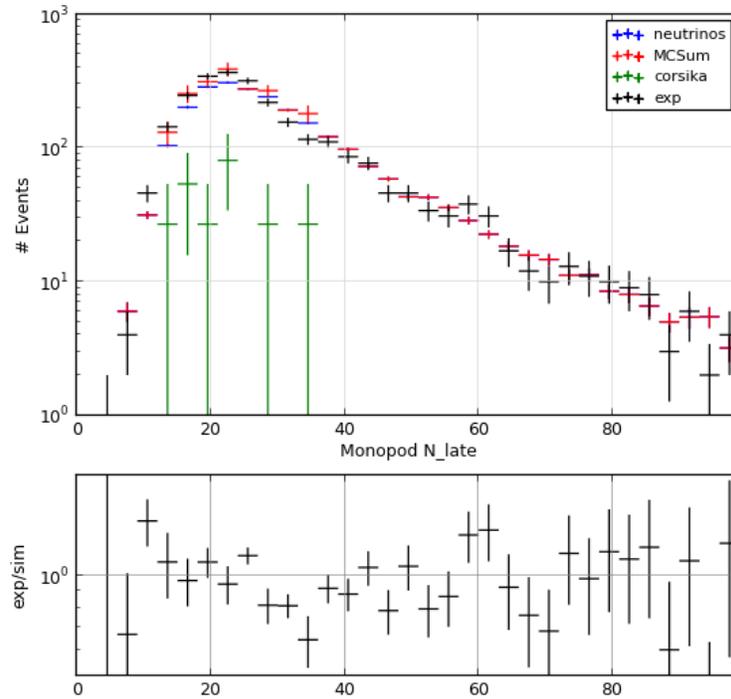


(a) Number of early pulses relative to the Monopod reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

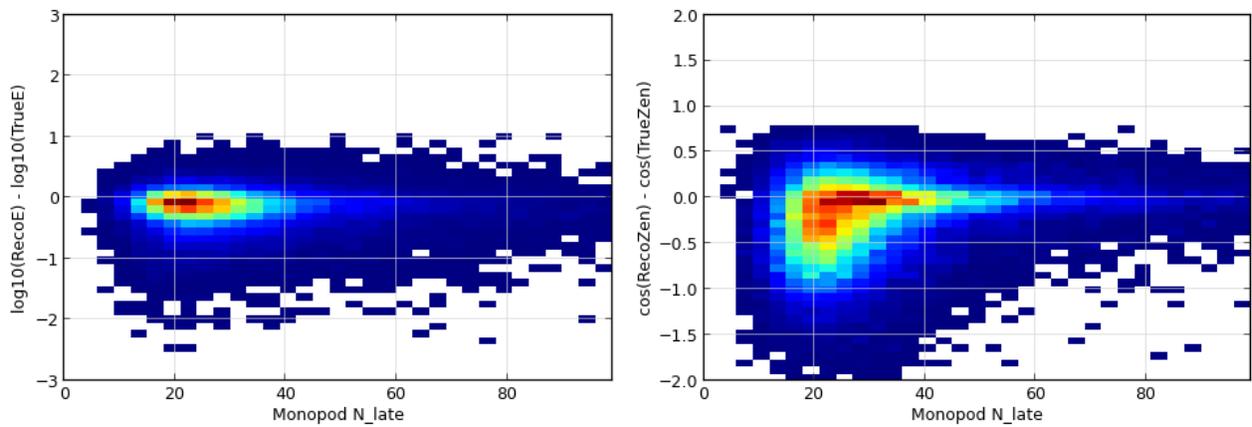


(b) Number of early pulses relative to the Monopod reconstructed interaction vertex versus Monopod energy reconstruction error (c) Number of early pulses relative to the Monopod reconstructed interaction vertex versus Igefit zenith angle reconstruction error

Figure C.10: Number of late pulses relative to the Monopod reconstructed interaction vertex, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1. Late pulses are those that arrive 250 nsec or more after the expected light arrival time.

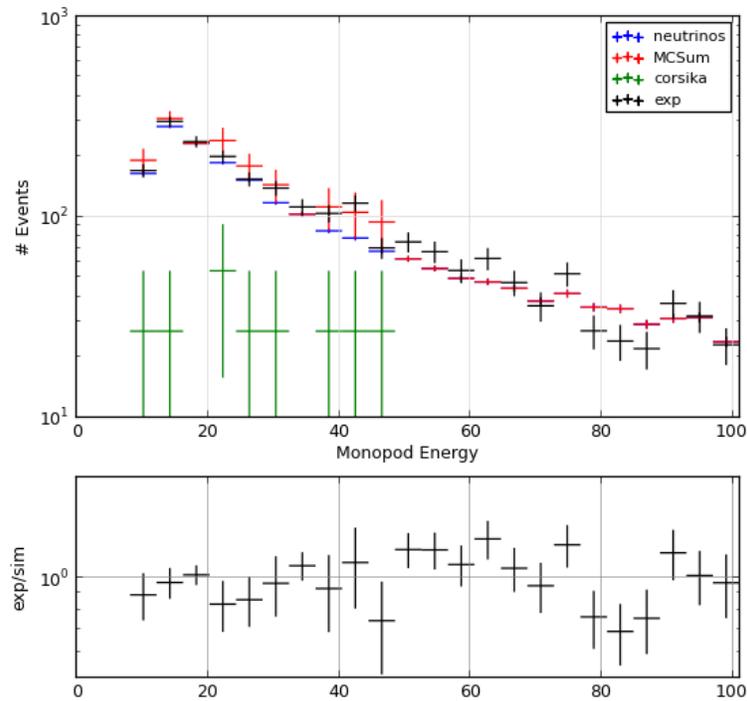


(a) Number of late pulses relative to the Monopod reconstructed interaction vertex versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

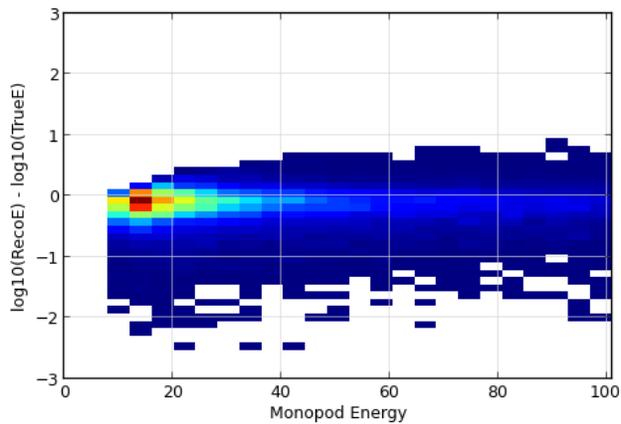


(b) Number of late pulses relative to the Monopod reconstructed interaction vertex versus Monopod energy reconstruction error (c) Number of late pulses relative to the Monopod reconstructed interaction vertex versus Igefit zenith angle reconstruction error

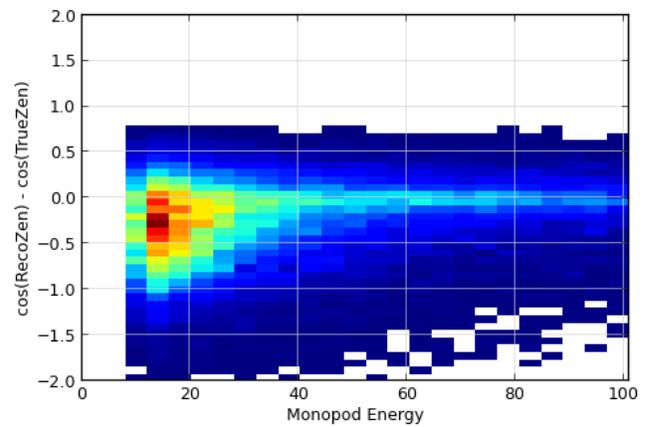
Figure C.11: Reconstructed energy, using the Monopod reconstruction algorithm, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1.



(a) Reconstructed Energy versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

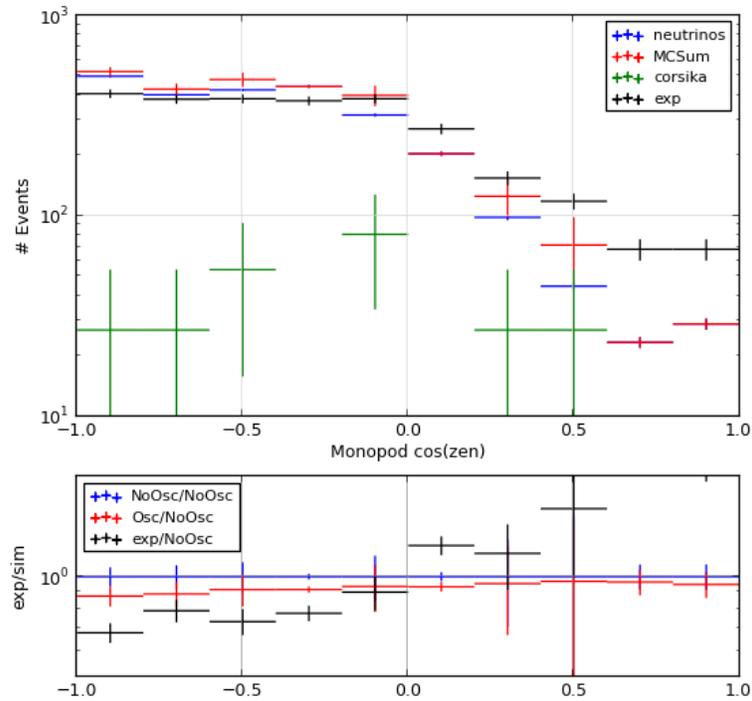


(b) Reconstructed Energy versus Monopod energy reconstruction error

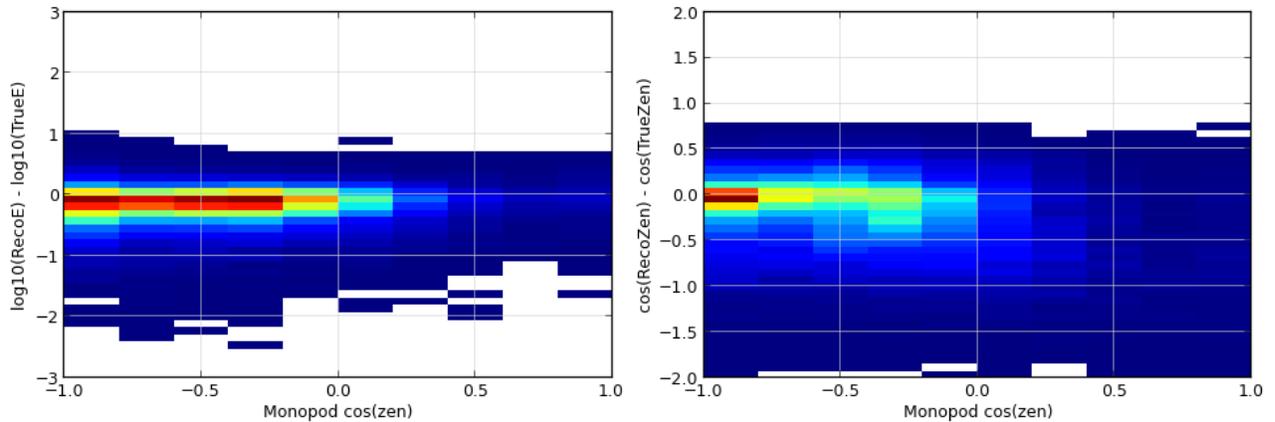


(c) Reconstructed Energy versus Igelfit reconstruction error

Figure C.12: Reconstructed Zenith, using the Monopod reconstruction algorithm, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1. Monopod is not expected to be an accurate reconstruction for zenith angle, but the plot is included here for comparison to the Igefifit reconstructed zenith.

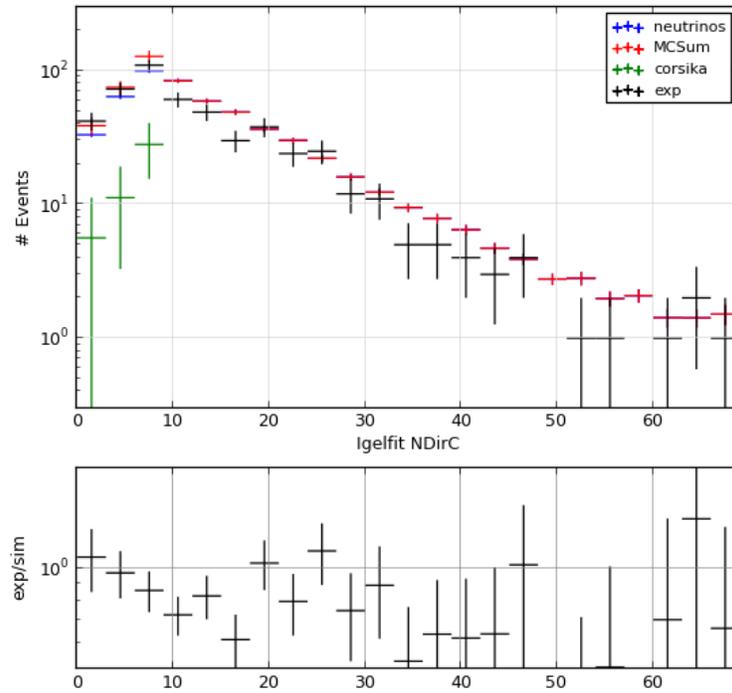


(a) Reconstructed zenith angle versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

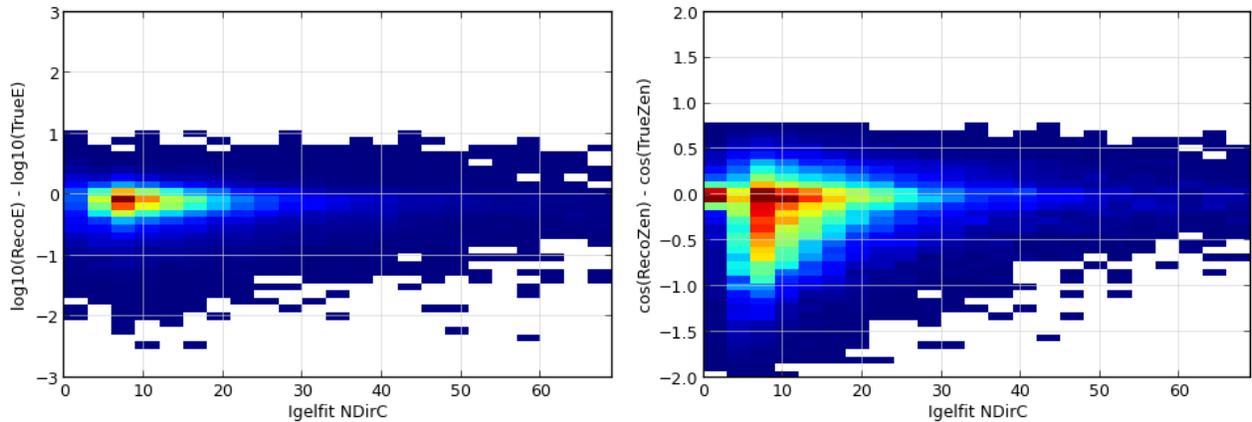


(b) Monopod reconstructed zenith angle versus Monopod energy reconstruction error (c) Monopod reconstructed zenith angle versus Igefifit zenith angle reconstruction error

Figure C.13: Number of direct hits relative to the Igelfit reconstruction, at cut level 10. For more information about the Igelfit reconstruction algorithm, see §5.11.3. Direct hits are those that arrive between 15 nsec before the expected light arrival time and 75 nsec after the expected time.



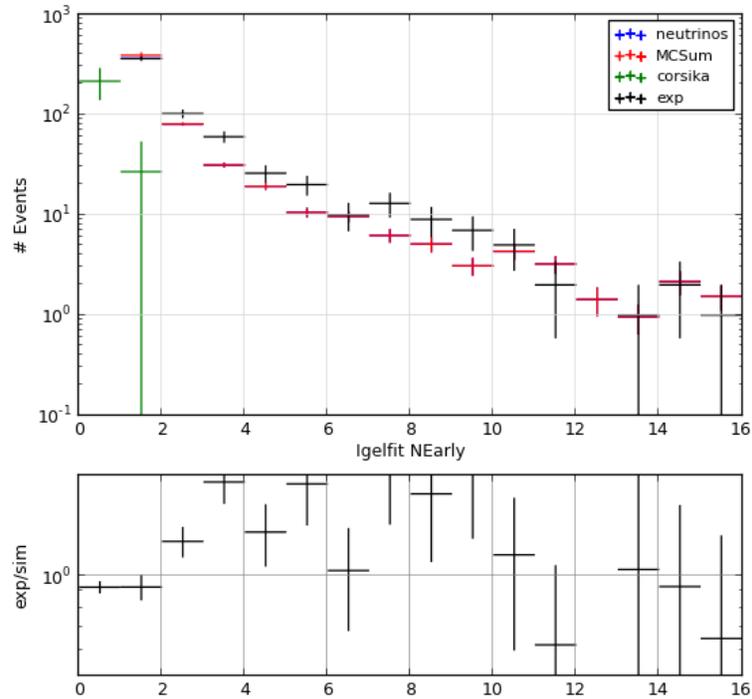
(a) Number of direct hits relative to Igelfit versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



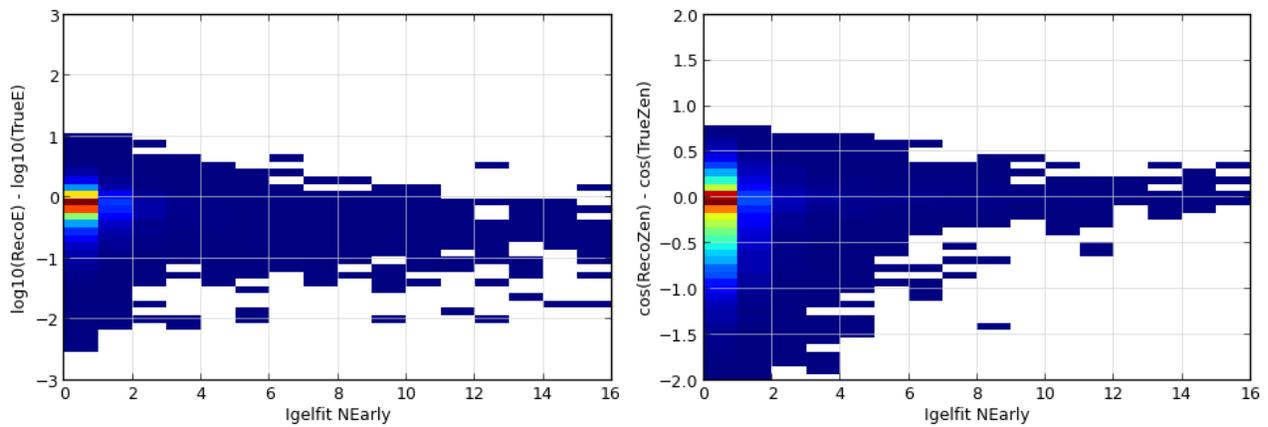
(b) Number of direct hits relative to Igelfit versus Monopod energy reconstruction error

(c) Number of direct hits relative to Igelfit versus Igelfit zenith angle reconstruction error

Figure C.14: Number of early pulses relative to the Igelfit reconstruction, at cut level 10. For more information about the Monopod reconstruction algorithm, see §5.11.1. Early pulses are those that arrive 15 nsec or more before the expected light arrival time.

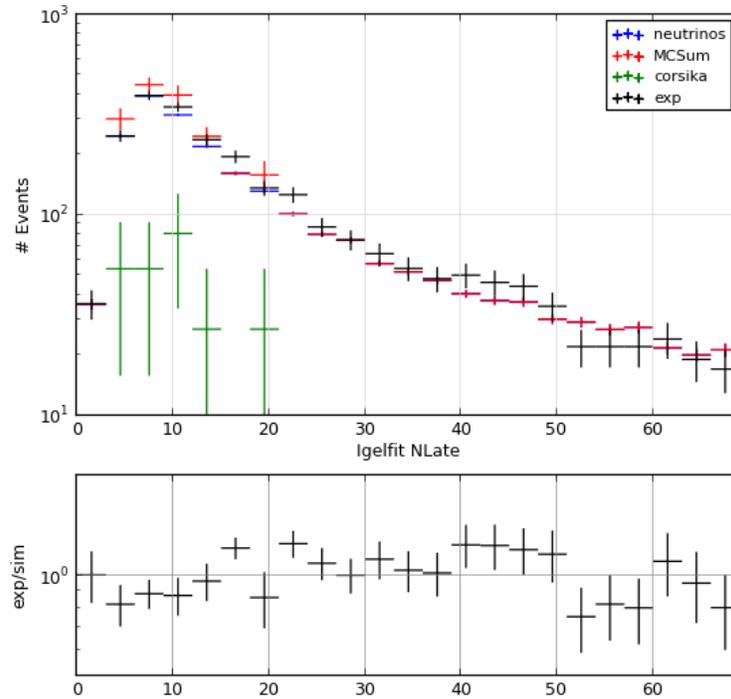


(a) Number of early pulses relative to the Igelfit reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

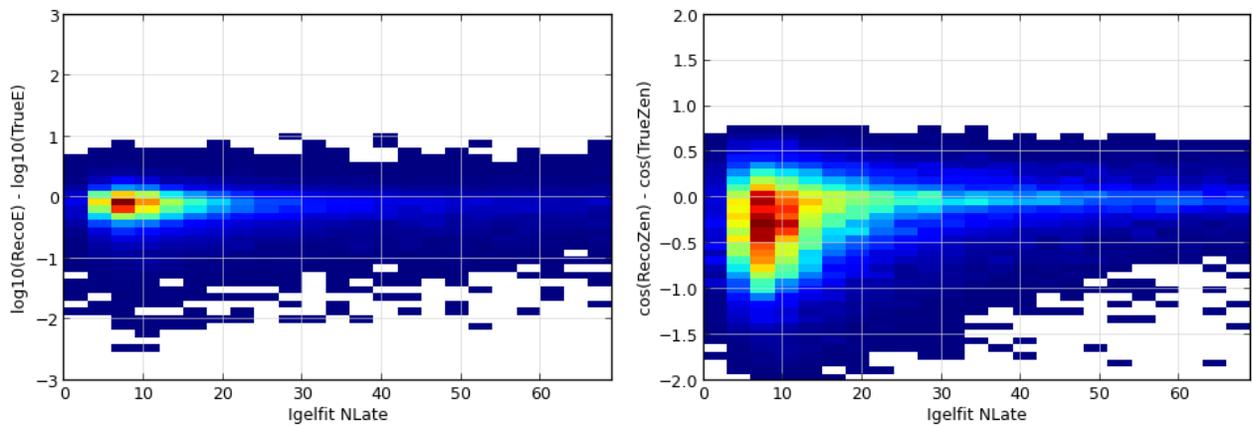


(b) Number of early pulses relative to the Igelfit reconstruction versus Monopod energy reconstruction error (c) Number of early pulses relative to the Igelfit reconstruction versus Igelfit zenith angle reconstruction error

Figure C.15: Number of late pulses relative to the Igelfit reconstruction, at cut level 10. For more information about the Igelfit reconstruction algorithm, see §5.11.3. Late pulses are those that arrive 250 nsec or more after the expected light arrival time.

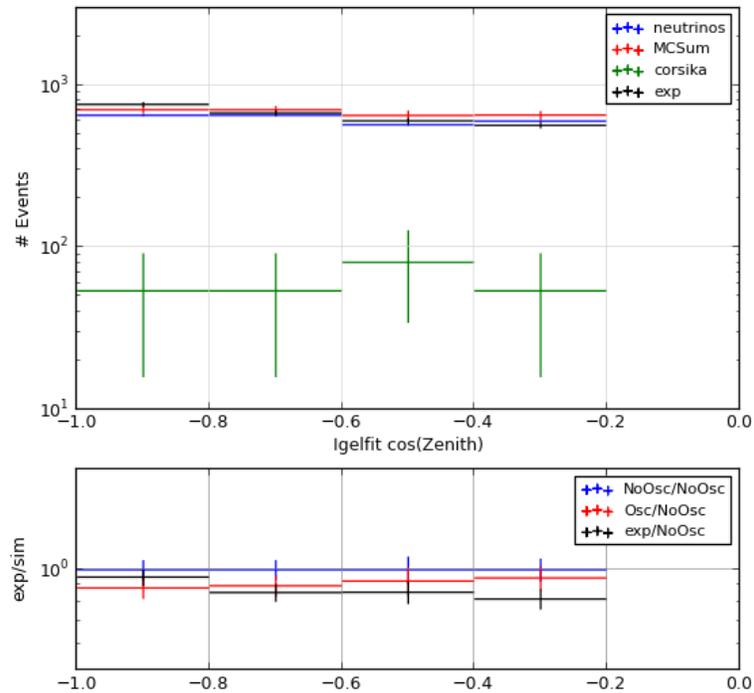


(a) Number of late pulses relative to the Igelfit reconstruction versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).

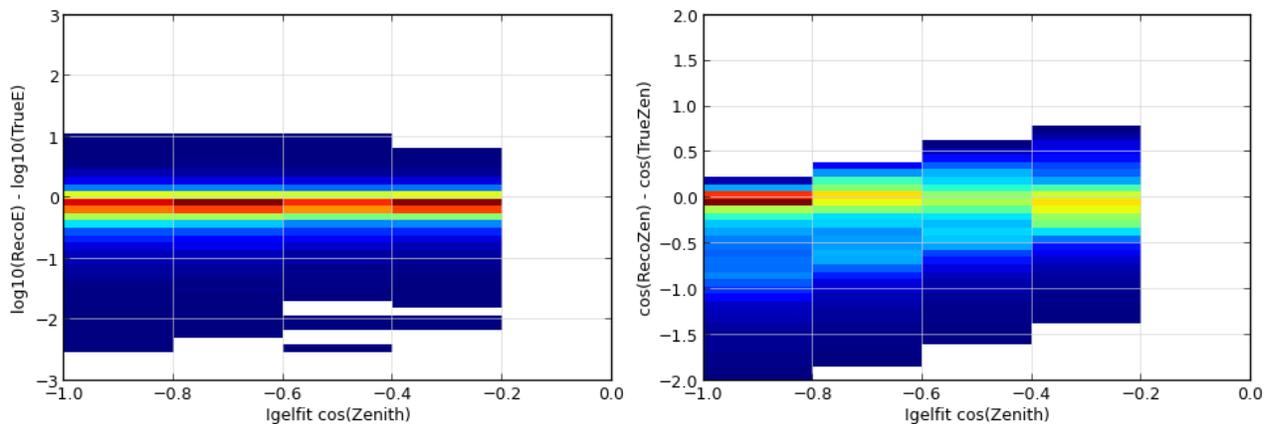


(b) Number of late pulses relative to the Igelfit reconstruction versus Monopod energy reconstruction error (c) Number of late pulses relative to the Igelfit reconstruction versus Igelfit zenith angle reconstruction error

Figure C.16: Reconstructed Zenith, using the Igelfit reconstruction algorithm, at cut level 10. For more information about the Igelfit reconstruction algorithm, see §5.11.3.



(a) Igelfit reconstructed zenith angle versus rate, with simulation compared to data. The simulation sample (red) is the sum of several components: CORSIKA-generated cosmic ray muons (green), and three flavors of neutrinos added together (blue).



(b) Igelfit reconstructed zenith angle versus Monopod energy reconstruction error (c) Igelfit reconstructed zenith angle versus Igelfit zenith angle reconstruction error

Appendix D: The IceCube Moon Shadow Paper

For more information about the IceCube moon shadow, see the following paper and also Chapter 6.

Observation of the cosmic-ray shadow of the Moon with IceCube

M. G. Aartsen,² R. Abbasi,²⁷ Y. Abdou,²² M. Ackermann,⁴¹ J. Adams,¹⁵ J. A. Aguilar,²¹ M. Ahlers,²⁷ D. Altman,⁹ J. Auffenberg,²⁷ X. Bai,^{31,‡} M. Baker,²⁷ S. W. Barwick,²³ V. Baum,²⁸ R. Bay,⁷ J. J. Beatty,^{17,18} S. Bechet,¹² J. Becker Tjus,¹⁰ K.-H. Becker,⁴⁰ M. Bell,³⁸ M. L. Benabderrahmane,⁴¹ S. BenZvi,²⁷ J. Berdermann,⁴¹ P. Berghaus,⁴¹ D. Berley,¹⁶ E. Bernardini,⁴¹ A. Bernhard,³⁰ D. Bertrand,¹² D. Z. Besson,²⁵ G. Binder,^{8,7} D. Bindig,⁴⁰ M. Bissok,¹ E. Blaufuss,¹⁶ J. Blumenthal,¹ D. J. Boersma,^{39,*} S. Bohaichuk,²⁰ C. Boehm,³⁴ D. Bose,¹³ S. Böser,¹¹ O. Botner,³⁹ L. Brayeur,¹³ H.-P. Bretz,⁴¹ A. M. Brown,¹⁵ R. Bruijn,²⁴ J. Brunner,⁴¹ M. Carson,²² J. Casey,⁵ M. Casier,¹³ D. Chirkin,²⁷ A. Christov,²¹ B. Christy,¹⁶ K. Clark,³⁸ F. Clevermann,¹⁹ S. Coenders,¹ S. Cohen,²⁴ D. F. Cowen,^{38,37} A. H. Cruz Silva,⁴¹ M. Danninger,³⁴ J. Daughhete,⁵ J. C. Davis,¹⁷ C. De Clercq,¹³ S. De Ridder,²² P. Desiati,²⁷ M. de With,⁹ T. DeYoung,³⁸ J. C. Díaz-Vélez,²⁷ M. Dunkman,³⁸ R. Egan,³⁸ B. Eberhardt,²⁸ J. Eisch,²⁷ R. W. Ellsworth,¹⁶ S. Euler,¹ P. A. Evenson,³¹ O. Fadiran,²⁷ A. R. Fazely,⁶ A. Fedynitch,¹⁰ J. Feintzeig,²⁷ T. Feusels,²² K. Filimonov,⁷ C. Finley,³⁴ T. Fischer-Wasels,⁴⁰ S. Flis,³⁴ A. Franckowiak,¹¹ R. Franke,⁴¹ K. Frantzen,¹⁹ T. Fuchs,¹⁹ T. K. Gaisser,³¹ J. Gallagher,²⁶ L. Gerhardt,^{8,7} L. Gladstone,²⁷ T. Glusenkamp,⁴¹ A. Goldschmidt,⁸ G. Golup,¹³ J. G. Gonzalez,³¹ J. A. Goodman,¹⁶ D. Góra,⁴¹ D. T. Grandmont,²⁰ D. Grant,²⁰ A. Groß,³⁰ C. Ha,^{8,7} A. Haj Ismail,²² P. Hallen,¹ A. Hallgren,³⁹ F. Halzen,²⁷ K. Hanson,¹² D. Heereman,¹² D. Heinen,¹ K. Helbing,⁴⁰ R. Hellauer,¹⁶ S. Hickford,¹⁵ G. C. Hill,² K. D. Hoffman,¹⁶ R. Hoffmann,⁴⁰ A. Homeier,¹¹ K. Hoshina,²⁷ W. Huelsnitz,^{16,§} P. O. Hulth,³⁴ K. Hultqvist,³⁴ S. Hussain,³¹ A. Ishihara,¹⁴ E. Jacobi,⁴¹ J. Jacobsen,²⁷ K. Jagielski,¹ G. S. Japaridze,⁴ K. Jero,²⁷ O. Jlelati,²² B. Kaminsky,⁴¹ A. Kappes,⁹ T. Karg,⁴¹ A. Karle,²⁷ J. L. Kelley,²⁷ J. Kiryluk,³⁵ F. Kislat,⁴¹ J. Kläs,⁴⁰ S. R. Klein,^{8,7} J.-H. Köhne,¹⁹ G. Kohnen,²⁹ H. Kolanoski,⁹ L. Köpke,²⁸ C. Kopper,²⁷ S. Kopper,⁴⁰ D. J. Koskinen,³⁸ M. Kowalski,¹¹ M. Krasberg,²⁷ K. Krings,¹ G. Kroll,²⁸ J. Kunnen,¹³ N. Kurahashi,²⁷ T. Kuwabara,³¹ M. Labare,¹³ H. Landsman,²⁷ M. J. Larson,³⁶ M. Lesiak-Bzdak,³⁵ M. Leuermann,¹ J. Leute,³⁰ J. Lünemann,²⁸ J. Madsen,³³ R. Maruyama,²⁷ K. Mase,¹⁴ H. S. Matis,⁸ F. McNally,²⁷ K. Meagher,¹⁶ M. Merck,²⁷ P. Mészáros,^{37,38} T. Meures,¹² S. Miarecki,^{8,7} E. Middell,⁴¹ N. Milke,¹⁹ J. Miller,¹³ L. Mohrmann,⁴¹ T. Montaruli,^{21,||} R. Morse,²⁷ R. Nahnauer,⁴¹ U. Naumann,⁴⁰ H. Niederhausen,³⁵ S. C. Nowicki,²⁰ D. R. Nygren,⁸ A. Obertacke,⁴⁰ S. Odrowski,³⁰ A. Olivas,¹⁶ M. Olivo,¹⁰ A. O'Murchadha,¹² L. Paul,¹ J. A. Pepper,³⁶ C. Pérez de los Heros,³⁹ C. Pfendner,¹⁷ D. Pieloth,¹⁹ E. Pinat,¹² N. Pirk,⁴¹ J. Posselt,⁴⁰ P. B. Price,⁷ G. T. Przybylski,⁸ L. Rädcl,¹ M. Rameez,²¹ K. Rawlins,³ P. Redl,¹⁶ R. Reimann,¹ E. Resconi,³⁰ W. Rhode,¹⁹ M. Ribordy,²⁴ M. Richman,¹⁶ B. Riedel,²⁷ J. P. Rodrigues,²⁷ C. Rott,^{17,¶} T. Ruhe,¹⁹ B. Ruzybayev,³¹ D. Ryckbosch,²² S. M. Saba,¹⁰ T. Salameh,³⁸ H.-G. Sander,²⁸ M. Santander,^{27,†} S. Sarkar,³² K. Schatto,²⁸ M. Scheel,¹ F. Scheriau,¹⁹ T. Schmidt,¹⁶ M. Schmitz,¹⁹ S. Schoenen,¹ S. Schöneberg,¹⁰ A. Schönwald,⁴¹ A. Schukraft,¹ L. Schulte,¹¹ O. Schulz,³⁰ D. Seckel,³¹ Y. Sestayo,³⁰ S. Seunarine,³³ C. Sheremata,²⁰ M. W. E. Smith,³⁸ D. Soldin,⁴⁰ G. M. Spiczak,³³ C. Spiering,⁴¹ M. Stamatikos,^{17,**} T. Stanev,³¹ A. Stasik,¹¹ T. Stezelberger,⁸ R. G. Stokstad,⁸ A. Stöbl,⁴¹ E. A. Strahler,¹³ R. Ström,³⁹ G. W. Sullivan,¹⁶ H. Taavola,³⁹ I. Taboada,⁵ A. Tamburro,³¹ A. Tepe,⁴⁰ S. Ter-Antonyan,⁶ G. Tešić,³⁸ S. Tilav,³¹ P. A. Toale,³⁶ S. Toscano,²⁷ M. Usner,¹¹ D. van der Drift,^{8,7} N. van Eijndhoven,¹³ A. Van Overloop,²² J. van Santen,²⁷ M. Vehringer,¹ M. Voge,¹¹ M. Vraeghe,²² C. Walck,³⁴ T. Waldenmaier,⁹ M. Wallraff,¹ R. Wasserman,³⁸ Ch. Weaver,²⁷ M. Wellons,²⁷ C. Wendt,²⁷ S. Westerhoff,²⁷ N. Whitehorn,²⁷ K. Wiebe,²⁸ C. H. Wiebusch,¹ D. R. Williams,³⁶ H. Wissing,¹⁶ M. Wolf,³⁴ T. R. Wood,²⁰ K. Woschnagg,⁷ C. Xu,³¹ D. L. Xu,³⁶ X. W. Xu,⁶ J. P. Yanez,⁴¹ G. Yodh,²³ S. Yoshida,¹⁴ P. Zarzhitsky,³⁶ J. Ziemann,¹⁹ S. Zierke,¹ and M. Zoll³⁴

(IceCube Collaboration)

¹*III. Physikalisches Institut, RWTH Aachen University, D-52056 Aachen, Germany*²*School of Chemistry and Physics, University of Adelaide, Adelaide, South Australia 5005, Australia*³*Department of Physics and Astronomy, University of Alaska Anchorage, 3211 Providence Drive, Anchorage, Alaska 99508, USA*⁴*CTSPS, Clark-Atlanta University, Atlanta, Georgia 30314, USA*⁵*School of Physics and Center for Relativistic Astrophysics, Georgia Institute of Technology, Atlanta, Georgia 30332, USA*⁶*Department of Physics, Southern University, Baton Rouge, Louisiana 70813, USA*⁷*Department of Physics, University of California, Berkeley, California 94720, USA*⁸*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*⁹*Institut für Physik, Humboldt-Universität zu Berlin, D-12489 Berlin, Germany*¹⁰*Fakultät für Physik & Astronomie, Ruhr-Universität Bochum, D-44780 Bochum, Germany*¹¹*Physikalisches Institut, Universität Bonn, Nussallee 12, D-53115 Bonn, Germany*¹²*Science Faculty CP230, Université Libre de Bruxelles, B-1050 Brussels, Belgium*¹³*Vrije Universiteit Brussel, Dienst ELEM, B-1050 Brussels, Belgium*¹⁴*Department of Physics, Chiba University, Chiba 263-8522, Japan*

- ¹⁵*Department of Physics and Astronomy, University of Canterbury, Private Bag 4800, Christchurch 8140, New Zealand*
- ¹⁶*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*
- ¹⁷*Department of Physics and Center for Cosmology and Astroparticle Physics, Ohio State University, Columbus, Ohio 43210, USA*
- ¹⁸*Department of Astronomy, Ohio State University, Columbus, Ohio 43210, USA*
- ¹⁹*Department of Physics, TU Dortmund University, D-44221 Dortmund, Germany*
- ²⁰*Department of Physics, University of Alberta, Edmonton, Alberta T6G 2E1, Canada*
- ²¹*Département de physique nucléaire et corpusculaire, Université de Genève, CH-1211 Genève, Switzerland*
- ²²*Department of Physics and Astronomy, University of Gent, B-9000 Gent, Belgium*
- ²³*Department of Physics and Astronomy, University of California, Irvine, California 92697, USA*
- ²⁴*Laboratory for High Energy Physics, École Polytechnique Fédérale, CH-1015 Lausanne, Switzerland*
- ²⁵*Department of Physics and Astronomy, University of Kansas, Lawrence, Kansas 66045, USA*
- ²⁶*Department of Astronomy, University of Wisconsin, Madison, Wisconsin 53706, USA*
- ²⁷*Department of Physics and Wisconsin IceCube Particle Astrophysics Center, University of Wisconsin, Madison, Wisconsin 53706, USA*
- ²⁸*Institute of Physics, University of Mainz, Staudinger Weg 7, D-55099 Mainz, Germany*
- ²⁹*Université de Mons, 7000 Mons, Belgium*
- ³⁰*T.U. Munich, D-85748 Garching, Germany*
- ³¹*Bartol Research Institute and Department of Physics and Astronomy, University of Delaware, Newark, Delaware 19716, USA*
- ³²*Department of Physics, University of Oxford, 1 Keble Road, Oxford OX1 3NP, United Kingdom*
- ³³*Department of Physics, University of Wisconsin, River Falls, Wisconsin 54022, USA*
- ³⁴*Oskar Klein Centre and Department of Physics, Stockholm University, SE-10691 Stockholm, Sweden*
- ³⁵*Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794-3800, USA*
- ³⁶*Department of Physics and Astronomy, University of Alabama, Tuscaloosa, Alabama 35487, USA*
- ³⁷*Department of Astronomy and Astrophysics, Pennsylvania State University, University Park, Pennsylvania 16802, USA*
- ³⁸*Department of Physics, Pennsylvania State University, University Park, Pennsylvania 16802, USA*
- ³⁹*Department of Physics and Astronomy, Uppsala University, P.O. Box 516, S-75120 Uppsala, Sweden*
- ⁴⁰*Department of Physics, University of Wuppertal, D-42119 Wuppertal, Germany*
- ⁴¹*DESY, D-15735 Zeuthen, Germany*

(Received 29 May 2013; published 28 May 2014)

We report on the observation of a significant deficit of cosmic rays from the direction of the Moon with the IceCube detector. The study of this “Moon shadow” is used to characterize the angular resolution and absolute pointing capabilities of the detector. The detection is based on data taken in two periods before the completion of the detector: between April 2008 and May 2009, when IceCube operated in a partial configuration with 40 detector strings deployed in the South Pole ice, and between May 2009 and May 2010 when the detector operated with 59 strings. Using two independent analysis methods, the Moon shadow has been observed to high significance ($> 6\sigma$) in both detector configurations. The observed location of the shadow center is within 0.2° of its expected position when geomagnetic deflection effects are taken into account. This measurement validates the directional reconstruction capabilities of IceCube.

DOI: [10.1103/PhysRevD.89.102004](https://doi.org/10.1103/PhysRevD.89.102004)

PACS numbers: 96.50.S-, 95.85.Ry, 96.20.-n, 29.40.Ka

*Corresponding author.
boersma@icecube.wisc.edu

†Corresponding author.
santander@icecube.wisc.edu

‡Present address: Physics Department, South Dakota School of Mines and Technology, Rapid City, South Dakota 57701, USA.

§Present address: Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA.

|| Also at Sezione INFN, Dipartimento di Fisica, I-70126 Bari, Italy.

¶Present address: Department of Physics, Sungkyunkwan University, Suwon 440-746, Korea.

**Present address: NASA Goddard Space Flight Center, Greenbelt, Maryland 20771, USA.

I. INTRODUCTION

IceCube is a km³-scale Cherenkov detector deployed in the glacial ice at the geographic South Pole. Its primary goal is to search for astrophysical sources of high-energy neutrinos. A major background for this search is the high rate of atmospheric muons produced when cosmic rays with energies above a few TeV interact with the Earth's atmosphere. The rate of muon events in IceCube above several hundred GeV dominates the total trigger rate of the detector, and is approximately six orders of magnitude higher than the rate of neutrino-induced events.

The incoming direction of multi-TeV cosmic muons is on average within 0.1° of the arrival direction of the primary cosmic-ray particle [1]. This implies that the distribution of incoming muons should mimic the almost isotropic distribution of TeV cosmic rays in the sky [2,3]. An important feature of the angular distribution of cosmic rays is the presence of a relative deficit in the flux of cosmic rays coming from the direction of the Moon. This effect, due to the absorption of cosmic rays by the Moon, was first predicted by Clark in 1957 [4], and its observation has been used by several experiments as a way of calibrating the angular resolution and the pointing accuracy of their particle detectors (see [5–8], or [9] for recent results.)

For IceCube, the Moon shadow analysis is a vital and unique verification tool for the track reconstruction algorithms that are used in the search for point-like sources of astrophysical neutrinos [10], among other analyses. In this paper we will report on the observation of the Moon shadow using data taken between April 2008 and May 2010, before the completion of the IceCube Neutrino Observatory in December 2010.

Two independent analysis methods were used in the search for the Moon shadow. The first analysis performs a binned, one-dimensional search for the Moon shadow that compares the number of events detected from the direction of the Moon to the number of background events recorded at the same declination as the Moon but at a different right ascension. The second method uses an unbinned, two-dimensional maximum-likelihood algorithm that retrieves the best fit value for the total number of events shadowed by the Moon.

Both methods show consistent results, and constitute the first statistically significant detection of the shadow of the Moon using a high-energy neutrino telescope.

II. DETECTOR CONFIGURATION AND DATA SAMPLE

A. The IceCube detector

The IceCube neutrino telescope uses the deep Antarctic ice as a detection medium. High-energy neutrinos that interact with nucleons in the ice produce relativistic leptons that emit Cherenkov radiation as they propagate through the detector volume. This Cherenkov light is detected by a

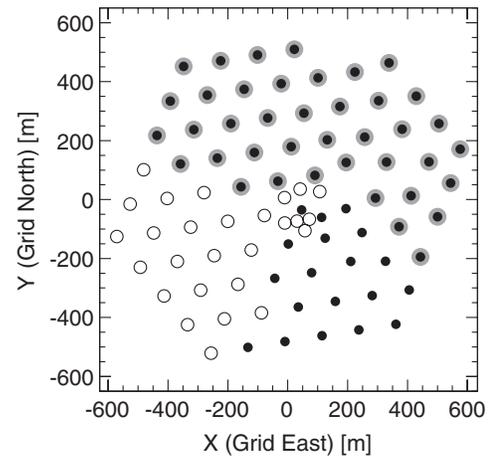


FIG. 1. Layout of the two detector configurations considered in this analysis. IC40 (gray) operated between 2008 and 2009. The deployment of more strings initiated the IC59 configuration (black) operated between 2009 and 2010. The remaining strings that form the final 86-string configuration, the last of them installed in December 2010, are shown as open circles. The y axis (Grid North) is aligned with the Greenwich Prime Meridian.

volumetric array of 5160 digital optical modules (DOMs) deployed at depths between 1450 and 2450 m below the ice surface. Each DOM consists of a 25 cm diameter photomultiplier tube (PMT) [11] and the electronics for signal digitization [12] housed inside a pressure-resistant glass sphere.

The DOMs are attached to 86 strings that provide mechanical support, electrical power, and a data connection to the surface. Consecutive DOMs in each string are vertically separated by a distance of about 17 m, while the horizontal spacing between strings is about 125 m. A compact group of eight strings with a smaller spacing between DOMs is located at the bottom of the detector and forms DeepCore [13], which is designed to extend the energy reach of IceCube to lower neutrino energies. The IceTop surface array, devoted to the detection of extensive air showers from cosmic rays with energies between 300 TeV and 1 EeV, completes the instrumentation of the observatory.

The construction of IceCube began in 2005 and was completed in December 2010. During construction, the detector operated in several partial configurations. Data from two different configurations were used in this paper: between 2008 and 2009 the detector operated with 40 strings deployed in the ice (IC40), and between 2009 and 2010 the detector operated in its 59-string configuration (IC59). The layout of the two detector configurations used in this work can be seen in Fig. 1.

B. Data sample

In order to reduce the rate of noise-induced events, IceCube DOMs are operated in a coincidence mode called hard local coincidence (HLC). During the operation of

IC40 and IC59 the HLC requirement was met if photon hits were detected within a $\pm 1 \mu\text{s}$ window in the two nearest neighbor or next-to-nearest neighbor DOMs. The detection of HLC hits leads to a full readout and transmission to the surface of the digitized PMT signals. A trigger condition is then used to combine these photon hits into a candidate event. The main trigger in IceCube is a simple multiplicity trigger called SMT8 that requires HLC hits in eight DOMs within $5 \mu\text{s}$. For each trigger, all HLC hits within a $\pm 10 \mu\text{s}$ window are recorded and merged into a single event.

The majority of events detected by IceCube are due to down-going muons produced in the interaction of high-energy cosmic rays with the Earth's atmosphere. During the operation of IC40, the cosmic muon-induced trigger rate was about 1.1 kHz, which increased to about 1.7 kHz during the IC59 data-taking period. This high rate of cosmic-ray muon events provides a high-statistics data set that can be used to search for the Moon shadow.

Since the rate of data transfer from the South Pole via the South Pole archival and data exchange satellite communication system is limited to about 100 Gb per day, only a limited number of muon events can be transmitted north over the satellite. For this reason, the data used in this analysis were taken using a dedicated online filter that selects only events passing minimum quality cuts and reconstructed within a predefined angular acceptance window around the Moon.

A fast likelihood-based muon track reconstruction [14] is performed at the South Pole to obtain the arrival direction of each event. The reconstructed direction of the muon track is then compared to the position of the Moon in the sky, which is calculated using the publicly available SLALIB library of astronomical routines [15].

An event satisfies the Moon filter selection criterium if at least 12 DOMs in 3 different strings record photon hits, and if the reconstructed direction is within 10° of the Moon position in declination and $40^\circ / \cos(\delta_\mu)$ in right ascension (where δ_μ is the declination of the event and the cosine factor accounts for projection effects).

The filter is enabled when the Moon is at least 15° above the horizon. Due to the particular geographic location of IceCube at the South Pole, the Moon rises above this threshold only once per month, as its elevation above the horizon changes slowly over the course of days. Since the number of muon events recorded by IceCube is a strong function of the elevation angle, the rate of events that pass the acceptance window condition changes during this period as this window follows the apparent motion of the Moon at the South Pole. The strong correlation between the Moon elevation and rate of events passing the Moon filter is shown in Fig. 2. The maximum event rate is also modulated over a longer time scale of 18.6 years (known as the lunar draconic period [16]) in which the maximum elevation of the Moon above the horizon at the South Pole oscillates between the extreme values of 18.4° and 28.4° . The maximum Moon

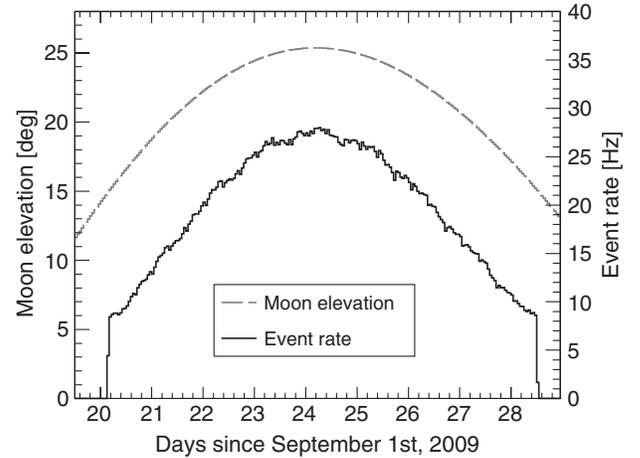


FIG. 2. Rate of muon events passing the Moon filter during the month of September 2009, when IceCube was operating in its IC59 configuration. The correlation between the Moon elevation (dashed line) and event rate (solid line) is clearly visible.

elevation during the IC40 data-taking period was 26.9° , while for IC59 it was 25.6° . Approximately 1.29×10^8 muon events passing the Moon filter condition were recorded during the IC40 data-taking period, and about 1.77×10^8 events were recorded during the operation of the IC59 configuration.

Once these events have been transferred from the South Pole, an iterative maximum-likelihood reconstruction algorithm is applied to the data set to obtain a more precise track direction [14]. The algorithm also determines the angular uncertainty in the reconstructed track direction by mapping the likelihood space around the best track solution and fitting it with a paraboloid function [17]. A narrow paraboloid indicates a precise reconstruction, while a wide paraboloid indicates a larger uncertainty in the reconstructed direction of the muon track. The 1σ contour line of the paraboloid function defines an error ellipse for the reconstructed direction of the track. In this analysis, a single, one-dimensional estimator of the uncertainty is obtained by calculating the root-mean-square (rms) value of the semimajor axes of that error ellipse.

The likelihood-based track reconstruction algorithm used in this work is based on the leading-edge times of the first light pulses recorded by each DOM. For the fast track reconstruction at the South Pole the single-photoelectron (SPE) fit is used. In this fit, the likelihood that the first photon arrived at the pulse leading-edge times is maximized. The photons arriving at later times are ignored.

Neutrino point source searches rely on the multiphotoelectron (MPE) fit. In the MPE fit, the total number of photoelectrons N_d in each DOM, d , is taken into account by multiplying the likelihood that a photon was detected at the first leading-edge time with the probability that the remaining $N_d - 1$ photons arrived later [14]. For bright, i.e. high-energy, events in simulated data the MPE fit results in a

slightly better angular resolution than the SPE fit. Also, the number of direct (unscattered) photons associated with the reconstructed track tends to be larger with the MPE fit than with the SPE fit. This makes this quantity as well as related quantities more effective for selecting well-reconstructed events. The MPE fit is discussed further in Sec. VB 2.

The track reconstruction algorithms use the local detector coordinate system and the direction of a reconstructed track is given as a zenith and azimuth angle. Using the event times as recorded by the data acquisition system, these are transformed into a right ascension, α_μ , and declination, δ_μ , which are the more natural variables for searches of neutrino point-like sources.

III. SIMULATION

A. Cosmic-ray energy and composition

The muons produced in the interaction between the cosmic rays and the atmosphere must traverse several kilometers of ice before reaching the IceCube detector, losing energy in the process. This sets a lower limit of several hundred GeV on the energy of the muons at the ground level that would trigger the detector. By extension, the primary cosmic-ray particle needed to produce this kind of muon should have an energy of at least several TeV. In the following, we will refer to the energy of the primary cosmic ray, not the muons, unless specified otherwise.

Given that this analysis deals with cosmic-ray showers near the energy threshold of the detector, the number of muons produced in each shower that reaches the detector is small. Most events in the Moon data sample are composed of one or two energetic muons, and only 2% of the events have muon multiplicities higher than ten.

The detailed energy scale for the IC40 and IC59 data sets was determined using simulated cosmic-ray air showers created with the CORSIKA Monte Carlo code [18] using the SIBYLL model of high-energy hadronic interactions [19]. The chemical composition and spectral shape of the cosmic rays generated in this simulation follow the polygono model [20].

From these simulations, we estimate that the median energy of the primary cosmic rays that trigger the IceCube detector is 20 TeV, while the median energy of events that satisfy the Moon filter condition is about 40 TeV for both IC40 and IC59, with 68% of the events between 10 and 200 TeV. The increased median energy of the filtered sample is due to the greater average zenith angles of the cosmic rays that pass the filter, which requires primary particles with enough energy to produce muons able to traverse more ice and trigger IceCube. The muons produced by cosmic rays passing the Moon filter have a mean energy of about 2 TeV at the ground level and reach the detector with a mean energy of 200 GeV. The mean muon energy also depends on the zenith angle, and increases from 2.5 TeV for a zenith angle of 65° (the maximum elevation

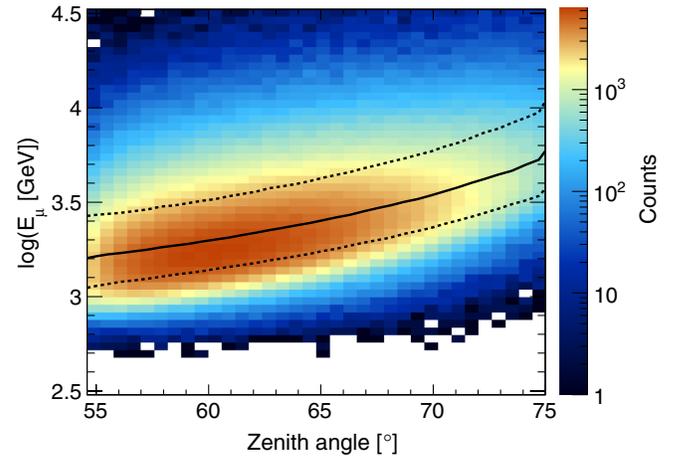


FIG. 3 (color online). Muon energy at the ground level as a function of the zenith angle for events passing the Moon filter condition. The solid line indicates the median muon energy in each zenith bin, while the 68% containing interval is defined between the two dashed lines. The monotonic increase in muon energy with the zenith angle is a consequence of the larger ice overburden for very inclined events, which raises the initial muon energy necessary to reach the detector.

of the Moon) to 5.6 TeV for a zenith angle of 75° (the minimum elevation of the Moon for which data is recorded.) This dependence is shown in Fig 3.

The energy spectrum of all primary cosmic rays triggering the IceCube detector is shown in Fig. 4 and compared to the spectrum of those that pass the Moon filter. Also shown in the figure are the five main chemical elements (protons, He, C, O, and Fe) that make more than 95% of the Moon filter data sample assuming the polygono composition model. The two main components of the sample are proton (68% of the events) and helium (23%).

As will be described in the following subsection, an important quantity for calculating the deflection of cosmic rays in the magnetic field of the Earth is the particle rigidity $R \sim E/Z$, for a particle with the energy E and electric charge Z . The distribution of magnetic rigidities for the sample is given in Fig. 4 for reference.

B. Geomagnetic field effects

Cosmic rays with TeV energies should experience a small deflection in their trajectories due to the influence of the magnetic field of the Earth as they propagate towards the detector. This deflection would appear in the Moon shadow analysis as a shift in the position of the shadow with respect to the true Moon position, which could be wrongly interpreted as a systematic offset produced by the event reconstruction.

In order to quantify this offset and compare it with any possible shift observed in the data, we have developed a particle propagation code that can be used to trace cosmic rays in the geomagnetic field. Using this code, particles are

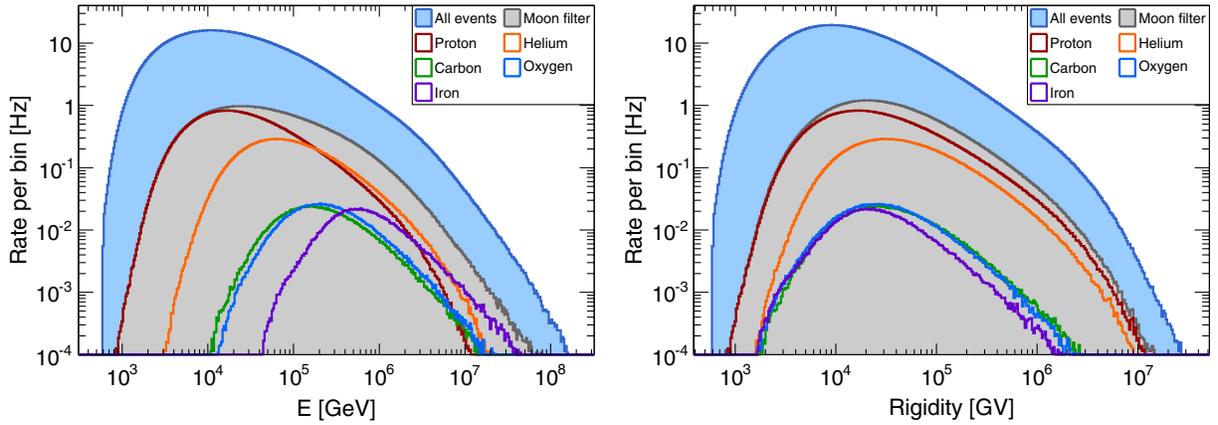


FIG. 4 (color online). Differential event rate as a function of cosmic-ray primary energy (left) and rigidity (right) for all events in IC59 (light blue) and for only those passing the Moon filter (gray) as determined from simulation studies. The main chemical elements that make up the events passing the Moon filter are shown with lines of different color. The width of the histogram bins is 0.014 in $\log_{10}(\text{energy, rigidity})$. The IC40 configuration shows a similar energy response.

propagated radially outwards from the South Pole up to a distance of 30 Earth radii from the center of the Earth at which point the opening angle between the initial and final velocity vectors is computed. This angle gives the magnitude of the deflection in the geomagnetic field.

We use the international geomagnetic reference field (IGRF) model [21] to calculate deflections. In this model, the field is calculated using a truncated multipole series expansion. The current revision of the model, IGRF-11, can be used to calculate B-field values through 2015, providing a good coverage of the time range over which the data were taken. The model is accessible through a library of FORTRAN routines called GEOPACK developed by N. Tsyganenko.¹

The IGRF-11 model describes what is known as the *internal* magnetic field of the Earth, which is presumably produced by electric currents in the outer core of the planet and accounts for most of the total magnetic field. A weaker component, known as the *external* field, is produced by electrical currents in the ionosphere. The external component is not included in our calculation since it only modifies the total angular deflection by a few percent while significantly increasing the computation time needed to perform the simulation.

In our simulation, primary cosmic rays are propagated in the direction of the Moon as seen from the South Pole for different times during the data-taking period. The cosmic-ray energy and chemical composition is sampled from the event distributions that pass the Moon filter, shown in Fig. 4. The resulting total deflection $\Delta\lambda$ is shown in Fig. 5 as a function of energy for 10^5 simulated cosmic-ray particles for the five main chemical elements that contribute to the Moon data set. The energy and charge dependence of the deflection angle is evident in the plot. Different bands in

the plot correspond to different chemical elements. The width of each band is due to particles that were propagated in different directions in the sky (i.e. through different regions of the Earth's magnetic field) experiencing different deflections. A power-law fit to the simulation results has been performed to estimate the deflection angle as a function of energy and charge. The fit gives a good agreement for the following expression:

$$\Delta\lambda[^\circ] = 1.9^\circ \frac{Z}{E[\text{TeV}]} \sim \frac{1.9^\circ}{R[\text{TV}]}, \quad (1)$$

where Z is the charge of the cosmic ray (CR) in units of the elementary charge e , E is its energy in TeV, and $\Delta\lambda$ is given in degrees. This expression has the same functional form as the one found in [22] with a higher normalization in our simulation, which could be due to the difference in geographic location and other simulation details. The

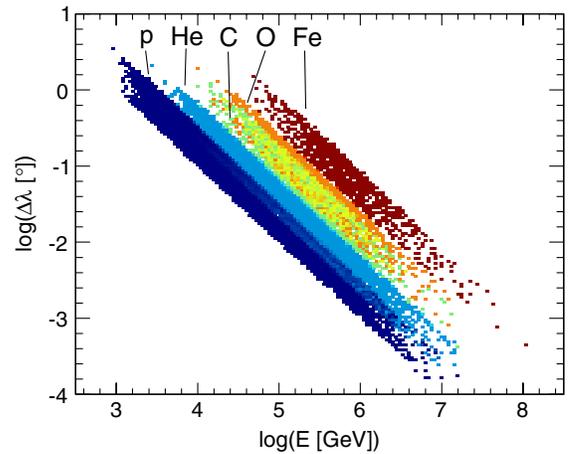


FIG. 5 (color online). Angular deflection as a function of energy for the different chemical elements simulated using the particle propagation code described in Sec. III B.

¹<http://geo.phys.spbu.ru/~tsyganenko/modeling.html>.

deflection angle is also given as a function of the rigidity R of the cosmic ray in teravolts (TV).

The deflection of each cosmic ray with the arrival direction (α_μ, δ_μ) in sidereal coordinates is calculated with respect to the position of the Moon at the time of the event $(\alpha_{\text{Moon}}, \delta_{\text{Moon}})$. The two coordinates that characterize the position of an event in this system are a right ascension difference, $\Delta\alpha = (\alpha_\mu - \alpha_{\text{Moon}}) \cos \delta_\mu$, and a declination difference, $\Delta\delta = \delta_\mu - \delta_{\text{Moon}}$, with respect to the nominal Moon position. The median shift in the right ascension $\Delta\alpha$ for all CR particles in our simulation is 0.08° , with 68% of the particles having deflection angles in the interval $0.02^\circ < \Delta\alpha < 0.24^\circ$. The median shift in the declination $\Delta\delta$ is consistent with 0° , with 68% of the events contained in the interval $|\Delta\delta| < 0.04^\circ$. Propagation tests performed using only the dipole term of the geomagnetic field instead of the full IGRF model produce similar results.

The cosmic-ray muons that ultimately trigger IceCube are also deflected by the geomagnetic field. However, given the typical energies of the muons (about 2 TeV), their total track length (in the 50–100 km range), and their charge distribution, their contribution to the total deflection angle is typically below $\sim 0.01^\circ$. For this reason, the muon contribution has been ignored in calculating the expected total deflection angle.

The direction of muons propagating through the ice is smeared due to multiple Coulomb scattering (MCS). Given the typical muon energies and the thickness of the ice overburden, we estimate that the average rms opening angle for deflections due to MCS is about 0.05° . This smearing is much smaller than the angular resolution of IceCube for muons in this energy range (typically between 0.5 and 0.9°) and is already included in the estimate of the angular resolution from simulation studies given in Secs. IV and V.

IV. BINNED ANALYSIS

A. Description of the method

The main goal of the binned analysis is to obtain a profile view of the Moon shadow and measure its width, which can be used as a direct estimator of the angular resolution of the event reconstruction. This is accomplished by comparing the observed number of events as a function of angular distance from the Moon to an estimate of how many events would have been observed if there was no shadow.

For this comparison, the angular distance between the reconstructed muon tracks and the expected position of the Moon is binned in constant increments of 0.2° up to a maximum angular distance of 5° . This defines the so-called *on-source* distribution of events. The same binning procedure is applied to eight *off-source* regions centered around points located at the same declination as the Moon, but offset from it in the right ascension by $\pm 5^\circ, \pm 10^\circ, \pm 15^\circ$, and $\pm 20^\circ$, where it is assumed that the shadowing effect is negligible. The average number of counts as a function of

radius for these eight off-source regions represents the expectation in the case of no Moon shadow.

The relative difference between the number of events in the i th bin in the on-source region N_i^{on} , and the average number of events in the same bin in the off-source regions $\langle N_i^{\text{off}} \rangle$ is calculated using the following expression:

$$\frac{\Delta N_i}{\langle N \rangle_i} = \frac{N_i^{\text{on}} - \langle N_i^{\text{off}} \rangle}{\langle N_i^{\text{off}} \rangle}. \quad (2)$$

The uncertainty in the relative difference is given by

$$\sigma_{\Delta N/\langle N \rangle} = \frac{N_i^{\text{on}}}{\langle N_i^{\text{off}} \rangle} \sqrt{\frac{1}{N_i^{\text{on}}} + \frac{1}{s \langle N_i^{\text{off}} \rangle}}, \quad (3)$$

where $s = 8$ is the number of off-source regions. The distribution of relative differences as a function of angular radius from the Moon constitutes a profile view of the shadow.

Simulation studies indicate that the point spread function (PSF) of the detector can be approximated with a two-dimensional Gaussian function. We use this approximation to obtain an estimate of the angular resolution of the track reconstruction by fitting the distribution of $\Delta N_i / \langle N \rangle_i$ for the events in the Moon data set.

Following [23], we treat the Moon as a point-like cosmic-ray sink that removes $\Phi \pi R_M^2$ events from the muon sample, where R_M is the angular radius of the Moon ($R_M \sim 0.26^\circ$) and Φ is the cosmic-ray flux at the location of the Moon in units of events per square degree. This deficit is smeared by the PSF of our detector, resulting in a radially symmetric two-dimensional Gaussian distribution of shadowed events. The differential density of shadowed events per solid angle, Ω , can be expressed as

$$\frac{dN}{d\Omega} = -\frac{\Phi R_M^2}{2\sigma^2} e^{-\psi^2/2\sigma^2}, \quad (4)$$

where ψ is the radial distance from the center of the Moon. The number of shadowed events in the i th bin of the width $\Delta\psi$ can be calculated by integrating the differential event density over the bin annulus in polar coordinates (ψ, ϕ) :

$$N_s(\psi_i) = \int_0^{2\pi} \int_{\psi_i - \Delta\psi/2}^{\psi_i + \Delta\psi/2} \psi \frac{dN}{d\Omega} d\psi d\phi \quad (5)$$

$$\approx -\frac{\Phi \pi R_M^2 \Delta\psi}{\sigma^2} \psi_i e^{-\psi_i^2/2\sigma^2}. \quad (6)$$

The number of events, N_e , that would have been observed in the same bin with no shadowing is $2\pi\Phi\psi_i\Delta\psi$. The ratio of Eqs. (6) and N_e gives us the expected distribution of relative differences $\Delta N_i / \langle N \rangle_i$ for a detector with a Gaussian PSF of the angular resolution σ :

TABLE I. Optimal bin radius (ψ_b), the number of observed events in the on-source (N_{on}^b) and off-source (N_{off}^b) bins, the event deficit in the on-source bin (ΔN), and the statistical significance of the deficit for the binned analysis of IC40 and IC59 data sets.

	IC40	IC59
ψ_b	0.75°	0.79°
N_{on}^b	52967	96412
N_{off}^b	54672	100442
ΔN	-1705	-4030
Significance	6.9 σ	12.1 σ

$$\frac{N_s}{N_e}(\psi_i) = -\frac{R_M^2}{2\sigma^2} e^{-\psi_i^2/2\sigma^2}. \quad (7)$$

This expression is used to fit the experimental data. The only free parameter in the fit is σ , which is used as the estimator for the angular resolution of the experimental data set. The value obtained from the fit can be compared to the expected angular resolution obtained from simulation studies. Following previous Moon shadow studies [8,23], we use the Gaussian σ parameter as the estimator instead of a 68% containing radius used elsewhere in the literature.

Our treatment ignores the finite angular size of the lunar disc, which may affect the result of the fit. However, since the expected angular resolution (of the order 1°) is several times larger than the angular radius of the Moon, we expect that the effect should influence the fit value of σ only at the few-percent level.

A set of cuts was developed to optimize for the statistical significance of the detection of the Moon shadow. Under the assumption of Poisson statistics, the relation between the significance S , the fraction η of events passing the cuts, and the resulting median angular resolution Ψ_{med} after cuts is

$$S \propto \frac{\sqrt{\eta}}{\Psi_{\text{med}}}. \quad (8)$$

The optimization of the cuts was performed on the CORSIKA-simulated air showers described in Sec. III.

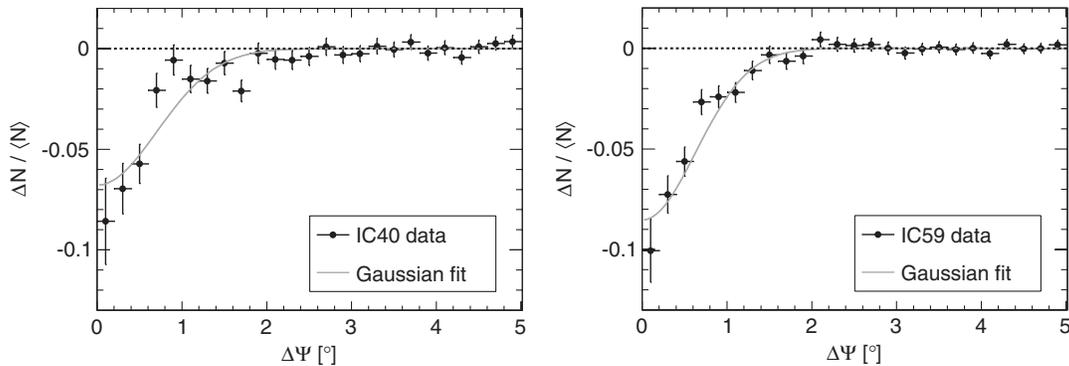


FIG. 6. Relative difference between the number of events in the *on-source* and the average *off-source* region as a function of the angular distance from the nominal position of the Moon for the IC40 (*left*) and IC59 (*right*) data sets. A Gaussian fit to the deficits is shown in gray.

Two cut variables were used in this analysis: the angular uncertainty σ_i in the reconstruction of the muon track direction estimated individually for each event, and the reduced log likelihood $rlogl$, which is the log likelihood for the best track solution divided by the number of degrees of freedom in the fit. The number of degrees of freedom in the track fit is equal to the number of DOMs triggered by the event minus the number of free parameters in the fit (five for this fit). Both $rlogl$ and σ are standard cut variables used in the search for point-like sources of astrophysical neutrinos [10], the search for a diffuse flux of high-energy neutrinos [24], and several other analyses of IceCube data.

Once the cuts have been determined, the number of events falling inside a circular search bin around the Moon is compared to the number of events contained in a bin of the same angular radius for the average off-source region. The statistical significance of an observed deficit in the number of events in the search bin is calculated using the method given by [25].

The optimal radius of the search bin ψ_b can be found by maximizing the S' parameter in the following expression:

$$S'(\psi_s) \propto \frac{\int_0^{\psi_s} \psi' \text{PSF}(\psi') d\psi'}{\psi_s}, \quad (9)$$

where ψ_s is the radius of the bin and $\text{PSF}(\psi')$ is the point spread function of the detector after cuts obtained from simulations. Due to its symmetry, the PSF has already been integrated over the azimuthal coordinate and only the radial dependence remains. The optimization of the search bin radius is also performed using simulated CORSIKA showers generated for each detector configuration.

B. Results

A set of cuts was determined independently for both the IC40 and IC59 detector configurations using the optimization procedure described above on simulated data. For IC40, only events with $rlogl < 9$ and $\sigma_i < 1.01^\circ$ were used in the analysis, with 26% of the events surviving the cuts.

TABLE II. Gaussian angular resolution σ obtained from the fit to the Moon shadow profile shown in Fig. 6. The χ^2/dof of the fit is also given for the two results.

	IC40	IC59
σ	$0.71^\circ \pm 0.07^\circ$	$0.63^\circ \pm 0.04^\circ$
χ^2/dof	31.4/24	13.0/24

After cuts, the median angular resolution of the reconstruction was estimated from simulation to be 0.93° , with 68% of the events having angular uncertainties, σ_i , between 0.38 and 2.18° . A two-dimensional fit to the simulated data shows that for the Gaussian approximation the corresponding resolution σ is about 0.74° .

In the case of IC59, the events selected for the analysis were those with $rlog < 8.8$ and $\sigma_i < 1.04^\circ$, which resulted in a passing rate of 34%. The median resolution after cuts was 0.78° , with the 68% containing interval located between 0.33° and 1.78° , with a Gaussian width, σ , of about 0.71° .

After the cuts were applied to both data sets, the radius of the optimal search bin (ψ_b) and the number of events contained in that bin for both the on-source (N_{on}^b) and off-source (N_{off}^b) windows was calculated. In both detector configurations, a deficit in the number of events in the on-source bin when compared to the off-source bin was observed at high statistical significance ($> 6\sigma$), as expected due to the shadowing effect of the Moon. A complete list of the number of events observed on each bin, the observed deficit in the on-source bin, as well as the statistical significance associated with such a deficit, is given in Table I.

The Moon shadow profile shown in Fig. 6 was fit using the expression given in Eq. (7), where σ is the only free parameter. A list of fit results is given in Table II. In both cases, the observed angular resolution shows good agreement with the one obtained from the above-mentioned simulation studies.

V. UNBINNED ANALYSIS

A. Description of the method

The second algorithm used to search for the Moon shadow is based on an unbinned maximum-likelihood method analogous to that used in the search for point-like sources of high-energy neutrinos [26]. This kind of likelihood analysis was first proposed in [27], and was applied for the first time to a Moon shadow search in [28].

The goal of the unbinned analysis is to determine the most likely location of the Moon shadow to compare it with the expected location after accounting for magnetic deflection effects. An agreement between the observed and expected positions of the shadow center will serve as an important confirmation of the absolute pointing accuracy of the detector.

The analysis is also used to obtain the most likely number of events shadowed by the Moon, which can be

compared to the expectation. An essential ingredient in the unbinned analysis is an event-wise estimation of the angular error. Both systematic underestimation and overestimation of this error would lead to a shallower apparent shadow than expected. The number of shadowed events is a free parameter in this analysis and the comparison with the expected number of shadowed events is effectively a test of the angular uncertainty estimate.

In this analysis [29,30], the position of each muon event is defined with respect to the Moon position in the coordinate system $(\Delta\alpha, \Delta\delta)$ that was defined in Sec. III B. Only events with $|\Delta\delta| \leq 8^\circ$ and $|\Delta\alpha + \alpha_{\text{off}}| \leq 8^\circ$ were considered in the analysis, where $\alpha_{\text{off}} = 0^\circ$ defines the on-source region, and $\alpha_{\text{off}} = \pm 18^\circ$ defines two off-source regions.

A set of quality cuts was determined for this analysis using the same simulation data set as in the one-dimensional binned case. The same variables were used in the optimization of the cuts: the angular reconstruction uncertainty σ_i , and the reduced log likelihood of each event $rlogl$.

The analysis method assumes that the data can be described as a linear combination of signal and background components, where the relative contribution from each component is established by a maximum-likelihood fit to the data. For a data set containing N events, the log-likelihood function is defined as

$$\log \mathcal{L}(n_s, \vec{x}_s) = \sum_{i=1}^N \log \left[\frac{n_s}{N} \mathcal{S}(\vec{x}_i, \sigma_i; \vec{x}_s) + \left(1 - \frac{n_s}{N} \right) \mathcal{B}(\vec{x}_i) \right], \quad (10)$$

where \mathcal{S} and \mathcal{B} are the signal and background probability density functions (PDFs), n_s is the unknown number of signal events, or in this case the total number of shadowed events, and \vec{x}_s is the unknown central position of the

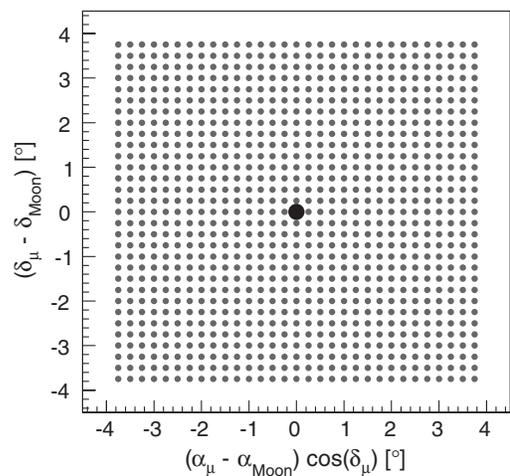


FIG. 7. Search grid in $(\Delta\alpha, \Delta\delta)$ as used in the unbinned likelihood analysis. A value of n_s is determined for each one of the points in the grid. The nominal location of the Moon is shown as a black square at $\vec{x}_s = (0, 0)$.

TABLE III. Description of the cuts used in the unbinned analysis. Percentages indicate the relative fraction of events that survive the cut with respect to the previous selection criterium.

	IC40	IC59
Events before cuts	18.8×10^6	22.2×10^6
Cut 1: $0.075^\circ < \sigma_i < 1.5^\circ$	50%	58%
Cut 2: $6.5 < \text{rlogl} < 8$	89%	91%
Events after cuts	8.4×10^6	11.7×10^6

shadow of the Moon, relative to the nominal position of the Moon. Note that since the expected signal in the case of the Moon is a deficit in the muon flux rather than an excess, n_s should be negative. In the absence of a geomagnetic field, the shadow should occur exactly on the nominal position of the Moon, i.e. $\vec{x}_s = (0, 0)$, but according to the estimates described in Sec. III B we expect the shadow to be shifted by about 0.1° .

The signal PDF for each event is modeled using a two-dimensional Gaussian distribution around the reconstructed direction \vec{x}_i of the muon track:

$$\mathcal{S}(\vec{x}_i, \sigma_i; \vec{x}_s) = \frac{1}{2\pi\sigma_i^2} e^{-\frac{|\vec{x}_i - \vec{x}_s|^2}{2\sigma_i^2}}, \quad (11)$$

where the width of the Gaussian distribution σ_i is the angular reconstruction uncertainty obtained on an event-by-event basis by the paraboloid algorithm described in Sec. II B.

The background PDF is assumed to depend only on $\Delta\delta$, and is derived from the distribution of reconstructed declination angles for the muon tracks contained in the two off-source regions.

The best fit values for the number of signal events in the data n_s and the shift of the shadow center \vec{x}_s are determined by maximizing the log-likelihood function (10). Besides \vec{x}_s and n_s , the width and overall shape of the shadow are also of interest. In searches for point sources of high-energy neutrinos [26], for all points \vec{x}_s on a fine grid covering the

sky, the value of n_s is determined which maximizes the likelihood function. Similarly, in the Moon shadow analysis we determine the value of n_s that maximizes the likelihood function (10) on a rectangular grid of 961 values for $\vec{x}_s = (\Delta\alpha_s, \Delta\delta_s)$. This 31×31 grid is defined inside a window with a size of $|\Delta\delta| \leq 4^\circ$ and $|\Delta\alpha| \leq 4^\circ$ shown in Fig. 7.

In order to avoid edge effects, all events in the $8^\circ \times 8^\circ$ on-source region are taken into account in the maximum-likelihood calculation.

The statistical significance associated with each value of n_s can then be calculated by applying the same likelihood analysis to the two off-source regions. The rms spread of the resulting distribution of n_s values for those regions gives an estimate of the 1σ spread expected in the case of a null detection. Using this estimate, each point in the on-source region can be given a statistical significance by taking the ratio between the value of n_s at that point and the 1σ estimate from the off-source regions.

The observed value of n_s is compared to an estimate of the true number of CRs shadowed by the Moon. This estimate is obtained by counting the number of events that fall within a circular window with the same radius as the Moon but located in the off-source region.

B. Results

1. SPE analysis

The cuts used in the unbinned analysis are listed in Table III. The resulting median angular resolution of the IC40 and IC59 data sets was estimated by applying those same cuts to simulated cosmic-ray events. In the case of IC40, the median angular resolution is 1.13° , with 68% of the events having angular uncertainties, σ_i , between 0.48° and 2.63° . For IC59, the median resolution is 0.98° , with a 68% containing interval defined between 0.38 and 2.23° .

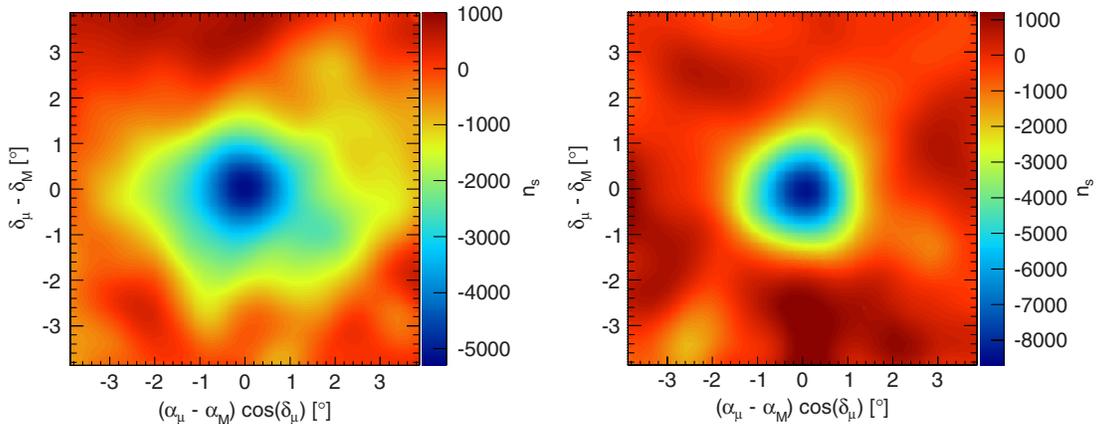


FIG. 8 (color online). Contour plot of the value of n_s in the $(\Delta\alpha, \Delta\delta)$ coordinate system for on-source regions of the IC40 (left) and IC59 data sets (right).

TABLE IV. Unbinned analysis results detailing the observed and expected deficit counts from the Moon for IC40 and IC59. The observed deficits and the $(\Delta\alpha, \Delta\delta)$ offsets are given for the most likely position of the Moon shadow as determined by the maximum-likelihood fit.

	IC40	IC59
Observed deficit	5320 ± 501	8700 ± 550
Expected deficit	5734 ± 76	8192 ± 91
Off-source rms	521	627
Significance	10.2σ	13.9σ
$\Delta\alpha$	-0.02°	0.06°
$\Delta\delta$	0.08°	0.00°

As described in the previous section, the maximum-likelihood values of n_s were calculated on a grid around the position of the Moon for both sets. The contour maps of the n_s values obtained for IC40 and IC59 are shown in Fig. 8, where the shadowing effect of the Moon is visible as a strong deficit in the central regions of the maps. The deepest deficit observed with both detector configurations is in good agreement with the expected number of shadowed events, listed in Table IV. Using the rms spread of the off-source regions as a 1σ estimator in the case of a null detection, we calculated the statistical significance of the observation by taking the ratio of the largest deficit observed to the rms spread, which is also shown in Table IV. The shadow of the Moon is observed in both the IC40 and IC59 data sets to a high statistical significance ($> 10\sigma$).

In order to obtain a better estimate of the position of the minimum of the shadow, a finer grid with a spacing of about 0.016° was used in the central $\pm 0.4^\circ \times \pm 0.4^\circ$ region around the Moon. Using this grid, we obtain the positions indicated in Table IV as offsets in right ascension ($\Delta\alpha$) and declination ($\Delta\delta$) with respect to the nominal position of the

Moon in the sky. The shadow positions for both detector configurations are shown in Fig. 9 together with 1σ , 2σ , and 3σ contours. The expected location of the minimum after accounting for geomagnetic deflection effects is also given for comparison. In both detector configurations, the observed position of the minimum is consistent with its expected location to within statistical fluctuations. These measurements imply that, on average, the absolute pointing accuracy of the detector during the IC40 and IC59 data-taking periods was better than about 0.2° .

2. MPE analysis

The unbinned analysis was also applied to IC59 events reconstructed using the MPE algorithm and its corresponding angular error estimate described in Sec. II B. In simulations, the MPE fit performs better than the SPE reconstruction thanks to its more realistic description of the arrival times of multiple photons at each DOM. However, at high energies the algorithm can be confused by stochastic energy losses that occur along the muon track and are not described in the likelihood function implemented in the MPE algorithm. This usually results in an underestimation of the angular uncertainty on the reconstructed direction of the track. In practice, this problem can be solved by rescaling the average pull (the ratio between the real and estimated angular errors as obtained from simulation studies) to unity. The MPE version of the unbinned analysis was used as a verification of this correction technique.

Simulation studies indicate an average pull of 1.55 for the MPE reconstruction, versus 1.0 for SPE. Without correcting for this underestimation of the angular error in the MPE fit, the Moon shadow analysis resulted in a minimum value for n_s of 3574 ± 434 shadowed events, differing by more than 5 standard deviations from the expectation of 6373 ± 80 . Redoing this analysis with the

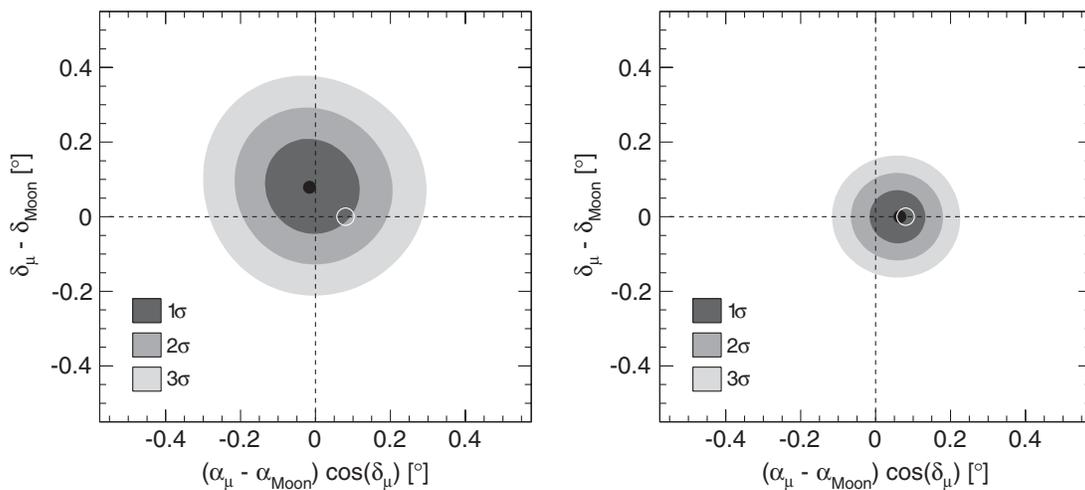


FIG. 9. Contour plot for the position of the minimum of the Moon shadow in the IC40 (left) and IC59 data (right) in the $(\Delta\alpha, \Delta\delta)$ coordinate system. The reconstructed position for the Moon shadow from the maximum-likelihood analysis is shown as a black point, while the expected position of the Moon shadow after accounting for magnetic deflection is shown as a white circle.

angular error estimates rescaled by a factor of 1.55 resulted in a fitted n_s value compatible with expectation, validating the pull correction method.

In neutrino analyses, where the range of muon energies is much larger than in the Moon analysis sample, the applied MPE pull correction is energy dependent, instead of using only the average value of the pull for the entire data set.

VI. CONCLUSIONS

The shadow of the Moon in TeV cosmic rays has been detected to a high significance ($> 6\sigma$) using data taken with the IC40 and IC59 configurations of the IceCube neutrino observatory. For both detector configurations, the observed positions of the shadow minimum show good agreement with expectations given the statistical uncertainties. An important implication of this observation is that any systematic effects introduced by the detector geometry and the event reconstruction on the absolute pointing capabilities of IceCube are smaller than about 0.2° .

The average angular resolution of both data samples was estimated by fitting a Gaussian function to the shadow profile. In both cases, the 1σ width of the Moon shadow was found to be about 0.7° , which is in good agreement with the expected angular resolution based on simulation studies of down-going muons.

The total number of shadowed events estimated using the unbinned analysis is also consistent with expectations for IC40 and IC59. This provides an indirect validation of the angular uncertainty estimator obtained from the reconstruction algorithm. This is especially relevant for the MPE analysis, where simulation studies indicate that the uncertainty estimator has to be rescaled in order to avoid underestimating the true angular error. Applying this correction factor to the data results in a number of shadowed events compatible with expectation.

Note that the value of the average angular resolution determined in this analysis is *not* a direct measurement of the point spread function to be used in searches for point sources of high-energy neutrinos. Rather, the agreement of this value with the value estimated from our simulations

should be seen as an experimental verification of our simulation and the methods used to estimate the angular uncertainty of individual track reconstructions. This angular uncertainty depends on several factors, in particular on the energy with which the muon traverses the detector. As the energy distribution for neutrino analyses differs from that of the Moon shadow analysis, the average angular resolution may be better or worse, but can reliably be estimated from our simulation.

ACKNOWLEDGMENTS

We acknowledge the support from the following agencies: the U.S. National Science Foundation, Office of Polar Programs, the U.S. National Science Foundation, Physics Division, the University of Wisconsin Alumni Research Foundation, the Grid Laboratory of Wisconsin (GLOW) grid infrastructure at the University of Wisconsin, Madison, the Open Science Grid (OSG) grid infrastructure; the U.S. Department of Energy, the National Energy Research Scientific Computing Center, the Louisiana Optical Network Initiative (LONI) grid computing resources; the Natural Sciences and Engineering Research Council of Canada, WestGrid and Compute/Calcul Canada; the Swedish Research Council, the Swedish Polar Research Secretariat, the Swedish National Infrastructure for Computing (SNIC), and the Knut and Alice Wallenberg Foundation, Sweden; the German Ministry for Education and Research (BMBF), Deutsche Forschungsgemeinschaft (DFG), the Helmholtz Alliance for Astroparticle Physics (HAP), the Research Department of Plasmas with Complex Interactions (Bochum), Germany; the Fund for Scientific Research (FNRS-FWO), FWO Odysseus programme, the Flanders Institute to encourage scientific and technological research in industry (IWT), the Belgian Federal Science Policy Office (Belspo); the University of Oxford, United Kingdom; the Marsden Fund, New Zealand; the Australian Research Council; the Japan Society for Promotion of Science (JSPS); and the Swiss National Science Foundation (SNSF), Switzerland.

-
- [1] R. Abbasi *et al.* (IceCube Collaboration), *Phys. Rev. D* **87**, 012005 (2013).
 - [2] R. Abbasi *et al.* (IceCube Collaboration), *Astrophys. J. Lett.* **718**, L194 (2010).
 - [3] R. Abbasi *et al.* (IceCube Collaboration), *Astrophys. J. Lett.* **740**, 16 (2011).
 - [4] G. W. Clark, *Phys. Rev.* **108**, 450 (1957).
 - [5] M. Ambrosio *et al.* (MACRO Collaboration), *Astropart. Phys.* **20**, 145 (2003).
 - [6] P. Achard, O. Adriani, M. Aguilar-Benitez, M. van den Akker, J. Alcaraz, G. Alemanni, J. Allaby, A. Aloisio, M. G. Alviggi, and H. Anderhub (L3 Collaboration), *Astropart. Phys.* **23**, 411 (2005).
 - [7] A. Oshima, S. R. Dugad, U. D. Goswami, S. K. Gupta, Y. Hayashi, N. Ito, A. Iyer, P. Jagadeesan, A. Jain, and S. Kawakami (GRAPES-3 Collaboration), *Astropart. Phys.* **33**, 97 (2010).
 - [8] P. Adamson *et al.* (MINOS Collaboration), *Astropart. Phys.* **34**, 457 (2011).

- [9] B. Bartoli *et al.* (ARGO-YBJ Collaboration), *Phys. Rev. D* **84**, 022003 (2011).
- [10] R. Abbasi *et al.* (IceCube Collaboration), *Astrophys. J.* **732**, 18 (2011).
- [11] R. Abbasi *et al.* (IceCube Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **618**, 139 (2010).
- [12] R. Abbasi *et al.* (IceCube Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **601**, 294 (2009).
- [13] R. Abbasi *et al.* (IceCube Collaboration), *Astropart. Phys.* **35**, 615 (2012).
- [14] J. Ahrens *et al.* (AMANDA Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **524**, 169 (2004).
- [15] P. Wallace, *ASP Conf. Ser.* **61**, 481 (1994).
- [16] *Encyclopedia of Astronomy and Astrophysics*, edited by R. A. Meyers and S. N. Shore (Academic Press, New York, 1989).
- [17] T. Neunhöffner, *Astropart. Phys.* **25**, 220 (2006).
- [18] D. Heck *et al.*, *CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers*, Technical Report No. FZKA 6019 (Forschungszentrum Karlsruhe, 1998).
- [19] E.-J. Ahn, R. Engel, T. Gaisser, P. Lipari, and T. Stanev, *Phys. Rev. D* **80**, 094003 (2009).
- [20] J. R. Hörandel, *Astropart. Phys.* **19**, 193 (2003).
- [21] C. C. Finlay *et al.*, *Geophys. J. Int.* **183**, 1216 (2010).
- [22] G. D. Sciascio and R. Iuppa, *Proceedings of the 2nd Roma International Conference on Astroparticle Physics (RI-CAP), Rome, Italy, 2009* [*Nucl. Instrum. Methods Phys. Res., Sect. A* **630**, 63 (2011)].
- [23] J. H. Cobb *et al.* (Soudan 2 Collaboration), *Phys. Rev. D* **61**, 092002 (2000).
- [24] R. Abbasi *et al.* (IceCube Collaboration), *Phys. Rev. D* **84**, 082001 (2011).
- [25] T.-P. Li and Y.-Q. Ma, *Astrophys. J.* **272**, 317 (1983).
- [26] J. Braun, J. Dumm, F. D. Palma, C. Finley, A. Karle, and T. Montaruli, *Astropart. Phys.* **29**, 299 (2008).
- [27] W. W. Kinnison *et al.*, *Phys. Rev. D* **25**, 2846 (1982).
- [28] M. Wascko, Ph. D. thesis, University of California, Riverside, 2001.
- [29] J. Blumenthal, Diploma thesis, Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, 2011.
- [30] R. Reimann, Diploma thesis, Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, 2011.