

# Antineutrino Oscillations in the Atmospheric Sector

Thesis by  
Alexander I. Himmel

In Partial Fulfillment of the Requirements  
for the Degree of  
Doctor of Philosophy



California Institute of Technology  
Pasadena, California

2011  
(Defended May 10<sup>th</sup>, 2011)



# Acknowledgments

While it is my name on the cover, the credit for this thesis must be shared with the many people who made it possible.

First, I would like to thank my advisor, Harvey Newman, who has ably guided me throughout my graduate studies. He was always ready with keen insights and helpful advice, both in physics and in the human enterprise of doing physics. In addition, I had the good fortune to be guided by two talented postdocs, Caius Howcroft and Ryan Patterson – they taught me more than I knew I had to learn about being an experimental physicist. I would also like to thank the rest of the Caltech group, Leon Mualem, Jason Trevor, and Mhair Orchanian, for all their help in these last few years.

I would also like to thank my MINOS colleagues who contributed so much to these analyses. In particular, I would like to thank the antineutrino group and its leaders Jeff Hartnell, Justin Evans, and Donna Naples. You were always ready with the big picture view that put my work into context. I would also like to thank the beam group, in particular Sacha Kopp and Zarko Pavlovic who got me started simulating the beam and Alfons Weber, Mary Bishai, and Phil Adamson who helped make sure the simulation was ‘correct.’

My study of experimental physics did not begin here, and I would be remiss not to acknowledge the support and guidance I received before joining Caltech. Specifically, I would like to thank Chris Hill, Dave Stuart, Joel Goldstein, and Lenny Spiegel who mentored me over many summers at Fermilab, and Richard Haskell who guided me through my undergraduate career at Harvey Mudd. Of course, I must also thank my first and continuing mentor in physics, my uncle Joe Incandela, who started me on this path in the first place.

Finally, I thank my family, my wife Tiffany, my parents, Scott and Deborah, and my sisters Chloe and Isabella. You have expected nothing less than excellence from me but have also given me the confidence and wherewithal to achieve it. I would not be here without you.



---

# Abstract

This thesis presents measurements of the oscillations of muon antineutrinos in the atmospheric sector, where world knowledge of antineutrino oscillations lags well behind the knowledge of neutrinos, as well as a search for  $\nu_\mu \rightarrow \bar{\nu}_\mu$  transitions. Differences between neutrino and antineutrino oscillations could be a sign of physics beyond the Standard Model, including non-standard matter interactions or the violation of  $\mathcal{CPT}$  symmetry. These measurements leverage the sign-selecting capabilities of the magnetized steel-scintillator MINOS detectors to analyze antineutrinos from the NuMI beam, both when it is in neutrino-mode and when it is in antineutrino-mode. Antineutrino oscillations are observed at  $|\Delta\bar{m}_{\text{atm}}^2| = (3.36_{-0.40}^{+0.46}(\text{stat}) \pm 0.06(\text{syst})) \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860_{-0.12}^{+0.11}(\text{stat}) \pm 0.01(\text{syst})$ . The oscillation parameters measured for antineutrinos and those measured by MINOS for neutrinos differ by a large enough margin that the chance of obtaining two values as discrepant as those observed is only 2%, assuming the two measurements arise from the same underlying mechanism, with the same parameter values. No evidence is seen for neutrino-to-antineutrino transitions.



# Contents

<b>Acknowledgments</b>	<b>iii</b>
<b>Abstract</b>	<b>v</b>
<b>List of Tables</b>	<b>xi</b>
<b>List of Figures</b>	<b>xii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Physics of Neutrinos and Antineutrinos</b>	<b>3</b>
2.1 Neutrino History . . . . .	3
2.1.1 The Beginning of the Neutrino . . . . .	3
2.1.2 Early Neutrino Experiments . . . . .	4
2.1.3 Neutrinos and Symmetry . . . . .	6
2.2 Neutrinos in the Standard Model . . . . .	7
2.2.1 The Symmetry Group . . . . .	7
2.2.2 Gauge Bosons and their Masses . . . . .	9
2.2.3 Fermion Masses . . . . .	11
2.2.4 Neutrino Interactions . . . . .	13
2.3 Neutrino Oscillations . . . . .	18
2.3.1 Neutrino Masses . . . . .	21
2.3.1.1 Dirac Mass . . . . .	21
2.3.1.2 Majorana Mass . . . . .	22
2.3.2 Oscillations in Vacuum . . . . .	25
2.3.3 The Mixing Matrix . . . . .	27
2.3.4 Two Neutrino Mixing . . . . .	31
2.3.5 Matter Effects . . . . .	32
2.3.6 Measurements . . . . .	34
2.4 Antineutrino Oscillations . . . . .	39

2.5	Differing Neutrino and Antineutrino Oscillations . . . . .	40
2.5.1	Non-standard Interactions . . . . .	40
2.5.2	$CPT$ Violation . . . . .	41
2.6	Neutrino-to-Antineutrino Transitions . . . . .	42
<b>3</b>	<b>The MINOS Experiment</b>	<b>45</b>
3.1	The NuMI Beam . . . . .	46
3.1.1	Making Neutrinos . . . . .	46
3.1.2	Making Antineutrinos . . . . .	51
3.1.3	MINOS Beam Data . . . . .	53
3.1.4	Beam Simulation with Flugg . . . . .	54
3.2	The MINOS Detectors . . . . .	57
3.2.1	A MINOS Plane . . . . .	59
3.2.2	The Magnetic Field . . . . .	61
3.2.3	Detector Readout and Data Acquisition . . . . .	63
3.3	Calibration . . . . .	65
3.4	Reconstruction . . . . .	68
3.4.1	Event topologies . . . . .	68
3.4.2	Tracks . . . . .	69
3.4.3	Showers . . . . .	71
3.5	Simulation . . . . .	72
<b>4</b>	<b>An Oscillation Analysis in Brief</b>	<b>73</b>
<b>5</b>	<b>Antineutrinos in a Neutrino Beam</b>	<b>75</b>
5.1	Selection . . . . .	77
5.2	Near Detector Spectrum . . . . .	81
5.3	Extrapolation (the Beam Matrix) . . . . .	82
5.4	Systematics . . . . .	86
5.4.1	Energy Reconstruction Systematics . . . . .	86
5.4.2	Background Systematics . . . . .	86
5.4.3	Extrapolation Systematics . . . . .	87
5.4.4	Downstream Parents . . . . .	88
5.4.5	Effect on the Analyses . . . . .	90
5.5	Far Detector Data . . . . .	92
5.5.1	Cross-checks . . . . .	95
5.6	Oscillation Analysis . . . . .	98



5.7	Transition Analysis . . . . .	98
5.8	Feldman-Cousins Method . . . . .	101
<b>6</b>	<b>Antineutrinos in an Antineutrino Beam</b>	<b>105</b>
6.1	Selection . . . . .	105
6.2	Near Detector Spectrum . . . . .	109
6.3	Systematics . . . . .	111
6.3.1	Energy Reconstruction Systematics . . . . .	111
6.3.2	Background Systematics . . . . .	112
6.3.3	Extrapolation Systematics . . . . .	112
6.3.4	Effect on the Analysis . . . . .	113
6.4	Far Detector Data . . . . .	115
6.5	Oscillation Results . . . . .	118
6.5.1	Comparison to Neutrinos . . . . .	120
<b>7</b>	<b>Conclusion</b>	<b>125</b>
<b>A</b>	<b>Meson and muon decay kinematics</b>	<b>129</b>
A.1	Probability that the Neutrino Reaches the Detectors . . . . .	129
A.2	Energy of the Neutrino that Reaches the Detectors . . . . .	131
A.3	Helicity Suppression in Meson Decays . . . . .	132
<b>B</b>	<b>The Decay Pipe Systematic</b>	<b>135</b>
<b>C</b>	<b>Feldman-Cousins Method for the Antineutrino Analyses</b>	<b>141</b>
C.1	Introduction . . . . .	141
C.2	The Feldman-Cousins Method . . . . .	141
C.3	Generating Pseudo-Experiments . . . . .	142
C.4	Applying Systematics . . . . .	144
C.5	Feldman-Cousins Surfaces . . . . .	147
C.5.1	Oscillations . . . . .	147
C.5.2	Transitions . . . . .	147
<b>D</b>	<b>Detector Edge Study</b>	<b>151</b>
D.1	Introduction . . . . .	151
D.2	Plane Types . . . . .	152
D.3	U and V . . . . .	152
D.4	Physical Measurements . . . . .	156
D.5	Monte Carlo “Measurements” . . . . .	158

D.6	Conclusions . . . . .	158
<b>E</b>	<b>Validation of the Reconstruction Upgrade</b>	<b>161</b>
E.1	The Charge Sign Problem . . . . .	162
E.2	The Charge Sign Fix . . . . .	162
E.3	Dogwood3 . . . . .	164
E.4	Ancillary Problems . . . . .	164
<b>F</b>	<b>Technical Aspects of Flugg</b>	<b>167</b>
F.1	Introduction . . . . .	167
F.2	How to Run . . . . .	168
F.2.1	Setup and Installation . . . . .	168
F.2.2	Running the Simulation . . . . .	168
F.2.3	Inside the Scripts . . . . .	169
F.3	The Code . . . . .	171
F.3.1	Overview . . . . .	171
F.3.2	Importance Weighting and Thresholds . . . . .	172
F.3.3	Particle History Tracking . . . . .	173
F.3.4	Fluka Source Files . . . . .	173
F.3.5	GEANT Source Files . . . . .	175
F.4	The Output Ntuples . . . . .	176
F.4.1	Available Files . . . . .	176
F.4.2	File Contents . . . . .	178
F.4.3	Useful Codes . . . . .	180
F.5	Validation . . . . .	182
F.5.1	The Focusing Peak . . . . .	182
F.5.2	Helium and Other Downstream Production . . . . .	183
F.5.3	The Near Detector $\nu_e$ Flux . . . . .	184
F.5.4	Low Angle Scatters . . . . .	190
F.5.5	Multiple Scattering in the Chase . . . . .	191
<b>G</b>	<b>The Water-Scintillator Detector</b>	<b>193</b>
	<b>Bibliography</b>	<b>197</b>

# List of Tables

2.1	Particles of the standard model . . . . .	8
2.2	Electroweak charges of standard-model fermions. . . . .	14
3.1	The decays that produce neutrinos and their frequency at the Near Detector. . . . .	51
5.1	Components of the Near-to-Far normalization systematic uncertainty. . . . .	87
6.1	Components of the Near-to-Far normalization systematic uncertainty. . . . .	112
B.1	The systematic errors on the Near Detector . . . . .	137
C.1	Table of systematics and how they are implemented. . . . .	145
D.1	The correspondence between plots in Figure D.5 and the modules in Figure D.2. . . . .	156
F.1	Required environmental variables for <code>g4numi_fluka.sh</code> . . . . .	169
F.2	Optional environmental variables for <code>g4numi_fluka.sh</code> . . . . .	169
F.3	Output files found in the directory of a successful job . . . . .	170
F.4	Describes how all of the Fluka spare tracking variables <code>ISPUSR[]</code> and <code>SPAUSR[]</code> are used. . . . .	174
F.5	Available files at FNAL. You can access them at <code>/minos/data/flux/flugg/</code> . . . . .	177
F.6	The entries stored in the neutrino ntuple files . . . . .	178
F.7	The entries stored in the neutrino ntuple files . . . . .	179
F.8	The entries stored in the hadron ntuple files . . . . .	179
F.9	The particle codes across the three schemes used in MINOS. . . . .	180
F.10	The material codes as defined by Gnumi and used in the fluxfiles, old and current. . . . .	181
F.11	The decay codes stored in <code>Ndecay</code> . . . . .	181



# List of Figures

2.1	Feynman diagrams of neutrino-quark interactions . . . . .	15
2.2	Neutrino vs. antineutrino interaction cross sections . . . . .	17
2.3	Zenith angle distributions from Super-Kamiokande . . . . .	19
2.4	Solar neutrino fluxes measured by SNO . . . . .	20
2.5	Feynman diagrams of double beta decay . . . . .	25
2.6	Measurement of the $Z^0$ mass resonance at LEP . . . . .	28
2.7	Oscillation parameters allowed by KamLAND and Solar experiments . . . . .	36
2.8	Atmospheric neutrino oscillation contours from MINOS and Super-Kamiokande . . . .	37
2.9	Antineutrino vs. neutrino contours circa 2008 . . . . .	39
3.1	Schematic views of the components of the MINOS experiment . . . . .	45
3.2	Cross-sectional view of the NuMI beam complex and MINOS Near Detector Hall . . .	46
3.3	Drawing of a NuMI target and its housing. . . . .	47
3.4	Drawing and picture of a NuMI Horn . . . . .	48
3.5	Beam focusing for neutrinos . . . . .	49
3.6	The $p_T$ vs. $p_Z$ distribution of the $\pi^\pm$ parents of neutrinos and antineutrinos . . . . .	49
3.7	The neutrino spectrum in various beam configurations . . . . .	50
3.8	Beam focusing for antineutrinos . . . . .	52
3.9	The event rate for $\nu_\mu$ and $\bar{\nu}_\mu$ in neutrino-mode and antineutrino-mode beams . . . . .	52
3.10	Neutrino vs. antineutrino production and interaction cross sections . . . . .	52
3.11	The protons-on-target accumulated per week during NuMI operations . . . . .	53
3.12	The ratio of the Near Detector spectrum with and without helium in the decay pipe .	55
3.13	Near Detector $\nu_\mu$ spectra with and without beam tuning . . . . .	56
3.14	The Near Detector $\bar{\nu}_\mu$ spectrum with and without beam tuning . . . . .	57
3.15	Arrangement of MINOS scintillator planes . . . . .	58
3.16	Pictures of the MINOS detectors . . . . .	58
3.17	Cut-away view of a single MINOS scintillator strip . . . . .	59
3.18	The arrangement of scintillator modules into a $U$ plane in the Far Detector . . . . .	60

3.19	The arrangement modules into the four types of Near Detector scintillator planes . . .	60
3.20	Maps of the field strengths in the detectors . . . . .	62
3.21	The readout of WLS fibers . . . . .	64
3.22	An example of a Near Detector snarl . . . . .	69
3.23	The three event topologies relevant to the antineutrino analysis . . . . .	70
5.1	Simulated Far Detector spectrum showing the effect of transitions and oscillations . .	75
5.2	The $p_T$ vs. $p_Z$ distribution of the $\pi^\pm$ parents of neutrinos and antineutrinos . . . . .	76
5.3	Contamination in the positive sample after preselection only . . . . .	78
5.4	Charge-sign selection variables $(q/p) / \sigma_{q/p}$ and $ \text{Relative Angle} - \pi $ . . . . .	78
5.5	The three variables that form the PDFs in the $DpID$ CC/NC separator . . . . .	79
5.6	CC/NC separation parameter, $DpID$ . . . . .	80
5.7	Efficiency and contamination after selection . . . . .	80
5.8	Reconstructed $\mu^+$ track momentum shower energy in the Near Detector . . . . .	81
5.9	The selected $\bar{\nu}_\mu$ energy spectrum in the Near Detector . . . . .	82
5.10	Diagram of the neutrino parents in the NuMI decay pipe . . . . .	83
5.11	The relationship between the energies of events in the Near and Far Detectors . . . .	84
5.12	The beam matrix . . . . .	85
5.13	The Near Detector spectrum and Far/Near ratio for the various antineutrino parents .	88
5.14	Spectra for parents produced in the decay pipe and upstream regions . . . . .	89
5.15	Total systematic error band on the Far Detector prediction . . . . .	90
5.16	Systematic shifts in the oscillation parameters . . . . .	91
5.17	Systematic shifts in the transition probability . . . . .	91
5.18	Far Detector $\bar{\nu}_\mu$ data spectrum . . . . .	93
5.19	CC/NC separation parameter, $DpID$ , in the Far Detector . . . . .	93
5.20	Charge-sign selection variable $(q/p) / \sigma_{q/p}$ in the Far Detector . . . . .	94
5.21	Charge-sign selection variable $ \text{Relative Angle} - \pi $ in the Far Detector . . . . .	94
5.22	Selected antineutrino event vertex and end positions . . . . .	95
5.23	The distribution of the number of selected events in 10,000 fake experiments . . . . .	96
5.24	Effect of shifting the Near Detector selection cuts . . . . .	96
5.25	Far Detector antineutrino events vs. time . . . . .	97
5.26	Near Detector antineutrino events vs. time . . . . .	97
5.27	Antineutrino oscillation contours . . . . .	99
5.28	$ \Delta\bar{m}_{\text{atm}}^2 $ exclusion at maximal mixing . . . . .	100
5.29	Transitions exclusion curves . . . . .	100
5.30	Example $-2\Delta \ln L$ distributions, one near and one distant from a physical boundary .	101

5.31	Feldman-Cousins 90% coverage surface for oscillations . . . . .	102
5.32	Feldman-Cousins curves for transitions . . . . .	103
6.1	$q/p$ distribution of selected events before charge sign selection in the Near Detector . . . . .	106
6.2	Distribution of the 4 CC/NC separator input variables . . . . .	108
6.3	CC/NC separation parameter . . . . .	108
6.4	Efficiency and purity of the antineutrino selector . . . . .	109
6.5	Reconstructed $\mu^+$ track momentum and shower energy in the Near Detector . . . . .	110
6.6	Reconstructed Near Detector spectrum . . . . .	110
6.7	Total systematic error band on the Far Detector prediction . . . . .	114
6.8	Systematic shifts in the oscillation parameters . . . . .	114
6.9	The selected $\bar{\nu}_\mu$ energy spectrum at the Far Detector . . . . .	115
6.10	$q/p$ distribution of selected Far Detector events . . . . .	116
6.11	Inelasticity distribution of selected $\bar{\nu}_\mu$ and $\nu_\mu$ interactions . . . . .	116
6.12	Event vertices of selected $\bar{\nu}_\mu$ 's and $\nu_\mu$ 's . . . . .	117
6.13	Track end positions of selected $\bar{\nu}_\mu$ 's and $\nu_\mu$ 's . . . . .	117
6.14	Far Detector antineutrino spectrum . . . . .	118
6.15	Antineutrino oscillation contours compared to a global fit . . . . .	119
6.16	Antineutrino oscillation contours from both neutrino- and antineutrino-mode analyses . . . . .	120
6.17	Far Detector antineutrino spectrum with the prediction using $\nu_\mu$ parameters . . . . .	121
6.18	Antineutrino and neutrino contours . . . . .	121
6.19	Example $\nu_\mu$ and $\bar{\nu}_\mu$ fake experiment . . . . .	122
6.20	The empirical distribution of $-2\Delta \ln L$ 's between 4-parameter and 2-parameter fits . . . . .	123
7.1	Possible contours with $\bar{\nu}_\mu$ -like future data . . . . .	125
7.2	Possible contours with $\nu_\mu$ -like future data . . . . .	126
B.1	Production vertices of antineutrino parents . . . . .	135
B.2	The Near and Far Detector antineutrino spectra by parent production region . . . . .	136
B.3	The Near Detector spectrum with no systematics applied . . . . .	137
B.4	The Near Detector systematic error bands . . . . .	138
B.5	The Near Detector spectrum with all systematics applied . . . . .	138
B.6	The Far-over-Near ratios for the two calculated systematic errors . . . . .	139
C.1	Example $-2\Delta \ln L$ distributions for two transition probabilities . . . . .	142
C.2	Two Feldman-Cousins spectra compared with the prediction and the simulation . . . . .	143
C.3	Two Feldman-Cousins methods compared with each other . . . . .	146
C.4	The curve used to parameterize the antineutrino cross section uncertainty . . . . .	146

C.5	Two methods for applying the decay pipe systematic . . . . .	147
C.6	Feldman-Cousins 90% coverage surface for oscillations . . . . .	148
C.7	Feldman-Cousins curves for transitions . . . . .	149
C.8	The upper and lower 90% confidence intervals for transitions of 10,000 fake experiments	150
D.1	The track end point in $x$ in meters for positive tracks in the Near Detector . . . . .	151
D.2	The types of planes in the Near Detector with their component modules. . . . .	152
D.3	Figure D.1 broken up into full and partial U and V planes . . . . .	153
D.4	V and U planes broken up into calorimeter and spectrometer . . . . .	153
D.5	Position along a strip vs. position perpendicular to the strips . . . . .	154
D.6	A proposed explanation of the effects seen in Figure D.5 . . . . .	155
D.7	The top and bottom of the detector . . . . .	155
D.8	Measurements taken of the example plane in the MINOS surface building. . . . .	156
D.9	Figure D.5 marked with physical measurements . . . . .	157
D.10	True vertices in carbon . . . . .	158
D.11	True Monte Carlo vertex positions . . . . .	159
D.12	Figure D.1 with the DigiScintHits locations overlaid . . . . .	159
E.1	Event view of a catastrophic charge mis-id . . . . .	161
E.2	Distributions for events where the biased and unbiased fitters disagree . . . . .	163
E.3	Event-by-event comparison of track length and energy resolution . . . . .	165
E.4	Track length and $(q/p) / \sigma_{q/p}$ distributions when the reconstruction versions disagree .	165
F.1	The focusing peak for neutrinos with vertices in the chase ( $V_z < 45$ m) . . . . .	182
F.2	The focusing peak for neutrinos with vertices in the decay pipe ( $V_z > 45$ m) . . . . .	183
F.3	The low energy spectrum for $\nu_\mu$ 's whose parents were produced outside the target . .	183
F.4	The low energy spectrum for $\bar{\nu}_\mu$ 's whose parents were produced outside the target. . .	184
F.5	The ratio of the Near Detector spectrum with and without helium in the decay pipe .	184
F.6	Near Detector $\nu_e$ and $\bar{\nu}_e$ energy spectrum . . . . .	186
F.7	Energy spectrum for neutrinos from $\mu^\pm$ decay . . . . .	186
F.8	The Near Detector weighted $V_z$ distribution for neutrinos from muon decay . . . . .	186
F.9	The importance (non-detector) weighted $V_z$ distribution for neutrinos from muon decay	187
F.10	The Near Detector weighted $V_z$ distribution for $\nu_\mu$ 's . . . . .	187
F.11	The importance (non-detector) weighted $V_z$ distribution for $\nu_\mu$ 's. . . . .	187
F.12	The projected $x$ -position the neutrino parent would have passed through at $z = 45$ m	188
F.13	The ratio of integrated $\mu^+$ Near Detector fluxes from Flugg and Gnumi . . . . .	188
F.14	The ratio of integrated $\mu^-$ Near Detector fluxes from Flugg and Gnumi . . . . .	189



---

F.15	The <i>divy</i> vs. <i>divx</i> positions of parents produced in the target . . . . .	190
F.16	Scattering the chase . . . . .	191
G.1	The water-scintillator detector prototype at Caltech . . . . .	193
G.2	The water-scintillator detector readout and triggering . . . . .	194
G.3	Histograms of the detector response in each of the eight channels . . . . .	195
G.4	The summed response of all channels to a through-going muon trigger . . . . .	195



# Chapter 1

## Introduction

Neutrinos have fascinated scientists throughout their history. Since they interact only via the weak force, they were long considered undetectable. This very property, which makes them so difficult to study, also means they can provide unique insights into the behavior of the weak interactions. Neutrinos have played central roles in some of the most surprising discoveries in particle physics, such as the violation of parity symmetry. The discovery of neutrino oscillations, and with it neutrino mass, provided the first glimpses of physics beyond the Standard Model. The relationship between the neutrino and its antiparticle is a key piece of this unfolding mystery.

This thesis begins with a look at the history of the neutrino in Chapter 2, following it from its early days as a ‘desperate remedy’ to save the conservation of energy through the key experiments that determined its existence and properties. The neutrino was at the heart of understanding the weak force and its unexpected symmetry properties, and thus was a key element in creating the Standard Model of particle physics. The discovery of neutrino mass in 1998 once again put neutrinos at the forefront of science, providing the first evidence of physics beyond the Standard Model.

However, despite these successes, many mysteries about the neutrino remain. Among those mysteries is the relationship between the neutrino and its even more difficult to study partner, the antineutrino. This thesis looks in depth at the oscillation properties of antineutrinos, which might once again revolutionize our understanding of fundamental particles. If neutrinos and antineutrinos oscillate differently, it is a signature of more physics not predicted by the Standard Model. It could be a sign of new particles, or new interactions with matter, or even a violation of the most fundamental symmetries of quantum field theory: Lorentz and  $CPT$  invariance.

The MINOS experiment, with its intense beam and sign-selecting detectors described in Chapter 3, is uniquely suited to the study of neutrinos and antineutrinos. The detectors are optimized to study muon tracks, the signature of charged-current muon neutrinos and antineutrinos and the key feature that allows neutrinos and antineutrinos to be distinguished. Chapter 4 describes in a general way how the data taken at the two MINOS detectors can be leveraged to measure the oscillation parameters of antineutrinos in a way robust against mismodelling of the antineutrino flux

and cross-sections and other systematic uncertainties.

The first measurement of the antineutrino oscillation parameters, described in Chapter 5, was made with the small (7%) antineutrino component of the NuMI neutrino beam. This measurement demonstrated MINOS's ability to identify antineutrino events despite large backgrounds, but its sensitivity was limited by the properties of the antineutrino flux. In addition to being few in number, the antineutrinos in the neutrino beam are typically at energies too high to be sensitive to oscillation parameters close to the ones already measured for neutrinos. However, the limited low-energy intrinsic antineutrino component of the flux, combined with the ability to select a pure sample of antineutrinos despite large backgrounds, does provide an opportunity to search for the transitions of the much more numerous neutrinos, which are known to be disappearing, into antineutrinos.

A precise measurement of the antineutrino oscillation parameters around the atmospheric sector requires not just more antineutrinos but also requires them to be at the correct energies. Such a measurement is possible at MINOS thanks to the adaptability of the NuMI beam, which can be retuned as an antineutrino beam as first proposed by the Caltech MINOS group, producing a high flux of antineutrinos at an adjustable energy. The first direct, precision measurement of the antineutrino oscillation parameters in the atmospheric sector was made with data taken while running in antineutrino mode, and is described in Chapter 6.

The study of antineutrino oscillations, in the tradition of the study of neutrinos, reveals hints of physics just beyond our reach. As MINOS and its successors accumulate more antineutrino data, they may once again revolutionize particle physics.

## Chapter 2

# Physics of Neutrinos and Antineutrinos

### 2.1 Neutrino History

#### 2.1.1 The Beginning of the Neutrino

The neutrino began as a solution to a problem. When radioactivity was first discovered at the end of the 19<sup>th</sup> century, the  $\beta$ -decay process was believed to be solely the emission of an electron from a radioactive nucleus. The electron, then, was expected to emerge at a fixed energy corresponding to the change in binding energy of the nucleus. However, in 1914 James Chadwick showed that the energy spectrum of the emitted electron was continuous, not monoenergetic [1]. This apparent violation of the conservation of energy was understandably troubling to physicists at the time. One seemingly promising explanation was that the electron was losing energy into the medium containing the radioactive nuclei, but this hypothesis was disproved by measurements performed in a calorimeter by Ellis and Wooster in the 1920's [2]. Some physicists, including the distinguished Neils Bohr, went as far as suggesting that energy was not conserved in individual decays [3].

Amidst the confusion, Wolfgang Pauli wrote an open letter [4] to a 1930 conference on radioactivity in which he proposed “a desperate remedy to save the exchange theorem of statistics and the law of conservation of energy.” Pauli’s remedy was to posit the existence of a new particle he dubbed the ‘neutron’ that was being emitted simultaneously with the electron during  $\beta$  decay. This new particle needed to be electrically neutral, have spin  $\frac{1}{2}$ , and have a mass less than 1% of the proton mass (in order to agree with the end point of the spectrum). Pauli was so hesitant about his speculation, believing the particle should have been observed if it were being emitted, that he did not publish his idea until 4 years later. Meanwhile, the modern neutron was discovered by Chadwick in 1932 [5] (the same Chadwick who measured the  $\beta$ -decay spectrum), but it was immediately clear that this neutral particle was too heavy to be Pauli’s neutron.

Belief in the neutrino increased significantly after 1933 when Enrico Fermi included it in the first successful theory of  $\beta$ -decay, constructed analogous to the earlier theory of electromagnetic interactions [6]. Fermi's theory used both Chadwick's and Pauli's neutrons: he adopted the proton-neutron theory of the atom and, to avoid confusion, renamed Pauli's neutron the *neutrino*, or little neutron in Italian. In this theory, the electron and neutrino<sup>1</sup> are produced during the transition of a neutron into a proton,

$$n \rightarrow p + e^{-} + \nu. \quad (2.1)$$

This theory, along with its generalization by Gamow and Teller in 1936 [9], was able to describe all  $\beta$ -decay data that had been collected. Pauli's major concern, that the neutrino should already have been observed if it existed, was addressed in 1934 when Hans Bethe and Rudolf Peierls showed that the cross section for the neutrino to interact with a nucleus was billions of times smaller than the equivalent cross section for the electron [10]. The cross section was so small that for many years it was considered an undetectable particle.

### 2.1.2 Early Neutrino Experiments

By the middle of the 1950's the neutrino had not been detected directly, but there was some indirect experimental evidence of its existence. When pion and muon decays were examined using photographic emulsions, the decay products could be seen continuing on at large angles relative to the initial parent direction [11]. These observations suggested that an additional particle was being emitted in the decays but that this particle was not being observed in the emulsion.

The first direct observation of the neutrino was in an experiment performed by Frederick Reines and Clyde Cowan in 1956 [12]. Their detector consisted of a water target with dissolved  $\text{CdCl}_2$  sandwiched between tanks of liquid scintillator instrumented with photomultiplier tubes. The antineutrinos were detected via inverse  $\beta$ -decay:

$$\bar{\nu} + p \rightarrow e^{+} + n. \quad (2.2)$$

The neutrino detector was set up close to the Savannah River nuclear reactor which provided a large flux of 'neutrinos,' though they are now known to be electron antineutrinos. Neutrino events were identified by a distinct pattern of activity in the detector. First, the positron would quickly annihilate with an electron in the water, producing two 511 keV photons moving in opposite directions that would be picked up in coincidence in the scintillator above and below the target. Then, several microseconds later, the neutron would capture on the dissolved cadmium, releasing more photons

---

<sup>1</sup>This particle is now known to be a  $\bar{\nu}_e$ , but at the time neither antineutrinos nor neutrino flavors had even been considered. In fact, the first antimatter particle of any kind, the positron, had only just been discovered in that same year, 1933 [7], after being posited by Dirac in 1928 [8].

that would be picked up in delayed coincidence with the positron signature. Importantly, they demonstrated that the observed events were actually neutrinos from the reactor by showing that the rate of these interactions dropped to nearly nothing when the reactor was turned off. Reines received the Nobel Prize in 1995 for the discovery.

Shortly thereafter, Ray Davis also performed an experiment near a nuclear reactor to try and observe the process:



which had been first proposed by Bruno Pontecorvo in the mid-1940's [13]. Note the opposite sign on the electron in Equation 2.3 as opposed to Equation 2.2. Davis found no evidence of this reaction, demonstrating that while  $\bar{\nu}$ 's from reactors could produce positrons via inverse beta decay, they could not produce electrons, demonstrating that neutrinos are distinct from antineutrinos [14]. This distinction suggested that lepton number (which is opposite in sign for leptons and their antiparticles) is a conserved quantity.

Two years later, in 1958, Bruno Pontecorvo first suggested that neutrinos might have small masses and consequently oscillations similar to those seen in the neutral kaon system [15]. However, Pontecorvo's oscillations were between neutrinos and antineutrinos – as far as anyone knew at the time, there was only a single type of neutrino. The first conclusive evidence that there was more than one flavor of neutrino came in 1962 from an experiment at the Brookhaven AGS performed by Leon Lederman, Melvin Schwartz, and Jack Steinberger [16]. The neutrinos studied in this experiment came from the first high-energy neutrino beam. Protons were accelerated to 15 GeV and directed onto a beryllium target, producing numerous  $\pi^{\pm}$ . The pions were then allowed to decay in a 21 m channel where the vast majority produced neutrinos via

$$\pi^{+} \rightarrow \mu^{+} + \nu \quad (2.4)$$

$$\pi^{-} \rightarrow \mu^{-} + \bar{\nu}. \quad (2.5)$$

The alternative process, where  $e^{\pm}$  is produced instead of  $\mu^{\pm}$  is suppressed by a factor of  $10^4$  (see explanation in Appendix A.3). At the end of the decay channel was an absorber that stopped all the charged particles before they reached the aluminum spark chamber behind it. The neutrinos, of course, passed through the absorber to the detector. If there was only a single neutrino, then approximately half the interactions in the spark chamber would produce electrons and half would produce muons:

$$\nu + N \rightarrow \begin{cases} \mu^{-} + X \\ e^{-} + X \end{cases} . \quad (2.6)$$

However, if there were two neutrino ‘flavors,’  $\nu_\mu$  and  $\nu_e$ , only muons would be produced:

$$\nu_\mu + N \rightarrow \mu^- + X \quad (2.7)$$

$$\nu_\mu + N \not\rightarrow e^- + X, \quad (2.8)$$

which is precisely what was observed at the Brookhaven experiment. Lederman, Schwartz, and Steinberg received the 1988 Nobel Prize for this discovery, which demonstrated the doublet structure of the leptons and that particles are arranged into families.

### 2.1.3 Neutrinos and Symmetry

Around the same time, an important property of the neutrino, and the weak force in general, was discovered at Columbia. Two theoretical physicists, Tsung-Dao Lee and Chen-Ning Yang, were trying to understand the decays of two mesons, the  $\tau^+$  and the  $\theta^+$ , that appeared to be identical in every regard except that they underwent weak decays into states of opposite parity.<sup>2</sup> Lee and Yang made the revolutionary proposal that these two mesons were one and the same but that in this particular type of decay, the weak decay, parity was violated [17]. Up until this time, physicists had taken it as given that parity was a fundamental symmetry of the universe.

Lee and Yang convinced an experimentalist at Columbia, Chien-Shiung Wu, to test their hypothesis experimentally. Wu examined the  $\beta$ -decays of a polarized  $^{60}\text{Co}$  source and found a large asymmetry in the direction of the emitted  $\beta$ 's that could only be attributed to parity-violation in the weak decay [18]. The asymmetry was so large that it required the parity violation to be maximal. Shortly after the Wu experiment, two more experimentalists, Leon Lederman and Richard Garwin followed another of Lee and Yang's suggestions and found evidence of parity violation in the  $\pi \rightarrow \mu \rightarrow e$  decay chain [19]. The experimental evidence led Lee and Yang to propose the two-component theory of the neutrino where the neutrino had only one possible helicity state and the antineutrino had only the opposite helicity [20]. This idea was the first step towards the modern  $V - A$  theory of the weak interaction from the generic theory which included all possible operators that was proposed by Fermi, Gamow, and Teller in the 1930's. A year later, in 1958, an experiment by M. Goldhaber, L. Grodzins, and A. W. Sunyar showed that the neutrino is the left-handed helicity state [21]. The combined evidence of that experiment and a series of other experiments looking at nuclear recoil [22] and polarization in another type of beta decay [23, 24, 25] cemented the  $V - A$  nature of the theory.

The story of parity violation does not end, there however. Physicists were troubled by the violation of a symmetry that they had assumed was fundamental. In 1957, Lev Landau proposed that while parity alone was violated, the combination of parity and charge-conjugation,  $\mathcal{CP}$ , could

---

<sup>2</sup>It is now known that the  $\tau^+$  and the  $\theta^+$  are, indeed, the same particle, the  $K^+$  meson



still be a good symmetry [26]. However, the respite was brief. In 1964, James Cronin and Val Fitch showed that even  $\mathcal{CP}$  was violated by the decay of kaons [27]. They showed that the  $K_L^0$ , which at the time was believed to be the  $\mathcal{CP}$ -odd eigenstate, did have a small but non-zero branching ratio,  $(2.0 \pm 0.4) \times 10^{-3}$ , for decaying to 2 pions, a  $\mathcal{CP}$ -even decay channel.<sup>3</sup> Unlike parity violation, which is maximal, the violation of  $\mathcal{CP}$  is small.

After  $\mathcal{CP}$  violation was demonstrated, physicists were only left with the combination of  $\mathcal{CP}$  and time reversal,  $\mathcal{CPT}$ . It was shown by several physicists in the 1950's<sup>4</sup> that  $\mathcal{CPT}$  must be a good symmetry in order for a theory to be Lorentz-invariant. A consequence is that any search for  $\mathcal{CPT}$  violation is also a search for Lorentz violation, since the two come hand-in-hand. Tests of  $\mathcal{CPT}$  violation are discussed further in Section 2.5.2.

## 2.2 Neutrinos in the Standard Model

By the 1960's, the theory of the neutrino began to take on its modern form. Previous theories of the weak interaction, like the one written by Fermi in the 1930's, were “unrenormalizable”: they produced divergent predictions at anything but the lowest order of perturbation theory. Sheldon Glashow, Steven Weinberg, and Abdus Salam wrote the first renormalizable theory of the weak interaction by unifying it with the only known renormalizable theory: the theory of electromagnetic interactions [32, 33, 34].

Today, neutrinos are included as part of the Standard Model which describes the interactions of the 17 fundamental particles: the 12 spin- $\frac{1}{2}$  fermions and the 5 spin-0 and spin-1 bosons [35]. These particles are shown in Table 2.1. The Standard Model begins as a locally-gauge invariant theory of massless fields, using gauge group  $SU(3)_{\text{color}} \otimes SU(2)_L \otimes U(1)_Y$ , unifying the strong force ( $SU(3)_{\text{color}}$ ) with Glashow, Weinberg, and Salam's electroweak theory,  $(SU(2)_L \otimes U(1)_Y)$ . Particle masses are introduced by the Higgs mechanism of spontaneous symmetry breaking: the initial degrees of freedom associated with massless fields can be reorganized to produce massive fields thanks to a non-zero vacuum value of the Higgs field [36].

### 2.2.1 The Symmetry Group

Since neutrinos interact only via the weak force, I will leave off discussion of the full Standard Model and focus on the GWS electroweak part of the theory, following notation and procedures adapted from [37] and [38].  $SU(2)$  is chosen as the symmetry group of the weak interaction since the handedness of the weak interaction can be naturally included by identifying the left-handed fermion fields as doublets and the right-handed fields as singlets (hence the subscript  $_L$ ). The representations are,

<sup>3</sup>In the paper, they refer to the  $K_L^0$  as  $K_2^0$ , the  $\mathcal{CP}$ -odd eigenstate, since they were believed to be equivalent.

<sup>4</sup>First J. Schwinger in 1951 [28], followed by G. Lüders [29], W. Pauli [30], and J.S. Bell [31].

Fermions				
	Name	Charge	Spin	Mass (MeV)
Quarks	$u$	+2/3	1/2	1.7 – 3.3
	$c$	+2/3	1/2	$(1.27^{+0.07}_{-0.09}) \times 10^3$
	$t$	+2/3	1/2	$(172 \pm 1.58) \times 10^3$
	$d$	-1/3	1/2	4.1 – 5.8
	$s$	-1/3	1/2	$101^{+29}_{-21}$
	$b$	-1/3	1/2	$(4.19^{+0.18}_{-0.06}) \times 10^3$
Leptons	$e$	-1	1/2	$0.51 \pm 1.3 \times 10^{-8}$
	$\mu$	-1	1/2	$105.7 \pm 3.8 \times 10^{-6}$
	$\tau$	-1	1/2	$1,776.8 \pm 0.16$
	$\nu_e$	0	1/2	
	$\nu_\mu$	0	1/2	$< 2 \times 10^{-6}$
	$\nu_\tau$	0	1/2	

Bosons				
Name	Charge	Spin	Mass (GeV)	Force
$\gamma$	0	1	0	EM
$W^\pm$	$\pm 1$	1	$80.40 \pm 0.02$	Weak
$Z^0$	0	1	$91.19 \pm 0.002$	Weak
$g$	0	1	0	Strong
$H^0$	0	0	$> 114$ with $158 - 173$ excluded	

Table 2.1: Particles of the standard model. The particle masses are as currently measured and are obtained from [35]. The masses of the quarks are ‘current-quark masses,’ determined in a mass-independent subtraction scheme. The Higgs mass limits come from [35].

for the lepton doublets ( $E$ ) and singlets ( $e$ ),

$$E_{lL} = \begin{pmatrix} \nu_l \\ l^- \end{pmatrix}_L, \quad e_{lR} = l_R, \quad \text{where } l = \{e, \mu, \tau\}, \quad (2.9)$$

for the quark doublets ( $Q$ ),

$$Q_{uL} = \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad Q_{cL} = \begin{pmatrix} c \\ s \end{pmatrix}_L, \quad Q_{tL} = \begin{pmatrix} t \\ b \end{pmatrix}_L \quad (2.10)$$

and for the quark singlets ( $q$ ),

$$\begin{aligned} q_{uR} &= u_R, & q_{cR} &= c_R, & q_{tR} &= t_R, \\ q_{dR} &= d_R, & q_{sR} &= s_R, & q_{bR} &= b_R. \end{aligned} \quad (2.11)$$

Note that there are no right-handed neutrinos since the massless neutrinos in the GWS theory are maximally parity-violating.

$SU(2)_L$  alone, however, does not have a massless gauge boson (i.e. the photon), so in order

to unify the weak force with the electromagnetic force, the symmetry group must be expanded to include U(1) as well. The theory needs to be locally invariant under rotations of the combined symmetry group. A generic, local rotation under  $SU(2)_L \otimes U(1)$  of a field  $\psi$  requires four vector gauge fields, three from  $SU(2)$ ,  $\mathbf{A}_\alpha$ , and one from  $U(1)$ ,  $B_\alpha$ . This transformation has the form

$$\psi \rightarrow e^{ig\boldsymbol{\tau}\cdot\mathbf{A}_\alpha} e^{ig'\frac{1}{2}B_\alpha}\psi, \quad (2.12)$$

where  $g$  and  $g'$  are two different coupling constants corresponding to the two component symmetry groups and the  $\boldsymbol{\tau}$  matrices are  $SU(2)$  representations.

### 2.2.2 Gauge Bosons and their Masses

Before addressing the interactions, let us begin by examining the gauge bosons. As written, the theory now contains four massless gauge bosons, but these bosons are not observed in nature. The massive bosons which do exist can be introduced via the Higgs mechanism of spontaneous symmetry breaking [36]. Posit the existence of a Higgs field which is an  $SU(2)$  doublet with  $U(1)$  charge  $+\frac{1}{2}$ :

$$\phi(x) = \begin{pmatrix} \phi_+ \\ \phi_0 \end{pmatrix}, \quad (2.13)$$

where  $\phi_+$  is a scalar complex field of charged particles and  $\phi_0$  the same, but for neutral particles.

The free Lagrangian for this field is given by:

$$\mathcal{L} = D_\alpha\phi^\dagger D^\alpha\phi - V(\phi^\dagger\phi). \quad (2.14)$$

where the partial derivative  $\partial_\alpha$  has been replaced by the covariant derivative

$$D_\alpha\phi = \left( \partial_\alpha + ig\boldsymbol{\tau}\cdot\mathbf{A}_\alpha + ig'\frac{1}{2}B_\alpha \right) \phi, \quad (2.15)$$

which makes the Lagrangian invariant under the transformation given in Equation 2.12. The potential  $V(\phi^\dagger\phi)$  is the specially chosen [39] ‘‘mexican hat’’ potential given by

$$V(\phi^\dagger\phi) = -\mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2 = \lambda\left(\phi^\dagger\phi - \frac{\mu^2}{2\lambda}\right)^2 - \frac{\mu^4}{4\lambda}, \quad (2.16)$$

where  $\mu^2$  and  $\lambda$  are positive constants. Note that the characteristic feature of this potential is that it reaches a minimum not at zero but when

$$(\phi^\dagger\phi)_{\text{vac}} = \frac{v^2}{2}, \quad v^2 = \frac{\mu^2}{\lambda}. \quad (2.17)$$

From the conservation of charge, we know that  $\phi_+$  must have a vacuum expectation of 0, and so it will be neglected from here on. Thus, the vacuum expectation value of the field can be written as

$$\phi_{\text{vac}} = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}, \quad (2.18)$$

and the complex doublet  $\phi(x)$  as

$$\phi(x) = e^{i\frac{1}{2}\boldsymbol{\tau}\cdot\boldsymbol{\theta}(x)} \begin{pmatrix} 0 \\ \frac{v+H(x)}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} 0 \\ \frac{v+H(x)}{\sqrt{2}} \end{pmatrix}, \quad (2.19)$$

where  $\boldsymbol{\theta}(x)$  and  $H(x)$  are real functions. In the final step, the leading exponential has been removed by taking advantage of local gauge invariance to choose a convenient (in this case ‘unitary’) gauge to work in.

With this vacuum field and choice of gauge, the Lagrangian in Equation 2.14 becomes

$$\mathcal{L} = \frac{1}{2}\partial_\alpha H \partial^\alpha H + \phi^\dagger \left( \frac{g}{2}\boldsymbol{\tau} \cdot \mathbf{A}_\alpha + \frac{g'}{2}B_\alpha \right) \left( \frac{g}{2}\boldsymbol{\tau} \cdot \mathbf{A}^\alpha + \frac{g'}{2}B^\alpha \right) \phi - V. \quad (2.20)$$

We can explicitly evaluate this expression using the field in Equation 2.19 and identifying the  $\boldsymbol{\tau}$  matrices as  $\frac{1}{2}$  times the Pauli matrices.<sup>5</sup> This gives

$$\phi^\dagger \phi = \frac{1}{2}(v+H)^2 \quad (2.21)$$

$$V(\phi^\dagger \phi) = \frac{\lambda}{4}(2vH+H^2)^2 \quad (2.22)$$

$$\boldsymbol{\tau} \cdot \mathbf{A}_\alpha \boldsymbol{\tau} \cdot \mathbf{A}^\alpha = \mathbf{A}_\alpha \cdot \mathbf{A}^\alpha \quad (2.23)$$

$$\phi^\dagger \boldsymbol{\tau} \cdot \mathbf{A}_\alpha B^\alpha \phi = -\frac{1}{2}(v+H)^2 A_\alpha^3 B^\alpha \quad (2.24)$$

where only the third component of  $\mathbf{A}_\alpha$  appears in Equation 2.24 since only  $\sigma^3$  has diagonal elements. The Lagrangian thus becomes

$$\begin{aligned} \mathcal{L} = & \frac{1}{2}\partial_\alpha H \partial^\alpha H + \frac{g^2}{8}(v+H)^2 \mathbf{A}_\alpha \cdot \mathbf{A}^\alpha + \frac{g'^2}{8}(v+H)^2 B_\alpha B^\alpha \\ & - \frac{g g'}{4}(v+H)^2 A_\alpha^3 B^\alpha - \frac{\lambda}{4}(2vH+H^2)^2. \end{aligned} \quad (2.25)$$

Insight can be gained into the physics of the Lagrangian in Equation 2.25 by making the following

---

<sup>5</sup>We have chosen the representation  $\sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ ,  $\sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ ,  $\sigma^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ .

substitutions:

$$W_\alpha^{(\dagger)} = W_\alpha^\pm = \frac{1}{\sqrt{2}}(A_\alpha^1 \mp iA_\alpha^2) \quad (2.26)$$

$$Z_\alpha = \frac{g}{\sqrt{g^2 + g'^2}}A_\alpha^3 - \frac{g'}{\sqrt{g^2 + g'^2}}B_\alpha \quad (2.27)$$

$$A_\alpha = \frac{g'}{\sqrt{g^2 + g'^2}}A_\alpha^3 + \frac{g}{\sqrt{g^2 + g'^2}}B_\alpha \quad (2.28)$$

where  $A_\alpha$  is the combination of  $A_\alpha^3$  and  $B_\alpha$  orthogonal to  $Z_\alpha$ . Equation 2.25 thus becomes

$$\mathcal{L} = \frac{1}{2}\partial_\alpha H \partial^\alpha H + \frac{g^2}{4}(v + H)^2 W_\alpha^+ W^{-\alpha} + \frac{g^2 + g'^2}{8}(v + H)^2 Z_\alpha Z^\alpha - \frac{\lambda}{4}(2vH + H^2)^2. \quad (2.29)$$

From this Lagrangian the mass terms can be extracted,

$$\mathcal{L}_{\text{mass}} = \frac{v^2}{4} \left[ g^2 W_\alpha^+ W^{-\alpha} + (g^2 + g'^2) Z_\alpha Z^\alpha - 4\lambda H^2 \right] \quad (2.30)$$

which gives

$$m_W^2 = \frac{v^2 g^2}{4} \quad (2.31)$$

$$m_Z^2 = \frac{v(g^2 + g'^2)}{4} \quad (2.32)$$

$$m_A^2 = 0 \quad (2.33)$$

$$m_H^2 = 2v^2 \lambda \quad (2.34)$$

for the masses of the physical bosons,  $W^\pm$ ,  $Z^0$ ,  $A^0$ , and  $H$ .  $A^0$  can be identified as the massless photon since it has no mass term in  $\mathcal{L}_{\text{mass}}$ .

### 2.2.3 Fermion Masses

The fermion masses are somewhat more complicated. A mass term inherently connects left-handed and right-handed fields but the simplest such terms, for example  $\mathcal{L}_{\text{mass}} = -m_l(\bar{E}_{lL}e_{lR} + \bar{e}_{lR}E_{lL})$ , are forbidden by gauge invariance since they combine doublet and singlet representations. Mass terms can be introduced via spontaneous symmetry breaking if couplings are included between the fermions and the Higgs field. This process is demonstrated below in the quark sector. Lepton masses are addressed later in Section 2.3.1.

The coupling to the Higgs boson, when the vacuum expectation value in Equation 2.19 is taken, introduces mass terms of the form [38]:

$$\mathcal{L}_{\text{mass}} = -\bar{D}_L M^{\text{down}} D_R \left( 1 + \frac{H}{v} \right) + \text{h.c.} \quad (2.35)$$

where

$$D_{L,R} = \begin{pmatrix} d_{L,R} \\ s_{L,R} \\ b_{L,R} \end{pmatrix} \quad (2.36)$$

and  $M^{\text{down}}$  is a unitary, *not* diagonal, matrix. In order to get the physical masses, the matrix  $M^{\text{down}}$  as well as the corresponding  $M^{\text{up}}$  need to be diagonalized:

$$M^{\text{down}} = V_L^{\text{down}} m^{\text{down}} V_R^{\text{down} \dagger}, \quad M^{\text{up}} = V_L^{\text{up}} m^{\text{up}} V_R^{\text{up} \dagger}. \quad (2.37)$$

where  $m^{\text{down}} = \text{diag}(m_d, m_s, m_b)$  and  $m^{\text{up}} = \text{diag}(m_u, m_c, m_t)$  and  $V_{L,R}^{\text{up,down}}$  are unitary transformation matrices. Using these unitary transformations, mass eigenvector states can be introduced

$$D'_{L,R} = \begin{pmatrix} d'_{L,R} \\ s'_{L,R} \\ b'_{L,R} \end{pmatrix} = V_{L,R}^{\text{down} \dagger} D_{L,R}, \quad U'_{L,R} = \begin{pmatrix} u'_{L,R} \\ c'_{L,R} \\ t'_{L,R} \end{pmatrix} = V_{L,R}^{\text{up} \dagger} U_{L,R} \quad (2.38)$$

as well as combined helicity states

$$U' = U'_L + U'_R, \quad D' = D'_L + D'_R \quad (2.39)$$

wich allows standard mass terms for the quarks to be written

$$\mathcal{L}_{\text{mass}} = - \sum_{u_1=u,c,t} m_{u_1} \bar{u}'_1 u'_1 - \sum_{d_1=d,s,b} m_{d_1} \bar{d}'_1 d'_1. \quad (2.40)$$

While the primed quark fields correspond to definite masses, they do not have definite transformation properties under the gauge group. Consequently, when the quark fields appear in the interaction terms (described in the next section), they are actually mixtures of the quarks with definite masses:

$$d_{1L} = \sum_{d_1=d,s,b} V_{ud_1} d'_{1L} \quad (2.41)$$

$$s_{1L} = \sum_{d_1=d,s,b} V_{cd_1} d'_{1L} \quad (2.42)$$

$$b_{1L} = \sum_{d_1=d,s,b} V_{td_1} d'_{1L}, \quad (2.43)$$

where

$$V = V_L^{\text{up} \dagger} V_L^{\text{down}} \quad (2.44)$$

is a  $3 \times 3$  unitary matrix called the Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix [40, 41].

The masses of the charged leptons are introduced via the same procedure. However, in the GWS electroweak theory, neutrinos were massless and appeared only as left-handed particles. Consequently, there are only two mixing matrices,  $U_{L,R}^{\text{lep}}$ , rather than the four required for the quarks and the mixing introduced by the combination of the up and down transformations does not appear. Of course, if neutrinos do have mass, then mixing among the lepton states analogous to the quarks would be expected.

## 2.2.4 Neutrino Interactions

The couplings of the gauge bosons to the fermion fields, and hence the dynamics of the weak interaction, are uniquely determined by the covariant derivative along with the quantum numbers of the fields. The most generic form for the covariant derivative of a fermion with SU(2) representation,  $\mathbf{T}$ , and a U(1) charge  $Y$  is

$$D_\alpha = \partial_\alpha - ig\mathbf{A}_\alpha \cdot \mathbf{T} - ig'\frac{1}{2}YB_\alpha, \quad (2.45)$$

where the  $\mathbf{T}$  matrices are weak isospin representations (Pauli matrices are a convenient basis) and  $Y$  is called the weak hypercharge (and gives the  $Y$  subscript in  $U(1)_Y$ ). The covariant derivative can be rewritten in terms of the massive boson fields just developed,

$$D_\alpha = \partial_\alpha - i\frac{g}{\sqrt{2}}(W_\alpha^+T^+ + W_\alpha^-T^-) - i\frac{1}{\sqrt{g^2 + g'^2}}Z_\alpha(g^2T^3 - g'^2\frac{1}{2}Y) - i\frac{gg'}{\sqrt{g^2 + g'^2}}A_\alpha(T^3 + \frac{1}{2}Y) \quad (2.46)$$

where  $T^\pm = (T^1 \pm iT^2)$  indicate  $\pm 1$  changes in weak isospin. If the last term is identified as the standard electromagnetic interaction, then the coefficient must be equal to the charge on the electron,  $e$ ,

$$e = \frac{gg'}{\sqrt{g^2 + g'^2}} \quad (2.47)$$

and the gauge generator  $T^3 + \frac{1}{2}Y$  must be the electric charge quantum number

$$Q = T^3 + \frac{1}{2}Y \quad (2.48)$$

where  $T^3$  is the third component of weak isospin and  $Y$  is called the weak hypercharge; this expression is known as the Gell-Mann-Nishijima relation [42, 43]. The constants  $g$  and  $g'$  can also be related to one another by the weak mixing angle,  $\theta_w$ , which defines the rotation from the  $(A^3, B)$  basis to the  $(Z, A)$  basis,

$$\begin{pmatrix} Z \\ A \end{pmatrix} = \begin{pmatrix} \cos \theta_w & -\sin \theta_w \\ \sin \theta_w & \cos \theta_w \end{pmatrix} \begin{pmatrix} A^3 \\ B \end{pmatrix}, \quad (2.49)$$

Particle	Helicity	$Q$	$T^3$	$Y$
$u, c, t$	L	+2/3	+1/2	+1/3
$d, s, b$	L	-1/3	-1/2	+1/3
$e, \mu, \tau$	L	-1	-1/2	-1
$\nu_e, \nu_\mu, \nu_\tau$	L	0	+1/2	-1
$u, c, t$	R	+2/3	0	+4/3
$d, s, b$	R	-1/3	0	-2/3
$e, \mu, \tau$	R	-1	0	-2
$\nu_e, \nu_\mu, \nu_\tau$	R	0	0	0

Table 2.2: Electroweak charges of standard-model fermions.

which, when compared with Equations 2.27 and 2.28, define the angle as

$$\cos \theta_w = \frac{g}{\sqrt{g^2 + g'^2}}, \quad \sin \theta_w = \frac{g'}{\sqrt{g^2 + g'^2}}. \quad (2.50)$$

With these relations, we can now rewrite the covariant derivative in terms of only the familiar electron charge,  $e$ , and the new weak mixing angle,  $\theta_w$ ,

$$D_\alpha = \partial_\alpha - i \frac{e}{\sqrt{2} \sin \theta_w} (W_\alpha^+ T^+ + W_\alpha^- T^-) - i \frac{e}{\sin \theta_w \cos \theta_w} Z_\alpha (T^3 - \sin^2 \theta_w Q) - ie A_\alpha Q. \quad (2.51)$$

The relative masses of the  $W$  and  $Z$  bosons are even determined by this mixing angle

$$\frac{m_W}{m_Z} = \cos \theta_w. \quad (2.52)$$

Next the quantum numbers of the fermion fields must be determined. This can be done using the relation in Equation 2.48, the already known quantum numbers shown in Table 2.1, and their weak isospins. The left-handed fermions, the  $SU(2)$  doublets, have a weak isospin  $T = \frac{1}{2}$  with  $T^3$  projections of opposite signs for the components of the doublets. The right-handed  $SU(2)$  singlets all have  $T = 0$  and do not participate in weak interactions. Finally, by requiring that each doublet element has the same value of  $Y$ , all of the quantum numbers can be constrained. They are summarized in Table 2.2.

Now all the pieces are in place to write down the kinetic terms in the Lagrangian of the weak interaction. For this discussion the masses of the fermions will be ignored – effectively a high-energy approximation. With this approximation, the kinetic Lagrangian takes the form

$$\mathcal{L}_{kin} = \sum_{l=e,\mu,\tau} \left[ \bar{E}_{lL}(i\mathcal{D})E_{lL} + \bar{e}_{lR}(i\mathcal{D})e_{lR} \right] + \sum_{u_1=u,c,t} \bar{Q}_{u_1L}(i\mathcal{D})Q_{u_1L} + \sum_{d_1=d,s,b} \bar{q}_{d_1R}(i\mathcal{D})q_{d_1R} \quad (2.53)$$

where the covariant derivative is given by Equation 2.51 with  $\mathbf{T}$  and  $Y$  evaluated for the relevant fermion representations. By substituting in the covariant derivatives, the physics of the theory will






Figure 2.1: Feynman diagrams at lowest order for charged current (left) and neutral current (right) neutrino-quark interactions in MINOS.

become more apparent.  $\mathcal{L}_{kin}$  becomes

$$\begin{aligned} \mathcal{L}_{kin} = & \sum_{l=e,\mu,\tau} \left[ \bar{E}_{lL}(i\not{\partial})E_{lL} + \bar{e}_{lR}(i\not{\partial})e_{lR} \right] + \sum_{u_1=u,c,t} \bar{Q}_{u_1L}(i\not{\partial})Q_{u_1L} + \sum_{d_1=d,s,b} \bar{q}_{d_1R}(i\not{\partial})q_{d_1R} \\ & + \frac{g}{\sqrt{2}} \left[ W^+ J_{CC}^\alpha + W^- J_{CC}^{\alpha\dagger} \right] + \frac{g}{\cos\theta_w} Z_\alpha J_{NC}^\alpha + e A_\alpha J_{EM}^\alpha \end{aligned} \quad (2.54)$$

separated into free kinetic terms and interactions between the vector bosons and the following currents

$$J_{CC}^\alpha = \bar{\nu}_{eL}\gamma^\alpha e_L + \bar{\nu}_{\mu L}\gamma^\alpha \mu_L + \bar{\nu}_{\tau L}\gamma^\alpha \tau_L + \bar{u}_L\gamma^\alpha d_L + \bar{c}_L\gamma^\alpha s_L + \bar{t}_L\gamma^\alpha b_L \quad (2.55)$$

$$\begin{aligned} J_{NC}^\alpha = & \sum_{l=e,\mu,\tau} \bar{\nu}_{lL}\gamma^\alpha \left( \frac{1}{2} \right) \nu_{lL} + \sum_{l=e,\mu,\tau} \bar{l}_L\gamma^\alpha \left( -\frac{1}{2} \right) l_L \\ & + \sum_{u_1=u,c,t} \bar{u}_{1L}\gamma^\alpha \left( \frac{1}{2} \right) u_{1L} + \sum_{d_1=d,s,b} \bar{d}_{1L}\gamma^\alpha \left( -\frac{1}{2} \right) d_{1L} \\ & - \sin^2\theta_w J_{EM}^\alpha \end{aligned} \quad (2.56)$$

$$J_{EM}^\alpha = \sum_{l=e,\mu,\tau} \bar{l}\gamma^\alpha (-1)l + \sum_{u_1=u,c,t} \bar{u}_1\gamma^\alpha \left( \frac{2}{3} \right) u_1 + \sum_{d_1=d,s,b} \bar{d}_1\gamma^\alpha \left( -\frac{1}{3} \right) d_1. \quad (2.57)$$

where  $CC$  stands for ‘‘charged current,’’ since the interaction is with the charged vector boson, and  $NC$  stands for ‘‘neutral current,’’ since the interaction is with the neutral weak vector boson. There are several properties to note in this Lagrangian. First, notice that thanks to the  $T^+$  and  $T^-$  operators, the charged current only connects particles that differ by 1 in weak isospin (charged lepton to neutrino, up-type quark to down-type quark). Second,  $J_{EM}^\alpha$ , which couples to the photon field  $A_\alpha$ , is the standard electromagnetic current and does not differentiate between left- and right-handed particles. Finally, note that right-handed neutrinos never appear in these currents since they lack both electromagnetic charge and weak hypercharge, implying that they do not interact via the electroweak force and hence are called ‘‘sterile.’’

Even left-handed neutrinos, since they lack an electromagnetic charge, do not appear in the electromagnetic current and hence only interact via the two weak currents. In the MINOS experiment,

the neutrinos being studied interact with a nucleus in the detector (usually iron). They can undergo a charged current interaction by exchanging a  $W^\pm$  boson,

$$\nu_l + q \rightarrow l^- + q' \quad (2.58)$$

$$\bar{\nu}_l + q \rightarrow l^+ + q' \quad (2.59)$$

where  $q$  and  $q'$  are up- and down-type quarks in the same family (except for the mixing described in the previous section), or they can undergo a neutral current interaction by exchanging a  $Z^0$  boson,

$$\nu_l + q \rightarrow \nu_l + q \quad (2.60)$$

$$\bar{\nu}_l + q \rightarrow \bar{\nu}_l + q \quad (2.61)$$

where the quark does not change. Feynman diagrams for these processes at lowest order are shown in Figure 2.1.

Most neutrinos in MINOS are at high enough energies that they interact via deep inelastic scattering, where the neutrino interacts with an individual parton in the nucleus. The matrix element for this interaction has the form,

$$|\mathcal{M}|^2 = \left(\frac{4G_f}{\sqrt{2}}\right)^2 L_{\alpha\beta} W^{\alpha\beta} \quad (2.62)$$

where  $L_{\alpha\beta}$  and  $W_{\alpha\beta}$  correspond to the leptonic and hadronic vertices in Figure 2.1. The leptonic vertex can be calculated using the weak charged current and, when summed over initial and final spins, is a function of the initial and final lepton momenta,  $k$  and  $k'$ :

$$L_{\alpha\beta} = \frac{2}{m_\mu m_\nu} [k'_\alpha k_\beta + k'_\beta k_\alpha - k \cdot k' g_{\alpha\beta} \mp ik^\gamma k'^\delta \epsilon_{\gamma\delta\alpha\beta}]. \quad (2.63)$$

Note that the last term encapsulates the parity-violation of the weak force, changing sign between left-handed neutrinos and right-handed antineutrinos. The hadronic vertex, on the other hand, cannot be calculated directly but must instead be parameterized in terms of structure functions that describe the underlying structure of the nucleon. In the end, the differential neutrino cross section, as a function of the *Bjorken scaling variables*  $x$  and  $y$ ,<sup>6</sup> takes the form

$$\frac{d^2\sigma^{\nu(\bar{\nu})N}}{dx dy} = \frac{G_f^2 m_N E_\nu}{\pi(1 + Q^2/m_W^2)^2} \left[ xy^2 F_1 + \left(1 - y - \frac{m_N xy}{2E_\nu}\right) F_2 \pm y(1 - y/2) x F_3 \right], \quad (2.64)$$

where  $F_i = F_i(x, Q^2)$  are the structure functions that relate to the internal structure of the nucleon and are functions of the kinematical variables. Note that the parity-violating nature of the weak

---

<sup>6</sup> $x$  is the fraction of the total nucleon momentum carried by the quark involved in the interaction and  $y$  is the fraction of the momentum of the incident lepton that is transferred to the hadronic system.




Figure 2.2: The  $\nu_\mu$  and  $\bar{\nu}_\mu$  interaction cross sections as a function of energy as measured in many experiments, including MINOS. Figure taken from [44].

force is encapsulated in the  $\pm F_3$  term.

When the neutrino and antineutrino cross sections are measured, the  $\bar{\nu}N$  cross section turns out to be significantly lower than the  $\nu N$  cross section, as seen in Figure 2.2. This difference can be understood qualitatively using the naive parton model [45] of the nucleon where only the on-shell valence quarks are considered (argument from [46]). In this model the inelastic neutrino-nucleon scattering becomes elastic neutrino-parton scattering. The  $V - A$  nature of the weak interaction requires that the left-handed neutrinos interact only with left-handed particles and the right-handed antineutrinos interact only with right-handed particles. By also requiring that angular momentum is conserved, we find that neutrino-parton ( $\nu q$ ) interactions must always have total spin 0 and antineutrino-parton ( $\bar{\nu} q$ ) interactions must always have total spin 1 (the opposite would be true for  $\bar{q}$  partons, but in the naive parton model the sea which might contain antiquarks is neglected). Thus, the  $\frac{d^2\sigma^{\nu q}}{d\cos\theta^*}$  cross section is isotropic and depends only on the center-of-mass energy while the  $\frac{d^2\sigma^{\bar{\nu} q}}{d\cos\theta^*}$  cross section is reduced by a factor of  $\left(\frac{1+\cos\theta^*}{2}\right)^2$  ( $\theta^*$  is the angle of the outgoing lepton in the center-of-mass frame). Thus the antineutrino interaction cross section is reduced relative to the neutrino cross section.

Whatever the flavor of the outgoing quark, it cannot be detected by MINOS since color confinement will cause the quark to immediately hadronize. The energy transferred from the neutrino to the outgoing quark will cause that quark to begin to separate from the nucleon. However, as the quark separates the force between it and the nucleon does not diminish, which means as the separation increases it eventually becomes more energetically favorable to spontaneously produce more quarks and gluons out of the vacuum than to allow the original quark to get any further away. This

behavior prevents color-charged particles from existing individually and is called color confinement. As the new quarks and gluons that were just produced continue to separate from one another, the process can repeat itself until the quarks are low enough in energy to form bound meson and baryon states. It is these bound particles that are observed in the detector. The spontaneous appearance of quarks and gluons is called hadronization, and the collection of outgoing particles is called the hadronic shower.<sup>7</sup>

## 2.3 Neutrino Oscillations

By the late 1960's, physicists believed they had the neutrino well in hand. The neutrino was a massless fermion, it appeared only in a left-handed helicity state, and it interacted solely via the weak nuclear force that Glashow, Weinberg, and Salam had unified with the electromagnetic force. In 1970, Ray Davis put the neutrino back in turmoil. Back in the mid 1950's, Davis searched for neutrinos at a nuclear reactor using the reaction shown in Equation 2.3, which would be forbidden by the GWS theory since it violates lepton number. However, this radiochemical process,



would proceed if Davis could expose his chlorine-based detector to a significant source of  $\nu_e$ 's. The fusion reactions in the Sun turned out to be just such a source, and Davis was able to isolate the solar neutrinos by putting his detector 4,850 ft deep in the Homestake mine in South Dakota. The detector itself was a large tank filled with perchloroethylene, a chlorine-rich dry cleaning chemical. As the neutrinos interacted in the tank, they would produce argon which Davis could isolate and count.

Working with Davis was theoretical astrophysicist John Bahcall who calculated the expected flux of neutrinos from the Sun [47]. However, when the data from the detector was examined it was approximately a third lower than predicted [48]. This discrepancy became known as the 'solar neutrino problem,' and a number of neutrino experiments were built in the following years to solve it. Where the Homestake experiment used chlorine as the target, the newer experiments used a range of target nuclei. The different nuclei had different thresholds for interacting with a neutrino, and hence measured different parts of the solar neutrino energy spectrum. For example, SAGE in the Soviet Union and GALLEX in Italy used gallium, which has a very low threshold in the hundreds of keV. The chlorine reaction used in the Homestake experiment has a threshold a little below an MeV and the Kamiokande<sup>8</sup> experiment in Japan used water, which has a threshold of about 8 MeV.

---

<sup>7</sup>It is called a hadronic shower since it is hadronic in origin, even though it may have significant electromagnetic components, for example if a neutral pion is produced and it decays via  $\pi^0 \rightarrow 2\gamma$ .

<sup>8</sup>This experiment was actually built primarily to search for proton decay but could also measure solar neutrinos.

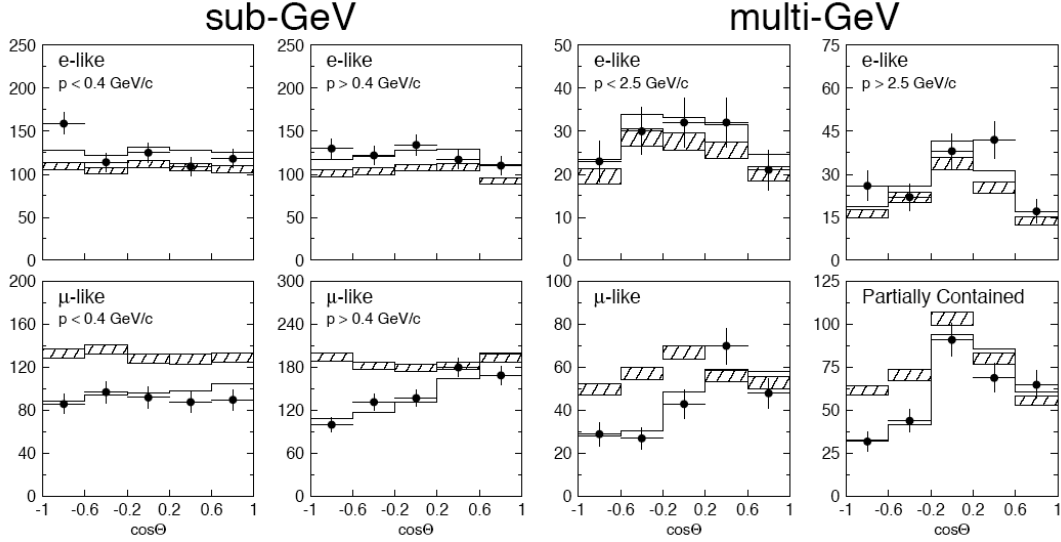


Figure 2.3: Zenith angle distributions from  $e$ -like events (upper row) and  $\mu$ -like events (lower row). The hatched region shows the expectation without oscillations and the line shows the best  $\nu_\mu \rightarrow \nu_\tau$  oscillation fit. Figure taken from [50]. The  $\mu$ -like samples (lower row) show deficits which depend on zenith angle while the  $e$ -like samples (upper row) show no deficit.

The measurements at different energies helped to vindicate the Bahcall solar model, even before a solution had been agreed upon. According to the model, the neutrino fluxes depend strongly on the interior temperature of the Sun, and the different neutrino sources have different dependencies, ranging from  $\Phi_\nu \propto T^{-1}$  for  $pp$  neutrinos to  $\Phi_\nu \propto T^{24}$  for  ${}^8\text{B}$  neutrinos [49]. In order to account for the deficit, the Sun would need to be significantly cooler than astronomers believed. However, the energy distribution measured by the various solar neutrino experiments required, if anything, a higher core temperature in the Sun. The solar model could not account for both of these observations simultaneously.

After nearly three decades of struggle, the solar neutrino problem was finally solved by two experiments: Super-Kamiokande [50] in Japan and the Sudbury Neutrino Observatory (SNO) [51] in Canada. The Super-Kamiokande detector [52] is a large tank of ultra pure water instrumented with phototubes to measure the Cherenkov light emitted by the outgoing products of a neutrino-nucleus interaction. Super-Kamiokande's revolutionary measurement was made not with solar neutrinos, but with atmospheric neutrinos produced in the air showers of cosmic rays, primarily  $\nu_\mu$ 's and  $\nu_e$ 's. Super-Kamiokande saw no deficit in the atmospheric  $\nu_e$ 's, but there was a significant deficit in the  $\nu_\mu$ 's. Furthermore, the deficit was not uniform: there was a strong dependence on zenith angle (see Figure 2.3). Neutrinos coming from directly above the detector, i.e. those that had traveled only a short distance, showed no deficit; however, the neutrinos that came from below and had traveled a long distance showed a significant deficit. This zenith-angle, or path length, dependence





Figure 2.4: The flux of  $\nu_\mu$  and  $\nu_\tau$  vs. the flux of  $\nu_e$  from  $^8\text{B}$  solar neutrinos as measured at SNO. The three colored bands correspond to different interactions which are sensitive to different mixtures of the flavors. Notice that the NC band, which is sensitive to the total rate of neutrinos of all flavors, is in excellent agreement with the standard solar model. All three samples are consistent with the interpretation that a fraction of the solar  $\nu_e$  flux has changed flavor. Figure taken from [53].

is a signature of neutrino flavor oscillations, similar to those proposed by Bruno Pontecorvo back in 1968. Since no excess was seen in the  $\nu_e$ 's, the  $\nu_\mu$ 's were presumed to be oscillating to  $\nu_\tau$ 's. The observation of neutrino oscillations had a significant consequence: as shown for the quarks in Section 2.2.3, in order for neutrinos to oscillate, they need to have mass (this will be demonstrated specifically for neutrinos in the next section).

While Super-Kamiokande was the first experiment to see neutrino oscillations, it was the SNO experiment that truly ended the solar neutrino problem in a model-independent way. The experiment consists of 1 kTon of heavy water, again instrumented with photomultiplier tubes and situated deep beneath the surface (2,092 m) [54]. Unlike Super-Kamiokande, SNO measured solar neutrinos, like many of the previous experiments. What made SNO unique was that it could measure three different neutrino interactions, each of which was sensitive to a different mix of flavors:

1. The charged-current process, sensitive only to electron neutrinos,



2. The neutral-current process, sensitive to all flavors equally,



3. The electron-scattering process, sensitive to all flavors, but approximately six times more sensitive to  $\nu_e$ 's,<sup>9</sup>

$$\nu_x + e^- \rightarrow \nu_x + e^-. \quad (2.68)$$

By measuring both the flux of electron neutrinos and of all neutrinos, SNO conclusively demonstrated that the solar neutrino deficit was, in fact, only a deficit of solar *electron* neutrinos – the flux of all neutrino species was in good agreement with the standard solar model. The flux constraints from the three measurements, and the consistency among them and the standard solar model, are shown in Figure 2.4. Today, with the vindicated standard solar model and precise measurements of the solar neutrino spectrum, Bahcall has been able to calculate the temperature of the Sun's core with better than 1% precision [55].

### 2.3.1 Neutrino Masses

All of the physics of neutrino mixing (oscillations) is defined by the neutrino mass matrix, just as down-quark mixing is defined by the quark mass matrix in Section 2.2.3. Neutrino mass is not part of the GWS electroweak theory and so neutrino mass terms need to be added to the Lagrangian. There are numerous theories which can introduce the neutrino mass terms, such as the various seesaw mechanisms, but they all must eventually produce terms with the same form. Two classes of mass terms that can be written because neutrinos are neutral particles: Dirac and Majorana. The two classes differ in the relationship between the neutrino and antineutrino.

#### 2.3.1.1 Dirac Mass

In a Dirac mass term, neutrinos are treated as a 4-component spinor like all other fermions, with left- and right-handed particles and antiparticles which are distinct from one another:

$$\nu_L, \nu_R, \bar{\nu}_L, \bar{\nu}_R. \quad (2.69)$$

The mass term in the Lagrangian has the form:

$$\mathcal{L}_{\text{mass}}^D = - \sum_{l', l} \bar{\nu}_{l'L} M_{l'l}^D \nu_{lR} + \text{h.c.} \quad (2.70)$$

where  $M^D$  is a  $3 \times 3$  complex matrix and  $l', l$  run over  $\{e, \mu, \tau\}$ , the neutrinos that couple to the weak force (the 'flavor basis'). This mass term preserves the invariance of the Lagrangian under

---

<sup>9</sup>Electron scattering is sensitive to all flavors of neutrino via  $Z^0$  exchange, but  $\nu_e$ 's can also interact via  $W^-$  exchange.

global lepton number transformations

$$\nu_{lL} \rightarrow e^{i\Lambda} \nu_{lL}, \quad \nu_{lR} \rightarrow e^{i\Lambda} \nu_{lR}, \quad l \rightarrow e^{i\Lambda} l, \quad (2.71)$$

$$\bar{\nu}_{lL} \rightarrow e^{-i\Lambda} \bar{\nu}_{lL}, \quad \bar{\nu}_{lR} \rightarrow e^{-i\Lambda} \bar{\nu}_{lR}, \quad \bar{l} \rightarrow e^{-i\Lambda} \bar{l}, \quad (2.72)$$

and consequently conserves total lepton number. In order to determine the physical (i.e. real, single-valued) masses, the matrix  $M^D$  must be diagonalized:

$$M^D = U^\dagger m V \quad (2.73)$$

where  $m$  is a real, positive diagonal matrix,  $m_{ij} = m_i \delta_{ij}$ , with  $m_i > 0$ . The neutrino fields can be rewritten in this ‘mass basis,’

$$\nu_{lL} = \sum_{i=1}^3 U_{li} \nu_{iL} \quad (2.74)$$

$$\nu_{lR} = \sum_{i=1}^3 V_{li} \nu_{iR} \quad (2.75)$$

where, again,  $l \in \{e, \mu, \tau\}$  and three neutrino mass states have been presumed. When the mass Lagrangian  $\mathcal{L}_{\text{mass}}^D$  is rewritten in the mass basis, it takes on the form of a standard mass term,

$$\mathcal{L}_{\text{mass}}^D = - \sum_{i=1}^3 m_i \bar{\nu}_i \nu_i. \quad (2.76)$$

The unitary matrix,  $U$ , is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix [15, 56].

### 2.3.1.2 Majorana Mass

Unlike a Dirac fermion, a Majorana fermion only has two components: the charge-conjugate of the left-handed particle is the right-handed particle. Put another way, the particle is its own antiparticle.

$$\nu_L, \quad \bar{\nu}_R \equiv (\nu_L)^c = C \bar{\nu}_L^T \quad (2.77)$$

where  $C$  is the charge-conjugation operator. Neutrinos are the only Standard Model fermions that can be Majorana particles since all other fermions have electric charges which distinguish the particles from the antiparticles.

By definition, a Lagrangian mass term is a Lorentz-invariant product of the left- and right-handed components of a field. Thus, it must be demonstrated that  $(\nu_L)^c$  is, in fact, a right-handed field.



The handedness of a field is defined by its behavior under multiplication by  $\gamma^5$ :

$$\gamma^5 \psi_L = -\psi_L, \quad \gamma^5 \psi_R = \psi_R. \quad (2.78)$$

Let us, therefore, examine the behavior of the field in question,  $(\nu_L)^c$ ,

$$\gamma^5 (\nu_L)^c = \gamma^5 C \bar{\nu}_L^T \quad (2.79)$$

$$= (\bar{\nu}_L C^T \gamma_5^T)^T \quad (2.80)$$

Using  $C^T = -C$ ,  $C \gamma_5^T C^{-1} = \gamma_5$ , the expression becomes

$$\gamma_5 (\nu_L)^c = -(\bar{\nu}_L C \gamma_5^T C^{-1} C)^T \quad (2.81)$$

$$= -(\bar{\nu}_L \gamma_5 C)^T \quad (2.82)$$

Finally, using the relation  $\bar{\nu}_L \gamma_5 = \nu_L$ ,

$$\gamma_5 (\nu_L)^c = -(\nu_L C)^T \quad (2.83)$$

$$= C \nu_L^T \quad (2.84)$$

$$= (\nu_L)^c \quad (2.85)$$

showing that  $(\nu_L)^c$  behaves like the right-handed component of a field and can be used to construct a mass term. That term has the form:

$$\mathcal{L}_{\text{mass}}^M = -\frac{1}{2} \sum_{l', l} \bar{\nu}_{l'} M_{l'l}^M (\nu_l)^c + \text{h.c.} \quad (2.86)$$

or, in matrix form,

$$\mathcal{L}_{\text{mass}}^M = -\frac{1}{2} \bar{\nu}_L M^M (\nu_L)^c + \text{h.c.}, \quad (2.87)$$

where  $M^M$  is a complex matrix and  $\nu_L^T = (\nu_{eL} \ \nu_{\mu L} \ \nu_{\tau L})$ . Unlike the Dirac mass term, this term is not invariant under global gauge transformations. While the left- and right-handed Dirac fields transformed in the same way (Equation 2.71), the relationship between the left- and right-handed Majorana neutrinos requires that

$$\nu_{lL} \rightarrow e^{i\Lambda} \nu_{lL}, \quad \nu_{lR} = (\nu_{lL})^c \rightarrow e^{-i\Lambda} (\nu_{lL})^c, \quad (2.88)$$

and Equation 2.86 is not invariant under this transformation. Consequently, a theory with Majorana neutrinos does not conserve total lepton number.

The relationship between the left- and right-handed fields in this mass term provide a constraint on the form of  $M^M$ :

$$\bar{\nu}_L M^M (\nu_L)^c = \bar{\nu}_L M^M C \bar{\nu}_L^T \quad (2.89)$$

$$= -\bar{\nu}_L (M^M)^T C^T \nu_L^T \quad (2.90)$$

$$= \bar{\nu}_L (M^M)^T C \nu_L^T \quad (2.91)$$

using the anticommutation properties of fermion fields. This implies that

$$M^M = (M^M)^T \quad (2.92)$$

or that  $M^M$  is symmetric. Thus, when the matrix is diagonalized to find the physical neutrino masses, only one unitary matrix,  $U$ , is needed:

$$M^M = U m U^T \quad (2.93)$$

where, again,  $U$  is a unitary matrix and  $m_{ij} = m_i \delta_{ij}$ , with  $m_i > 0$ . Substituting into the Lagrangian,

$$\mathcal{L}_{\text{mass}}^M = -\frac{1}{2} \bar{\nu}_L U m U^T C \nu_L^T + \text{h.c.} \quad (2.94)$$

$$= -\frac{1}{2} \overline{U^\dagger \nu_L} m (U^\dagger \nu_L)^c - \frac{1}{2} \overline{(U^\dagger \nu_L)^c} m U^\dagger \nu_L \quad (2.95)$$

$$= -\frac{1}{2} \bar{\nu}^M m \nu^M \quad (2.96)$$

where

$$\nu^M = U^\dagger \nu_L + (U^\dagger \nu_L)^c = \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (2.97)$$

once again producing the canonical mass term by transforming the fields into the mass basis (compare to Equation 2.76). It is clear from Equation 2.97 that

$$(\nu^M)^c = \nu^M, \quad (2.98)$$

and thus that

$$\nu_i^c = \nu_i. \quad (2.99)$$

This relation is known as the Majorana condition, and it states that in the mass basis the neutrino and antineutrino are the same particle.

As in the Dirac case (Equation 2.74), the mass and flavor bases are related to one another by




Figure 2.5: Feynman diagrams for  $2\nu$  (left) and  $0\nu$  (right) double beta decay. Note, the  $0\nu\beta\beta$  is only possible if the same particle can act as  $\nu_e$  and  $\bar{\nu}_e$ .

the mixing matrix  $U$

$$\nu_L = U\nu_L^M \quad \text{or} \quad \nu_{iL} = \sum_{j=1}^3 U_{ij}\nu_{jL}. \quad (2.100)$$

It is significant to note that whether neutrinos are Dirac or Majorana, in the end there are still three neutrinos of definite mass  $m_1, m_2, m_3$  which are related to the three neutrinos with definite  $SU(2)_L \otimes U(1)_Y$  transformation properties (the flavor neutrinos,  $\nu_e, \nu_\mu, \nu_\tau$ ) by the PMNS mixing matrix  $U$ . In fact, in either case only left-handed neutrinos and right-handed antineutrinos are ever observed since only those states couple to the electroweak force.

The only known way, experimentally, to distinguish between Dirac and Majorana neutrinos is to search for the rare neutrinoless double beta decay process ( $0\nu\beta\beta$ ). In a typical double beta decay, two neutrons transition to two protons by emitting two antineutrinos and two electrons (via two  $W^-$ 's),

$$2n \rightarrow 2p^+ + 2e^- + 2\bar{\nu}_e. \quad (2.101)$$

However, if the neutrino and antineutrino are the same particle, then instead of releasing two  $\bar{\nu}_e$ 's, one virtual  $\bar{\nu}_e$  could be emitted and then absorbed as a  $\nu_e$  all within the decay (since the emission of a  $\bar{\nu}_e$  and absorption of a  $\nu_e$  are equivalent processes), leaving

$$2n \rightarrow 2p^+ + 2e^-. \quad (2.102)$$

Note that in this process, total lepton number is violated by 2, which is forbidden if neutrinos are Dirac particles but is allowed if neutrinos are Majorana particles. Consequently, the observation of neutrinoless double beta decay is a sensitive test of lepton number conservation and hence the nature of the relationship between the neutrino and antineutrino.

### 2.3.2 Oscillations in Vacuum

Whatever the source of neutrino mass, and whatever type of particle the neutrino is, neutrino mixing takes on a common form. The neutrino flavor eigenstates are related to the neutrino mass eigenstates

via the unitary PMNS matrix (Equations 2.74 or 2.100). A weak interaction in nature will produce one of the flavor states, here denoted  $|\nu_\alpha\rangle$ , since these are the states which couple to the  $W^\pm$  and  $Z^0$  bosons. Let us examine the evolution of this state in time.

The Schrödinger equation tells us that the time evolution of the neutrino state will be given by

$$|\nu_\alpha(t)\rangle = e^{-iHt} |\nu_\alpha(0)\rangle \quad (2.103)$$

where  $H$  is the Hamiltonian and  $|\nu_\alpha(0)\rangle = |\nu_\alpha\rangle$  is the state of the neutrino at time  $t = 0$ . Since  $|\nu_\alpha\rangle$  is not an eigenstate of the Hamiltonian, this expression is difficult to evaluate. However, the mass basis states,  $|\nu_i\rangle$ , are eigenstates,

$$H |\nu_i\rangle = E_i |\nu_i\rangle \quad (2.104)$$

with eigenvalues  $E_i = \sqrt{p^2 + m_i^2}$ . Taking advantage of this relation and the mixing relationship (Equation 2.74) developed above, the time evolution of the neutrino state (Equation 2.103) becomes

$$|\nu_\alpha(t)\rangle = \sum_{i=1}^N e^{-iE_i t} U_{\alpha i}^* |\nu_i\rangle \quad (2.105)$$

where  $N$  is the number of neutrino mass states. However, when the neutrino state  $|\nu_\alpha(t)\rangle$  interacts again (the only way to observe the new state), it will be as a flavor eigenstate, not a mass eigenstate  $|\nu_i\rangle$ . So, thanks to  $U$ 's unitarity, by inverting Equation 2.74 we find

$$|\nu_\alpha(t)\rangle = \sum_{\beta} \sum_{i=1}^N e^{-iE_i t} U_{\beta i} U_{\alpha i}^* |\nu_\beta\rangle \quad (2.106)$$

where  $\beta$  sums over the flavor states.

The probability of observing a neutrino flavor state,  $|\nu_\beta\rangle$ , starting from state  $|\nu_\alpha\rangle$  after time  $t$  has elapsed is given by the square of the amplitude:

$$\begin{aligned} P_t(\nu_\alpha \rightarrow \nu_\beta) &= |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 \\ &= \left| \sum_{i=1}^N e^{-iE_i t} U_{\beta i} U_{\alpha i}^* \right|^2 \\ &= \sum_{i=1}^N \sum_{j=1}^N U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^* e^{-i(E_i - E_j)t}. \end{aligned} \quad (2.107)$$

Since neutrinos have very little mass, they are highly relativistic and their energy can be approximated as  $E_i \approx p + \frac{m_i^2}{2p}$ , and the time traveled can be converted to a distance,  $L$  (using natural units where  $c = 1$ ). Notice that  $p$  is the same for all neutrino mass states  $i$  since they are produced

coherently. The probability thus becomes

$$P_L(\nu_\alpha \rightarrow \nu_\beta) = \sum_{i=1}^N \sum_{j=1}^N U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^* e^{-i \frac{L}{2E} \Delta m_{ij}^2} \quad (2.108)$$

where  $\Delta m_{ij}^2 = m_i^2 - m_j^2$  and, given the small size of the absolute neutrino mass, the momentum  $p$  has been further approximated as  $E$ , the energy of the original neutrino state  $|\nu_\alpha\rangle$ . The expression can be written more conveniently as

$$\begin{aligned} P_L(\nu_\alpha \rightarrow \nu_\beta) = & \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\beta i} U_{\beta j}^* U_{\alpha i}^* U_{\alpha j}) \sin^2 \left( \frac{\Delta m_{ij}^2 L}{4E} \right) \\ & + 2 \sum_{i>j} \Im(U_{\beta i} U_{\beta j}^* U_{\alpha i}^* U_{\alpha j}) \sin \left( \frac{\Delta m_{ij}^2 L}{2E} \right). \end{aligned} \quad (2.109)$$

Some properties of neutrino oscillations are apparent directly from Equation 2.109. Oscillations are only possible if neutrinos are massive, and only if the mass eigenstates  $\nu_i$  have different masses. Otherwise, the sine terms go to zero and  $P_L(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta}$ . In fact, it is the sine terms that give the neutrino flavor transitions the name ‘‘oscillations.’’

It is interesting to examine some of the symmetry properties of this expression. Note that swapping  $\alpha$  and  $\beta$  (time-reversal) is equivalent to swapping  $U$  and  $U^*$ ,

$$P_L(\nu_\alpha \rightarrow \nu_\beta; U) = P_L(\nu_\beta \rightarrow \nu_\alpha; U^*), \quad (2.110)$$

while  $\mathcal{CPT}$ -invariance requires

$$P_L(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta) = P_L(\nu_\beta \rightarrow \nu_\alpha). \quad (2.111)$$

Putting these together we find

$$P_L(\nu_\alpha \rightarrow \nu_\beta; U) = P_L(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta; U^*). \quad (2.112)$$

Thus, if  $U = U^*$  ( $U$  is real) the transition probability is  $\mathcal{CP}$ -invariant and  $P_L(\nu_\alpha \rightarrow \nu_\beta) = P_L(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta)$ . Conversely, if it is observed that  $P_L(\nu_\alpha \rightarrow \nu_\beta) \neq P_L(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta)$ , then that is evidence that  $\mathcal{CP}$  is violated.

### 2.3.3 The Mixing Matrix

The previous discussion of neutrino oscillations is generic – it does not depend on a specific number of neutrinos. The Standard Model predicts three neutrinos corresponding to the three families of




Figure 2.6: Measurement of the hadron production cross section around the  $Z^0$  mass resonance at LEP. The curves show the prediction for this cross section with 2, 3, and 4 light active neutrino species. The results constrain  $N_\nu$  to  $2.984 \pm 0.008$ . Figure taken from [57].

quarks and charged leptons. The number of active neutrinos (specifically the number of neutrinos that couple to the  $Z^0$  boson) was measured with high precision by the Large Electron-Positron (LEP) collider at CERN in the 1990's. Several LEP experiments (ALEPH, DELPHI, L3, OPAL) made measurements of the hadron production cross section around the  $Z^0$  mass resonance and compared it with predictions assuming differing numbers of light neutrinos (see Figure 2.6). The combined results constrain  $N_\nu$  to  $2.984 \pm 0.008$  [35, 57].

The PMNS mixing matrix can be written in a standard form with three neutrinos. In general, an  $N \times N$  unitary matrix can be parameterized by

$$n_\theta = \frac{1}{2}N(N-1) = 3 \text{ (for } N=3\text{)} \quad (2.113)$$

angles and

$$n_\phi = \frac{1}{2}N(N+1) = 6 \text{ (for } N=3\text{)} \quad (2.114)$$

complex phases. In practice, not all of these phases will remain. The fermion fields can take on arbitrary phases and these phases can be chosen to eliminate phases in the mixing matrix.

Take, for example, the leptonic part of the charged current from Equation 2.55, which can be rewritten in terms of the neutrino mass states:

$$J_{CC}^{\alpha\dagger} = \sum_{l=e,\mu,\tau} \bar{l}_L \gamma^\alpha \nu_{lL} = \sum_{i=1}^N \bar{l}_L \gamma^\alpha U_{li} \nu_{iL}. \quad (2.115)$$

For Dirac neutrinos, each neutrino and charged lepton field can take on an arbitrary phase and still be physically equivalent:

$$\nu_{iL} \rightarrow \nu'_{iL} = e^{i\alpha_i} \nu_{iL}, \quad l_L \rightarrow l'_L = e^{i\beta_i} l_L. \quad (2.116)$$

The charged current in Equation 2.115 will remain unchanged if the mixing matrix,  $U$ , is also transformed using the phase matrix  $S$

$$U \rightarrow U' = S(\boldsymbol{\beta})US^\dagger(\boldsymbol{\alpha}) \quad (2.117)$$

where

$$S^\dagger(\boldsymbol{\alpha}) = e^{-i\alpha_N} \begin{pmatrix} e^{-i(\alpha_1 - \alpha_N)} \\ \vdots \\ 1 \end{pmatrix}, \quad S(\boldsymbol{\beta}) = e^{i\alpha_N} \begin{pmatrix} e^{i(\beta_1 - \alpha_N)} \\ \vdots \\ e^{i(\beta_N - \alpha_N)} \end{pmatrix}. \quad (2.118)$$

Note, the overall phase  $\alpha_N$  is pulled out of both  $S$  matrices to ensure the current does not pick up an overall phase. This leaves  $N + (N - 1)$  free phases which can be chosen to eliminate the phases in the PMNS matrix,  $U$ . The final number of phases, assuming three Dirac neutrinos, becomes,

$$n_\phi^D = \frac{N(N+1)}{2} - (2N-1) = \frac{(N-1)(N-2)}{2} = 1 \text{ (for } N=3\text{)}. \quad (2.119)$$

So the mixing of Dirac neutrinos is parameterized by 3 angles and 1 complex phase.

The Majorana fields can also taken on arbitrary phases, like the Dirac fields. However, the Majorana condition (Equation 2.99) already constrains the values of the  $\alpha_i$  neutrino phases, leaving only the charged lepton phases free. Thus, the Majorana mixing matrix,  $U^M$ , will include two phases,  $\alpha_1$  and  $\alpha_2$ , in addition to the mixing matrix  $U$  from the Dirac case

$$U^M = U \begin{pmatrix} \alpha_1 & 0 & 0 \\ 0 & \alpha_2 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.120)$$

and only  $S(\boldsymbol{\beta})$  enters the transformation of  $U^M$ ,

$$U^M \rightarrow U^{M'} = S(\boldsymbol{\beta})U^M \quad (2.121)$$

and only  $N$  free phases can be chosen. Thus, for three Majorana neutrinos,

$$n_\phi^M = \frac{N(N+1)}{2} - N = \frac{N(N-1)}{2} = 3, \quad (2.122)$$

and the mixing of Majorana neutrinos is parameterized by 3 angles and 3 complex phases, one that

was also in the Dirac matrix as well as the two constrained by the Majorana condition.

It is standard to identify each angle as an Euler rotation among two of the neutrino mass states. Therefore,  $\theta_{12}$  would be a rotation among  $|\nu_1\rangle$  and  $|\nu_2\rangle$  that leaves the  $|\nu_3\rangle$  alone:

$$|\nu_1\rangle' = \cos \theta_{12} |\nu_1\rangle + \sin \theta_{12} |\nu_2\rangle \quad (2.123)$$

$$|\nu_2\rangle' = -\sin \theta_{12} |\nu_1\rangle + \cos \theta_{12} |\nu_2\rangle \quad (2.124)$$

$$|\nu_3\rangle' = |\nu_3\rangle \quad (2.125)$$

or, in Matrix form,

$$|\nu\rangle' = U^{(12)} |\nu\rangle = \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} |\nu\rangle \quad (2.126)$$

where  $s_{12} = \sin \theta_{12}$  and  $c_{12} = \cos \theta_{12}$ , and equivalently for  $\theta_{13}$  and  $\theta_{23}$ . The complex phase,  $\delta$ , that occurs for both Dirac and Majorana neutrinos is included with the off-diagonal terms in the  $\theta_{13}$  rotation:

$$U^{(13)} = \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix}. \quad (2.127)$$

All told, the Dirac PMNS matrix takes the forms

$$|\nu\rangle' = U |\nu\rangle \quad (2.128)$$

$$= U^{(12)}U^{(13)}U^{(23)} |\nu\rangle \quad (2.129)$$

$$= \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} |\nu\rangle \quad (2.130)$$

$$= \begin{pmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{23}c_{12}s_{13}e^{i\delta} & c_{23}c_{12} - s_{23}s_{12}s_{13}e^{i\delta} & c_{13}s_{23} \\ s_{23}s_{12} - c_{23}c_{12}s_{13}e^{i\delta} & -s_{23}c_{12} - c_{23}s_{12}s_{13}e^{i\delta} & c_{13}s_{23} \end{pmatrix} |\nu\rangle. \quad (2.131)$$

The mixing matrix for Majorana neutrinos takes on the same general form, but with the two



additional phases included in Equation 2.120:

$$|\nu\rangle' = U^M |\nu\rangle \quad (2.132)$$

$$= U^{(12)}U^{(13)}U^{(23)} \begin{pmatrix} \alpha_1 & 0 & 0 \\ 0 & \alpha_2 & 0 \\ 0 & 0 & 1 \end{pmatrix} |\nu\rangle. \quad (2.133)$$

### 2.3.4 Two Neutrino Mixing

The oscillation probability among neutrino species, using the probability from Equation 2.109 and the mixing matrix from Equation 2.131, will generally be quite complicated when all three flavors are considered. If, however, the mixing can be approximated as between only two species, the probability can be simplified considerably. Let us consider transitions between  $\nu_\alpha$  and  $\nu_\beta$  which are mixtures of neutrino mass state  $\nu_1$  and  $\nu_2$ . With only two neutrinos, the mixing can be parameterized by one angle and no complex phases (neglecting diagonal Majorana phases which do not appear in the oscillation probabilities, see Equations 2.113, 2.119, 2.122). Thus, in Equation 2.109, there will be only one term in the sum  $i > j$  ( $i = 2$  and  $j = 1$ ) and since the matrix  $U$  is real,  $U = U^*$  and the imaginary term vanishes, leaving,

$$P_L(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} - 4U_{\alpha 2}U_{\alpha 1}U_{\beta 2}U_{\beta 1} \sin^2 \left( \frac{\Delta m_{21}^2 L}{4E} \right). \quad (2.134)$$

Substituting in the two-flavor mixing matrix,

$$U = \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} \\ -\sin \theta_{12} & \cos \theta_{12} \end{pmatrix}, \quad (2.135)$$

the probability becomes,

$$P_L(\nu_\alpha \rightarrow \nu_\alpha) = 1 - \sin^2 2\theta_{12} \sin^2 \left( \frac{\Delta m_{21}^2 L}{4E} \right) \quad (2.136)$$

$$P_L(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta_{12} \sin^2 \left( \frac{\Delta m_{21}^2 L}{4E} \right), \quad (2.137)$$

using the trigonometric identity  $2 \sin \theta \cos \theta = \sin 2\theta$ . Note that  $P_L(\nu_\alpha \rightarrow \nu_\alpha) + P_L(\nu_\alpha \rightarrow \nu_\beta) = 1$ , consistent with the conservation of probability.

Since  $U = U^*$  in the two-neutrino scenario, the time-reversal (Equation 2.110) and  $\mathcal{CP}$  inversion (Equation 2.112) expressions require:

$$P_L(\nu_\alpha \rightarrow \nu_\beta) = P_L(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta) = P_L(\nu_\beta \rightarrow \nu_\alpha) = P_L(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha). \quad (2.138)$$

Since these symmetry expressions assume only  $\mathcal{CPT}$  conservation and the expressions derived for  $P_L$ , any violation of them is evidence that either  $\mathcal{CPT}$  is not conserved or that the oscillation probability needs to be modified in a way that is different for neutrinos and antineutrinos.

In many experimental situations, including the one discussed in this thesis, the two-neutrino approximation is a good one. The experimental setup here is looking at muon neutrino or antineutrino survival at distances and energies corresponding to the atmospheric oscillation length. In practice, the approximation is actually two approximations: that the oscillations driven by the solar mass-splitting are unimportant and that only one mixing angle is relevant.

The first approximation is a good one since the measured neutrino mass-splittings (discussed later in Section 2.3.6) differ by more than an order of magnitude, giving them dramatically different oscillation lengths. The full three-neutrino  $\nu_\mu$  survival probability has two terms, one proportional to  $\sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right)$  and one proportional to  $\sin^2\left(\frac{\Delta m_{\text{sol}}^2 L}{4E}\right)$ . If  $L$  and  $E$  are chosen such that  $\sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right) \approx 1$ , then because  $\Delta m_{\text{atm}}^2/\Delta m_{\text{sol}}^2 \approx 30$ ,  $\sin^2\left(\frac{\Delta m_{\text{sol}}^2 L}{4E}\right) \approx 0.003$  and because the atmospheric mixing angle is large, the solar oscillations are negligible, independent of the size of the solar mixing angle.

The second approximation says, effectively, that only  $\theta_{23}$  drives the mixing and that  $\theta_{13}$  is not relevant. The full three-flavor coefficient to the  $\Delta m_{\text{atm}}^2$  term is  $\sin^2(2\theta_{23})\cos^2(\theta_{13}) + \sin^2(2\theta_{13})\cos^2(\theta_{23})$ . However, no experiment to date has yet been able to distinguish  $\sin^2(2\theta_{23})$  from 1 or  $\sin^2(2\theta_{13})$  from 0, so this term is effectively  $\sin^2(2\theta_{23})$ , corresponding to the two-neutrino approximation with mass states 2 and 3.

### 2.3.5 Matter Effects

The transitions among neutrino flavors depends only on the neutrino mass matrix when the neutrinos are travelling through vacuum. When the neutrinos are travelling through matter the effect of coherent forward scattering must be taken into account, as first observed by Lincoln Wolfenstein in 1978 [58]. Electron neutrinos can interact with electrons in the medium via charged-current interactions and all neutrinos can interact with electrons, protons, and neutrons via neutral-current interactions. However, this NC matter-effect, like all NC interactions, is  $e-\mu-\tau$  symmetric, meaning it only adds terms proportional to the identity, which cannot be observed in an oscillation experiment since they only introduce an overall phase to the neutrino state, not a phase difference between flavor components.

The effect of the charged-current scattering can be demonstrated in the simpler case of two-neutrino oscillations. Begin by writing the Schrödinger equation for the time evolution of the neutrino state in the mass basis

$$i \frac{d}{dt} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = H_{\text{vac}} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \quad (2.139)$$

where, for the mass eigenstates, the Hamiltonian is easily determined

$$H_{\text{vac}} = \begin{pmatrix} E_1 & 0 \\ 0 & E_2 \end{pmatrix} \approx p + \begin{pmatrix} m_1^2/2E & 0 \\ 0 & m_2^2/2E \end{pmatrix} \quad (2.140)$$

where  $p$  is the neutrino momentum,  $E$  is the neutrino energy, and  $m_1, m_2$  are the masses of  $\nu_1, \nu_2$  as was done in Section 2.3.2. Now, this expression can be transformed into the weak basis using the mixing matrix,  $U$ ,

$$\begin{pmatrix} \nu_e \\ \nu_\beta \end{pmatrix} = U \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \quad (2.141)$$

where one of the weak states has been assumed to be the electron neutrino. Substituting and rearranging, Equation 2.139 becomes

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\beta \end{pmatrix} = U H_{\text{vac}} U^\dagger \begin{pmatrix} \nu_e \\ \nu_\beta \end{pmatrix}. \quad (2.142)$$

The charged-current scattering of electron neutrinos off electrons in the medium adds an additional potential

$$V_m^e = \pm \sqrt{2} G_F n_e \quad (2.143)$$

where  $G_F$  is the Fermi constant,  $n_e$  is the density of electrons in the medium, and the term changes from positive to negative for antineutrinos. Since the interaction leaves the neutrino flavor unchanged and only effects  $\nu_e$ 's, it only appears in the  $e - e$  term of the Hamiltonian in the weak basis. The total effective Hamiltonian then becomes

$$H = U H_{\text{vac}} U^\dagger + H_{\text{CC}} = \frac{\Delta m^2}{2E} \begin{pmatrix} \sin^2 \theta + \frac{2E}{\Delta m^2} V_m^e & -\sin \theta \cos \theta \\ -\sin \theta \cos \theta & \cos^2 \theta \end{pmatrix}, \quad (2.144)$$

where terms proportional to the identity have been removed. This modified Hamiltonian can then be re-diagonalized to give a new mixing angle  $\theta^m$  and energy eigenvalues  $E_{1,2}^m$ ,

$$f_m = \sqrt{\sin^2 2\theta + \left( \cos 2\theta - \frac{2E}{\Delta m^2} V_m^e \right)^2} \quad (2.145)$$

$$E_{1,2}^m = \pm \frac{\Delta m^2}{4E} f_m \quad (2.146)$$

$$\cos 2\theta^m = \frac{1}{f_m} \left( \cos 2\theta - \frac{\Delta m^2}{2E} V_m^e \right) \quad (2.147)$$

$$\sin 2\theta^m = \frac{1}{f_m} \sin 2\theta. \quad (2.148)$$

In the end, the oscillation formula takes the same form, but with additional  $f_m$  factors to account

for the matter effects,

$$P(\nu_e \rightarrow \nu_e) = 1 - \frac{\sin^2 2\theta}{f_m^2} \sin^2 \left( \frac{\Delta m^2 f_m L}{4E} \right). \quad (2.149)$$

When multiplied out, the term  $\Delta m^2 f_m$  will have a term of the form  $(\Delta m^2 \cos 2\theta - 2EV_m^e)^2$ . Since the matter effects add or subtract from  $\Delta m^2$  directly, not its absolute value, the sign of  $\Delta m^2$  can be determined if matter effects can be observed.

Notice that neutrino oscillations in matter can be maximal, regardless of the value of the actual mixing angle, given the right electron density. When

$$n_e = \frac{\Delta m^2 \cos 2\theta}{2\sqrt{2}G_F E}, \quad (2.150)$$

$\sin^2 2\theta^m = 1$ , independent of the value of  $\theta$ . When the electron density is non-constant, as happens in the Sun, another resonance condition can occur called the MSW effect, after Stanislav Mikheyev and Alexei Smirnov, who first calculated the resonance, and Lincoln Wolfenstein from above [59]. In the Sun, the electron density decreases exponentially from the center, where neutrinos are produced, out to the edge. As the neutrinos travel along this density curve, they will reach a point where the matter effects become resonant, causing all the neutrinos to oscillate maximally into one of the eigenstates of the Hamiltonian. After undergoing this adiabatic transition, the solar neutrinos are in an effectively stationary state and their flavor content remains approximately unchanged for the rest of their trip to a detector on Earth.

### 2.3.6 Measurements

After the initial discovery of neutrino oscillations in 1998 by Super-Kamiokande, new experiments began to come online which could make precise measurements of neutrino mixing. There are six parameters that can be measured in neutrino oscillation experiments:  $\theta_{12}, \Delta m_{12}^2, \theta_{23}, \Delta m_{23}^2, \theta_{13}, \delta$ . The first four parameters have been measured<sup>10</sup> with precision, an upper limit has been placed on the fifth, and the sixth is not yet constrained.

The earliest solar neutrino experiments, the Homestake experiment discussed above [48] and the Kamiokande water Cherenkov experiment in Japan [60], were sensitive only to the highest energy solar neutrinos coming from the decay of  ${}^8\text{B}$ ,

$${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e (\approx 10 \text{ MeV}), \quad (2.151)$$


but these neutrinos are only a tiny fraction of the solar neutrino flux and their rate is model-dependent. They observed solar neutrino rates that were  $28\% \pm 5\%$  and  $46\% \pm 8\%$ , respectively, of those predicted by the standard solar model. The later radiochemical experiments, SAGE in Soviet

<sup>10</sup>With the notable exception of the sign of  $\Delta m_{23}^2$ , which remains unknown.

Union [61] and GALLEX-GNO in Italy [62], used the reaction



to observe solar neutrinos, which has a significantly lower threshold of 0.223 MeV. Consequently, these experiments were sensitive to the  $pp$  neutrinos produced in the beginning of the main solar fusion reaction



which are lower in energy but make up the bulk of the solar neutrino flux. Their rate, unlike that of the  ${}^8\text{B}$  neutrinos, is mostly model-independent. The standard solar model predicts a rate of 128 SNU on Gallium but Sage and Gallex measured rates of  $70.8_{-6.1}^{+6.5}$  SNU and  $77.5_{-7.8}^{+7.5}$  SNU respectively.<sup>11</sup> However, even when the results from all these experiments were combined, the oscillation parameters could not be constrained to a single region of  $\Delta m_{12}^2$ - $\theta_{12}$  parameter space (the 12 sector is believed to be responsible for transitions from  $\nu_e$  to  $\nu_\mu$  and  $\nu_\tau$ ). The isolation of a single pair of oscillation parameters came when the results of the SNO experiment were released [63], isolating the Large Mixing Angle (LMA) solution with MSW matter effects. The combined best oscillation fit to all solar neutrino data, circa 2005, gives [64]

$$\Delta m_{12}^2 = (6.5_{-2.3}^{+4.4}) \times 10^{-5} \text{ eV}^2, \quad \tan^2 \theta_{12} = 0.45_{-0.08}^{+0.09}. \quad (2.154)$$

Additionally, by measuring the sign of the MSW effect, it was established that  $m_2 > m_1$  [35].

Final confirmation of neutrino oscillations in the ‘solar sector’ (where the 12 oscillations are dominant) came from the KamLAND experiment. KamLAND is a  $\bar{\nu}_e$  detector situated in the Kamioka mine in Japan which uses 1 kTon of liquid scintillator surrounded by non-scintillating buffer oil and is instrumented with approximately 1,900 phototubes. It detects neutrinos via inverse beta decay with the prompt-delayed double coincidence used in the earliest neutrino experiments. It is exposed to a flux of  $\bar{\nu}_e$ ’s from 55 commercial nuclear reactors, which are an average of 180 km from the detector. KamLAND saw significant  $\bar{\nu}_e$  disappearance [65] which, when fit for oscillations, gave

$$\Delta m_{12}^2 = (7.66_{-0.20}^{+0.22}) \times 10^{-5} \text{ eV}^2, \quad \tan^2 \theta_{12} = 0.52_{-0.10}^{+0.16} \quad (2.155)$$

which are clearly consistent with the parameters measured by the solar experiments in Equation 2.154. The contours from KamLAND and the combined solar experiments can be seen in Figure 2.7. This independent measurement of similar oscillation parameters in a completely different experimental set up ( $\bar{\nu}_e$  vs.  $\nu_e$ , different energy and oscillation length, different detector technology) significantly increased confidence in neutrino oscillations. When the solar and KamLAND results

<sup>11</sup>SNU denotes a Solar Neutrino Unit which is defined as  $10^{-36}$  events/atom/second.




Figure 2.7: The oscillation parameter regions allowed by the KamLAND experiment (solid colors) and the combined solar experiments (black lines). The side plots show the two one-dimensional profiles for KamLAND (dashed), solar (dotted), and combined (solid). The consistency between the solar and reactor measurements lends significant support to the oscillation model of neutrino disappearance. Figure taken from [65].

are combined, the resulting measurement has a precision of better than 3.5% [35]:

$$\Delta m_{12}^2 = (7.59 \pm 0.20) \times 10^{-5} \text{ eV}^2, \quad \tan^2 \theta_{12} = 0.47_{-0.05}^{+0.06}. \quad (2.156)$$

After the initial discovery of oscillations, Super-Kamiokande continued to run and improve its measurements of oscillations in atmospheric neutrinos. The most recent result, based on data from phases SK-I, SK-II, and SK-III and including the possibility of sub-leading ( $\theta_{13}$ ) oscillation effects, find at 90% confidence assuming  $m_3 > m_{1,2}$  (the ‘normal hierarchy’) [66],

$$1.9 \times 10^{-3} \leq |\Delta m_{\text{atm}}^2| \leq 2.6 \times 10^{-3} \text{ eV}^2 \quad 0.407 \leq \sin^2 \theta_{23} \leq 0.583 \quad (2.157)$$

with central values  $\Delta m_{\text{atm}}^2 = 2.1 \times 10^{-3} \text{ eV}^2$ <sup>12</sup> and  $\sin^2 \theta_{23} = 0.5$ .

The atmospheric oscillations were confirmed by experiments using  $\nu_\mu$ ’s from terrestrial accelerators, just as the solar results were confirmed by the KamLAND reactor experiment. The experiments were designed with neutrino sources and detectors separated by approximately a quarter of the atmospheric oscillation length, the distance over which a neutrino with energy  $E$  will oscillate back to its original state. In a sense the oscillation length is the ‘wavelength’ of the oscillation, and it is

<sup>12</sup>The atmospheric mass-splitting is referred to as  $\Delta m_{\text{atm}}^2$  since it is, in fact, a combination of  $\Delta m_{23}^2$  and  $\Delta m_{13}^2$  which are too close to each other to be easily distinguished.




Figure 2.8: The allowed oscillation parameters at 90% confidence from MINOS in black [67], Super-Kamiokande in red and blue [68], and K2K in purple [69].

given by,

$$L_{\text{atm}} = 4\pi \frac{E\hbar c}{|\Delta m_{\text{atm}}^2|c^4} \approx 2.5 \frac{E}{|\Delta m_{\text{atm}}^2|} \text{ m}, \quad (2.158)$$

where  $E$  is the neutrino energy in MeV and  $|\Delta m_{\text{atm}}^2|$  is the neutrino mass-splitting in  $\text{eV}^2$ . The first confirmation came from the K2K experiment in Japan where  $\nu_\mu$ 's produced at approximately 1 GeV at the KEK accelerator were observed in the Super-Kamiokande detector 250 km away ( $L_{\text{atm}} \approx 1,000$  km for  $E = 1$  GeV with the  $|\Delta m_{\text{atm}}^2|$  measured by Super-Kamiokande).

The most precise accelerator measurement to date is from the MINOS experiment, which uses neutrinos with  $E \approx 3$  GeV and a baseline of 735 km and is described in great detail in the next chapter. After analyzing an exposure of  $7.2 \times 10^{20}$  protons-on-target, the atmospheric parameters were measured to [67]:

$$|\Delta m_{\text{atm}}^2| = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{ eV}^2 \quad \sin^2(2\theta_{23}) > 0.90 \text{ (90\% C.L.)} \quad (2.159)$$

where, again, the central value of the mixing angle is maximal.<sup>13</sup> MINOS has the most precise measurement (better than 5% precision) of the mass splitting because of its good energy resolution. Super-Kamiokande continues to have the best measurement of the mixing angle due to its large sample of oscillated neutrino events. Note that none of the atmospheric-sector experiments to date are sensitive to the sign of  $\Delta m_{\text{atm}}^2$ , only its absolute value  $|\Delta m_{\text{atm}}^2|$ , since matter effects do not alter the  $\nu_\mu - \nu_\tau$  transition.

<sup>13</sup>Maximal mixing, which means the disappearance probability can reach unity, occurs when  $\sin^2 \theta_{23} = 0.5$ ,  $\sin^2(2\theta_{23}) = 1$ ,  $\theta_{23} = 45^\circ$ .

The final mixing angle,  $\theta_{13}$ , has not yet been determined definitely to be nonzero in any experiments. If it is non-zero it will drive the disappearance of electron (anti)neutrinos at the atmospheric oscillation length (Equation 2.158) with survival probability given by:

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \sin^2 2\theta_{13} \sin^2 \left( 1.27 \Delta m_{\text{atm}}^2 \frac{L}{E} \right), \quad (2.160)$$

where  $L$  is in meters and  $E$  is in MeV. The CHOOZ experiment, in northern France, searched for the disappearance of reactor electron antineutrinos ( $E \approx 3.6$  MeV). The experiment searched for inverse beta-decay in a Gd-doped (to enhance neutron capture) liquid scintillator detector situated 1 km from the two reactors at the CHOOZ nuclear power plant. No electron neutrino disappearance was observed, and hence only an upper limit could be set on the value of the last mixing angle [35],

$$\sin^2 2\theta_{13} < 0.15 \text{ (90\% C.L.)}. \quad (2.161)$$

The last PMNS parameter accessible in oscillation experiments, the  $\mathcal{CP}$ -phase  $\delta$ , has not yet been constrained at all. There is great interest in measuring  $\delta$  since a non-zero value means that neutrinos are  $\mathcal{CP}$ -violating, meaning they might have played a role in producing the matter-antimatter asymmetry in the early universe, a process called ‘leptogenesis’ [70]. In Equation 2.131,  $\delta$  always appears connected to  $\sin \theta_{13}$  and so cannot be measured unless  $\theta_{13} > 0$ . Furthermore, the size of any  $\mathcal{CP}$ -violating effect will depend directly on the value  $\sin^2(2\theta_{13})$  since the terms always appear together. Measuring the sign of  $\Delta m_{\text{atm}}^2$  also depends on the value of  $\theta_{13}$  since this is the parameter governing  $\nu_e$  appearance over atmospheric oscillation lengths. It is necessary to observe  $\nu_e$  appearance since matter effects only influence oscillations involving  $\nu_e$ ’s, and only matter effects are sensitive to the sign of  $\Delta m_{\text{atm}}^2$ .

Consequently, all of the next-generation neutrino oscillation experiments commencing operations now or in the near future are focused on measuring  $\theta_{13}$ . There are three reactor experiments which, like CHOOZ, measure electron antineutrino disappearance, giving them direct access to the value of  $\sin^2(2\theta_{13})$ . These experiments are Double CHOOZ [71] in France, Daya Bay [72] in China, and RENO [73] in Korea. They all use the same basic detection technique: Gd-doped liquid scintillator detectors which detect antineutrinos via inverse beta decay, placed at both near and far locations. The experiments vary in the number detectors, their positions relative to the reactors, and the details of the detector construction. In addition to the reactor experiments, there are also two accelerator-based experiments searching for electron neutrino appearance rather than disappearance: T2K [74] in Japan and NO $\nu$ A [75] in the United States. T2K uses the existing Super-Kamiokande detector with a new neutrino beam from Tokai, approximately 300 km away. NO $\nu$ A, on the other hand, is building a new fully-active detector made up of cells of liquid scintillator optimized for observing  $\nu_e$  events 800 km away from the existing NuMI beamline at Fermilab. The appearance experiments,



rather than making a pure  $\sin^2(2\theta_{13})$  measurement, instead measure a more complicated appearance probability,  $P(\nu_\mu \rightarrow \nu_e)$ , that depends on several parameters, including the  $\mathcal{CP}$ -phase  $\delta$  and mass hierarchy. The different experiment types are thus complimentary, like KamLAND and the solar experiments or MINOS and Super-Kamiokande.

## 2.4 Antineutrino Oscillations

Oscillations in the solar sector have been studied in great detail with both  $\nu_e$ 's from the Sun and  $\bar{\nu}_e$ 's from nuclear reactors. As can be seen in Figure 2.7, the solar sector neutrino and antineutrino oscillations are in good agreement. The atmospheric sector measurements, on the other hand, have been performed primarily with  $\nu_\mu$ 's, in the case of accelerator neutrinos, or a mixed sample coming from the atmosphere. However, even the mixed sample is primarily neutrinos since they have twice the interaction cross section of antineutrinos at high energies and up to 4 times the cross section at low energies.

Before the MINOS measurements described in this thesis, there were no direct, precise measurements of atmospheric sector oscillations in antineutrinos. The NuMI beamline was running in neutrino mode and measurements were only made with neutrino-like events. Super-Kamiokande had antineutrino events, but they were the minority. Plus, whatever the fraction of antineutrino events, the Super-Kamiokande water Cherenkov detector is unable to distinguish individual neutrinos from

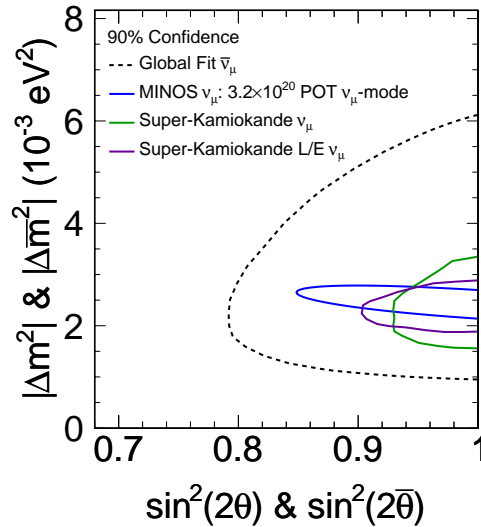


Figure 2.9: The black dashed line is a profile in  $|\Delta\bar{m}_{\text{atm}}^2|$  vs.  $\sin^2(2\bar{\theta}_{23})$  from a global fit to all neutrino and antineutrino data, from [76]. The colored  $|\Delta m_{\text{atm}}^2|$  vs.  $\sin^2(2\theta_{23})$  contours are from MINOS [77] and Super-Kamiokande [68] circa 2008. The combined global fit for antineutrinos is significantly broader than any individual measurement for the neutrinos. The global fit does not include the MINOS antineutrino results which are presented in the rest of this thesis.

antineutrinos, so measurements could not be made on a pure sample.

Even so, some constraints were able to be placed on the antineutrino oscillation parameters,  $|\Delta\bar{m}_{\text{atm}}^2|$  and  $\sin^2(2\bar{\theta}_{23})$ , in a global fit to all neutrino oscillation data. The fit includes data from many solar, reactor, atmospheric, and accelerator experiments, but most of the constraint on these parameters comes from the Super-Kamiokande data. The global fit 90% C.L. contour, made in 2008, can be seen in Figure 2.9, compared to three individual measurements of the neutrino parameters from contemporary experiments. Each individual neutrino experiment places significantly better limits than the combined global fit to the antineutrinos.

## 2.5 Differing Neutrino and Antineutrino Oscillations

This thesis focuses on measurements of flavor oscillations via the disappearance of muon antineutrinos at an  $L/E \approx 250$  km/GeV. At this distance, the disappearance probability is governed, to a good approximation, by two-neutrino oscillation with  $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ . With only two neutrinos, the mixing matrix  $U$  must be real and ‘standard’ neutrino  $\mathcal{CP}$ -violation via the phase  $\delta$  cannot influence the disappearance probabilities,  $P_L(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu)$  and  $P_L(\nu_\mu \rightarrow \nu_\mu)$ . Thus, as was shown in Equation 2.138, if  $P_L(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu) \neq P_L(\nu_\mu \rightarrow \nu_\mu)$  then either the neutrinos are not  $\mathcal{CPT}$  invariant or some other new physics needs to be introduced.

### 2.5.1 Non-standard Interactions

Standard matter-effects do introduce differences between neutrinos and antineutrinos since the additional components of the Hamiltonian have opposite signs. However, they cannot influence the  $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$  transition which does not involve  $\nu_e$ ’s. However, some other physics process could introduce additional Hamiltonian terms in the same way. Generically, these processes are called ‘non-standard interactions’ (NSI). A number of models have been proposed which introduce effects that, like matter effects, have opposite signs for neutrinos and antineutrinos. Examples include a flavor-dependent potential introduced by very light  $Z'$  bosons from a gauged  $L_\mu - L_\tau$  symmetry [78]; a new charged-current coupling between  $\nu_\tau$ ’s and leptons, constructed so that flavor-violating contributions to  $\tau$  decays are sufficiently suppressed [79]; and inclusion of a sterile neutrino, making the neutral-current matter effects non-trivial and/or introducing a new  $U(1)$  gauge force coupled to  $B - L$  [80].

A generic, model-independent formalism for NSI can be found in [81]. Non-standard interactions, whatever their source, will add a term to the Lagrangian of the form

$$\mathcal{L}_{\text{NSI}} = -\frac{G_F}{\sqrt{2}} \sum_{\substack{f=u,d,e \\ a=\pm 1}} \varepsilon_{\alpha\beta}^{fa} \left[ \bar{\nu}_\alpha \gamma^\mu (1 - \gamma^5) \nu_\beta \right] \left[ \bar{f} \gamma^\mu (1 + a \gamma^5) f \right] \quad (2.162)$$

where the various  $\varepsilon_{\alpha\beta}^{fa}$  coefficients give the strengths of the NSI effects. Analogous to standard matter effects, this NSI Lagrangian will add an additional term to the effective Hamiltonian with the form

$$H_{\text{NSI}} = V \begin{pmatrix} \varepsilon_{ee} & \varepsilon_{e\mu} & \varepsilon_{e\tau} \\ \varepsilon_{e\mu}^* & \varepsilon_{\mu\mu} & \varepsilon_{\mu\tau} \\ \varepsilon_{e\tau}^* & \varepsilon_{\mu\tau}^* & \varepsilon_{\tau\tau} \end{pmatrix} \quad (2.163)$$

where  $V = \sqrt{2}G_F N_e$  is the MSW potential and the NSI parameters have been summed over the various fermion contributions,

$$\varepsilon_{\alpha\beta} = \sum_{f,a} \varepsilon_{\alpha\beta}^{fa} \frac{N_f}{N_e} \quad (2.164)$$

where  $N_f$  is the number density of fermion  $f$  in the (unpolarized) medium.

There have been numerous analyses of the existing neutrino data searching for evidence of NSI. The existing constraints are summarized in [82]. The constraints are for neutrinos only, and the constraints get tighter (usually by an order of magnitude) when the NSI effects are included as part of a renormalizable theory and their effects on the charged leptons are accounted for. There has even been an analysis of the data presented in this thesis which finds a value for  $\varepsilon_{\mu\tau} = -0.12 \pm 0.21$  [83].

### 2.5.2 $CPT$ Violation

$CPT$  is a fundamental symmetry in quantum field theory.  $CPT$  invariance is required in order to have Lorentz-invariant models, and consequently any theory that includes  $CPT$  violation also, automatically, includes Lorentz violation [84].  $CPT$  invariance also requires that particles and antiparticles share certain properties, including charge and mass. If different mass-splittings are measured for neutrinos and antineutrinos, the naive interpretation that the neutrinos and antineutrinos have different masses thus requires  $CPT$  violation [85].

$CPT$  violation in the neutrino sector first came to the forefront as an explanation for the LSND anomaly. The Liquid Scintillator Neutrino Detector (LSND) at Los Alamos searched for  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  oscillations using  $\bar{\nu}_\mu$ 's from muon decays at rest. They reported evidence of a mass splitting  $\Delta m^2 \approx 1 \text{ eV}^2$ , inconsistent with the three-neutrino model and the two mass-splittings,  $\Delta m_{\text{atm}}^2$  and  $\Delta m_{\text{sol}}^2$ , that had been previously measured [86]. A number of  $CPT$ -violating theories were proposed to explain this discrepancy [87, 88, 89, 90], but the possible  $CPT$ -violating effects become much more tightly constrained once effects outside the neutrino sector are taken into account [91, 92]. Some theories endeavor to simultaneously explain LSND and the MINOS antineutrino measurements (discussed in Chapter 5) [93]. However, no theory has yet proven compelling and predictive.

Some theorists, primarily Alan Kostelecký and his colleagues at Indiana University, have studied

possible Lorentz/ $\mathcal{CPT}$ -violating models in great detail, including their possible influence on neutrino experiments [94]. They even have a model that explains MINOS in particular [95], although for now the model is more proof-of-principle and has too many parameters to be predictive. Using the Kostelecký parameterization of possible Lorentz/ $\mathcal{CPT}$  violation, numerous analyses have been performed on experiments in wide-ranging areas of physics including, for example, neutral meson oscillations [96, 97, 98], comparisons of clocks [99], muon spin precession [100], electron-positron g-2 [101], and many others. Searches have been performed in MINOS data for sidereal variations in the neutrino rate, another signature of Lorentz-violation in the Kostelecký model [102, 103]. There has been, to date, no convincing evidence of  $\mathcal{CPT}$  violation in any sector.

## 2.6 Neutrino-to-Antineutrino Transitions

In addition to standard oscillations, this thesis also presents a search for  $\nu_\mu \rightarrow \bar{\nu}_\mu$  transitions. Transitions between neutrinos and antineutrinos were Bruno Pontecorvo's original idea for neutrino oscillations in the 1960's [15]. The idea resurfaced in the literature around 1980 in the context of understanding the consequences of Majorana neutrinos [104, 105]. More recently, several possible models for neutrino-to-antineutrino transitions within the simplest Standard Model extensions for neutrino mass are catalogued in [106]. The transitions can be helicity-violating, mediated either by a Majorana mass term or by a magnetic moment, or they can be helicity-conserving, producing a sterile left-handed antineutrino. The analysis in this thesis is sensitive only to transitions like the ones mediated by the Majorana mass: the helicity-conserving transitions produce sterile neutrinos which cannot be observed in the MINOS detectors and the magnetic moment transition requires a change of flavor (e.g.  $\nu_\mu \rightarrow \bar{\nu}_e$ ). However, since the Majorana mass-mediated transition requires creating wrong-helicity states, its amplitude is suppressed by a factor proportional to  $(m_\nu/E_\nu)^2$  and consequently is expected to occur below the  $10^{-7}$  level in a high-energy  $\nu_\mu$  beam, making the process effectively unobservable [107]. Thus, any observation of  $\nu_\mu \rightarrow \bar{\nu}_\mu$  would require new physics.

Such transitions could be introduced by the V+A currents that arise in left-right symmetric models [108], but stringent limits have been set on leptonic V+A interactions by studies of the end point spectrum in polarized  $\mu^+$  decay [109], by studies of inverse muon decay [110], and by studies of the high- $y$  dependence of  $\nu_\mu/\bar{\nu}_\mu$ -nuclei interactions [111, 112]. Scalar (S) and pseudoscalar (P) interactions can also introduce spin-flips and change the lepton helicity [113], but contributions from these interactions were limited to less than 7% (95% C.L.) by investigations of the polarization of  $\mu^+$ 's in  $\bar{\nu}_\mu$  interactions at the CHARM experiment [114]. Neutrino-to-antineutrino transitions could also be introduced by certain  $\mathcal{CPT}$ -violating parameters in the Kostelecký parameterization [94].

In the early 1980's the BEBC bubble chamber in the CERN SPS neutrino beam was able to set limits on transitions to  $\bar{\nu}_e$  from  $\nu_\mu$  and  $\nu_e$  [115]. The best limit on anomalous  $\mu^+$  production

in a  $\nu_\mu$  beam comes from the CCFR experiment: the fraction of  $\mu^+$  relative to the  $CC\nu_\mu$  rate was limited to  $1.6 \times 10^{-4}$  for  $y < 0.5$  and  $3.1 \times 10^{-4}$  for  $y > 0.5$  [116]. However, there have been no measurements of transitions to  $\bar{\nu}_\mu$ 's at atmospheric oscillation length scales, or at such low energies: the CCFR experiment used 120 GeV neutrinos and a baseline of 1.1 km. Searching for  $\nu_\mu \rightarrow \bar{\nu}_\mu$  requires a detector capable of identifying individual  $\mu^+$ 's among many  $\mu^-$ 's produced by the muon neutrino beam, a unique capability of MINOS among long-baseline experiments. In fact, a recent analysis used the limit described later in this thesis to improve limits on the effective Majorana muon-neutrino mass,  $|\langle m_{\mu\mu} \rangle|$  [117].

The search for neutrino-to-antineutrino transitions in this thesis uses an empirical parameterization, based on the knowledge that  $\nu_\mu$ 's are known to be disappearing [67] in the NuMI beam with an energy-dependence described by Equation 2.136 and the supposition that some fraction,  $\alpha$ , of those  $\nu_\mu$ 's are transitioning to  $\bar{\nu}_\mu$ 's instead of oscillating to  $\nu_\tau$ 's. Thus, the appearance probability takes on the form

$$P(\nu_\mu \rightarrow \bar{\nu}_\mu) = \alpha \sin^2(2\theta_{23}) \sin^2\left(\Delta m_{\text{atm}}^2 \frac{L}{4E}\right) \quad (2.165)$$

Note that the oscillation parameters above are those for neutrinos, not those for antineutrinos.



## Chapter 3

# The MINOS Experiment

The Main Injector Neutrino Oscillation Search (MINOS) is a long-baseline neutrino-oscillation experiment. Its main components are the NuMI neutrino beam and two detectors. The NuMI beam is located at the Fermi National Accelerator Laboratory (Fermilab or FNAL), where 120 GeV protons from the Main Injector are directed at a fixed target to produce mesons which decay to produce the neutrino or antineutrino beam. Oscillations are measured by sampling the primarily  $\nu_\mu$  or  $\bar{\nu}_\mu$  beam at two locations: one close to the neutrino source before oscillations have occurred (the Near Detector) and one located approximately a quarter of the atmospheric oscillation length (see Equation 2.158) away, where oscillations will be near maximal (the Far Detector). Like the NuMI beam, the Near Detector is situated at Fermilab, 1 km downstream of the neutrino target. The Far Detector is situated 735 km from the neutrino source in the Soudan Underground Laboratory in northern Minnesota.

The two-detector design makes the measurement of oscillations less dependent on simulation and significantly more robust against a range of systematic uncertainties, especially the neutrino flux from NuMI and the poorly known low-energy neutrino interaction cross section. Since these




Figure 3.1: Schematic views of the components of the MINOS experiment: the NuMI beam and Near Detector at Fermilab and the Far Detector in Soudan, MN.




Figure 3.2: Cross-sectional view of the NuMI beam complex and MINOS Near Detector Hall at Fermilab. Figure taken from [118].

systematics affect both detectors in the same way, they are effectively cancelled out when both detectors are used to measure oscillations. Section 5.3 describes how this process works in practice.

The rest of this chapter gives detailed descriptions of the main components of the experiment. It describes the MINOS data, from the raw detector output through calibration and reconstruction. It also includes a description of how the data is simulated.

## 3.1 The NuMI Beam

### 3.1.1 Making Neutrinos

Neutrinos at the Main Injector (NuMI) [118] is a conventional [119] neutrino beam, where primary protons collide with a fixed target to produce secondary mesons which decay to tertiary neutrinos. The 120 GeV protons are extracted from the FNAL Main Injector (MI) by three horizontal kickers and three Lambertson magnets.<sup>1</sup> The MI loads five batches from the Booster and then loads six more into the gaps between the first five but with slightly different momenta. The second set of batches then “slips” relative to the first until they overlap in a technique called “multi-batch slip stacking” [120]. Typically, two batches are dedicated to the Tevatron collider [121] and the other nine are extracted to NuMI [122]. The whole process repeats every 2.2 seconds, decreased from 2.4 seconds or longer when NuMI began. Each set of batches has  $36 \times 10^{12}$  protons when NuMI is sharing the MI with the Tevatron, and more than  $40 \times 10^{12}$  protons per pulse when the MI is in NuMI-only mode.<sup>2</sup> The primary proton beam is then focused and bent downwards into the Earth at 58 mrad (towards Soudan, MN) by a series of quadrupole and bending magnets. As the beam is focused and directed towards the target, it passes through a series of position monitors, loss monitors, and proton

<sup>1</sup>Named for its inventor at Fermilab, a Lambertson is a special magnet with two bores, one with a bending field and one field-free, designed for separating adjacent beams.

<sup>2</sup>These proton intensities are almost double the original design intensity of  $25 \times 10^{12}$  protons per pulse.






Figure 3.3: Drawing of a NuMI target and its housing.

counters before reaching a graphite collimating baffle just before the target. The baffle protects the downstream beam components from being damaged by the beam halo. When the beam reaches the target, it has an rms spot size of 1.1 mm. A cross-sectional view of the whole Fermilab facility can be seen in Figure 3.2.

The target is made up of 47 vertical graphite fins, each 6.4 mm wide, 15 mm tall, and 20 mm long with 0.3 mm spacing. They are arranged longitudinally to form a total target length of 954 mm, approximately 1.9 hadronic interaction lengths. A 48<sup>th</sup> fin is mounted horizontally upstream of the main target to help align the beam in target scans. The graphite fins are water cooled by pipes running along the top and bottom of each fin. A drawing of the target and its housing can be seen in Figure 3.3. When the primary protons collide with the target, they produce a secondary beam of unfocused pions and kaons,

$$p + C \rightarrow h^{\pm} + X, \quad h = \{\pi, K\}. \quad (3.1)$$

The NuMI targets have a limited lifetime. Even if a target does not develop an obvious problem, such as a water leak, the graphite appears to delaminate over time, losing density in the region around the shower maximum, where the most primary beam energy is absorbed. This degradation gradually leads to a several percent drop in the neutrino flux in the focusing peak. NuMI target NT-02, which was installed between Runs I and II, began to show the effects of degradation about half way through Run II and throughout Run III.

The flux of neutrinos is enhanced by focusing the secondary beam using two magnetic focusing horns [123], shown in Figure 3.4. A current<sup>3</sup> of 185 kA passes along the outer and inner surfaces

<sup>3</sup>185 kA is the typical operating current in the LE configuration. The maximum possible current is 200 kA.





Figure 3.4: Above, a design drawing of Horn 2. Below, a picture of a NuMI horn being installed.




Figure 3.5: A cartoon showing how the horn focusing enhances the flux of neutrinos by bending positive mesons towards the beam axis. In addition to the focused mesons, there are also mesons of both charge signs that pass through the centers of the horns, avoiding focusing.

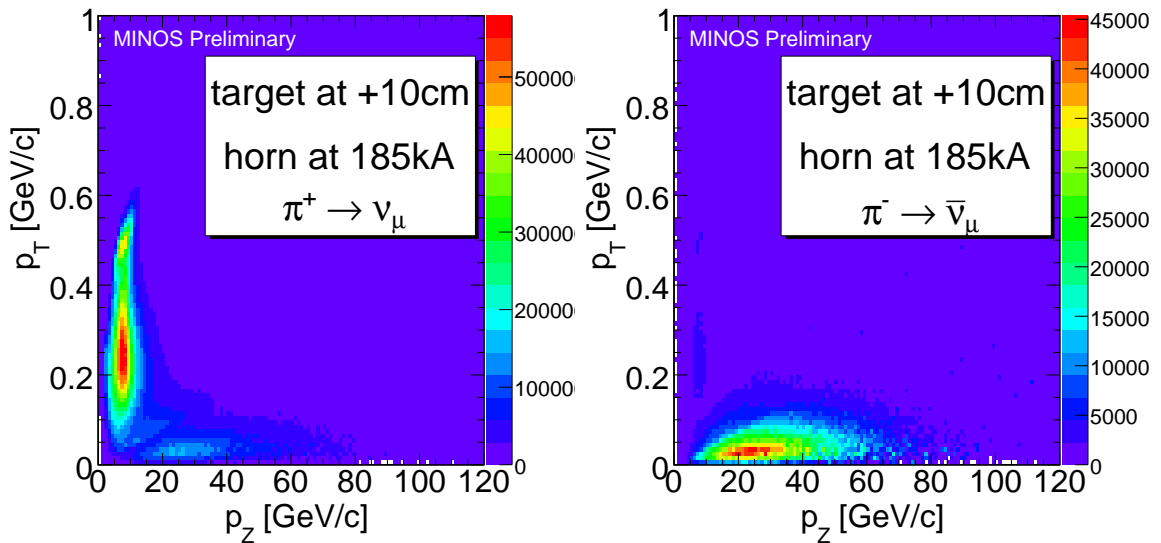


Figure 3.6: The  $p_T$  vs.  $p_Z$  distribution of the  $\pi^\pm$  parents that produce neutrinos (left) and antineutrinos (right) at the Near Detector when the beam is in low-energy neutrino-mode mode. The unfocused component has a broad range of  $p_Z$ , and hence total momentum, producing the diffuse high-energy tail.

(called ‘conductors’) of the horn in opposite directions, producing a toroidal magnetic field between the horn surfaces with peak intensity of 2.8 T. This magnetic field serves to sign- and momentum-select the mesons produced in the target by bending mesons toward or away from the beam axis. The horns act like point-to-parallel focusing lenses, with focal lengths proportional to particle momentum, thanks to the parabolic shape of the inner conductors. When a meson is bent towards the axis, the likelihood that its daughter neutrino will reach a MINOS detector is increased since the daughter neutrino will preserve much of the parent’s forward boost after the parent decays. Figure 3.5 shows a cartoon of the horn focusing process.

The focusing horns selectively focus mesons from a particular region of  $p_T - p_Z$  momentum space (see Figure 3.6, left), producing the focusing peak at a particular energy. The peak energy of the beam can be adjusted by changing either the relative distance between the two horns or the




Figure 3.7: The neutrino spectrum produced when the beam is in low energy (LE) mode with the target at its furthest downstream position, pseudo-medium energy (pME) mode with the target moved 1 m upstream, and pseudo-high energy (pHE) mode with the target moved 2.5 m upstream.

relative distance between the target and Horn 1. In practice, only the latter method is used since the target can be moved without opening up the shielding around the target hall. Figure 3.7 shows the neutrino spectrum produced when the beam is in low energy (LE) mode with the target at its furthest downstream position, pseudo-medium energy (pME) mode with the target moved 1 m upstream, and pseudo-high energy (pHE) mode with the target moved 2.5 m upstream.<sup>4</sup>

The focused mesons which make up the focusing peak are not the only component of the secondary beam. Some mesons with low transverse momentum, called ‘neck-to-neck,’ pass through the necks of both horns without being focused or defocused (also shown in Figure 3.5). As can be seen on the right in Figure 3.6, these mesons have a wide range of longitudinal, and hence total momenta, and they produce the diffuse high-energy tail in the neutrino spectrum. These neck-to-neck parents are the primary source of antineutrinos in the neutrino-mode beam since most other would-be antineutrino parents ( $\pi^-$ ,  $K^-$ ) are defocused when the beam is tuned for neutrinos.

After being focused by the horns, the mesons enter the decay volume: a 675 m long, 2 m diameter steel pipe embedded in concrete shielding. Its length is approximately the decay length of a 10 GeV pion. The mesons enter this volume and decay to produce neutrinos via a number of processes, which are listed in Table 3.1, though the two-body decays of pions are by far the most important (they make up more than 80% of the beam, and are even more prevalent at the lower energies sensitive to oscillations). Since the mesons are boosted strongly towards the detectors thanks to the focusing horns, the neutrinos are boosted in that direction as well (the kinematics of meson decay

<sup>4</sup>The pseudo in pME and pHE is because true ME and HE modes require moving Horn 2 further downstream from Horn 1 in addition to moving the target upstream.

Process		Frequency
$\pi^+ \rightarrow$	$\nu_\mu + \mu^+$	78.9%
$\pi^- \rightarrow$	$\bar{\nu}_\mu + \mu^-$	5.6%
$K^+ \rightarrow$	$\nu_\mu + \mu^+$	12.0%
	$\nu_e + \pi^0 + e^+$	0.53%
	$\nu_\mu + \pi^0 + \mu^+$	0.30%
$K^- \rightarrow$	$\bar{\nu}_\mu + \mu^-$	0.65%
	$\bar{\nu}_e + \pi^0 + e^-$	0.03%
	$\bar{\nu}_\mu + \pi^0 + \mu^-$	0.02%
$K_L^0 \rightarrow$	$\nu_e + \pi^- + e^+$	0.32%
	$\bar{\nu}_e + \pi^+ + e^-$	0.14%
	$\nu_\mu + \pi^- + \mu^+$	0.18%
	$\bar{\nu}_\mu + \pi^+ + \mu^-$	0.09%
$\mu^+ \rightarrow$	$\bar{\nu}_\mu + \nu_e + e^+$	1.2%
$\mu^- \rightarrow$	$\nu_\mu + \bar{\nu}_e + e^-$	0.07%

Table 3.1: The decays that produce neutrinos and their frequency at the Near Detector.

are discussed in detail in Appendix A). Originally, the decay pipe was evacuated to approximately 1 torr, giving the mesons the best chance to decay without interacting in the decay volume. Over time, however, the aluminum window at the upstream end of the decay pipe degraded, risking an implosion at a location that is impossible to repair since it had become highly radioactive. In order to alleviate this risk, the pressure on the window was reduced by filling the decay pipe with 0.9 atm of helium.<sup>5</sup> The addition of helium has two effects on the flux. One, some mesons that would have decayed to neutrinos interact with the helium instead, causing a few percent decrease in the focusing peak. In addition, it provides an additional target for meson, and hence neutrino, production, causing a small increase in flux above the focusing peak.

The end of the decay pipe is followed by a beam dump and 240 m of rock to stop the remaining primary and secondary beams, as well as the tertiary muons, leaving only neutrinos. There are a series of ionization chambers in this region that monitor the remaining hadrons and the tertiary muons produced with the neutrinos in the meson decays. In the end, the NuMI beam is 91.7%  $\nu_\mu$ . There is an additional 7%  $\bar{\nu}_\mu$  component (which is used in the analysis described in Chapter 5), and a 1.3%  $\nu_e/\bar{\nu}_e$  component. More details about the NuMI beam and the downstream monitoring devices can be found in [118].

### 3.1.2 Making Antineutrinos

The NuMI beam can be converted from a predominantly muon neutrino beam to a predominantly muon antineutrino beam by reversing the direction of the current flowing through the focusing horns. This change flips the direction of the toroidal magnetic field, focusing negative mesons instead of positive ones. Figure 3.8 shows a cartoon of focusing in antineutrino mode.

<sup>5</sup>The helium is at a slight under-pressure to avoid reversing the direction of the force on the window.




Figure 3.8: A cartoon showing how the horn focusing enhances the flux of antineutrinos by bending negative mesons towards the beam axis. In addition to the focused mesons, there are also mesons of both charge signs that pass through the centers of the horns, avoiding focusing.

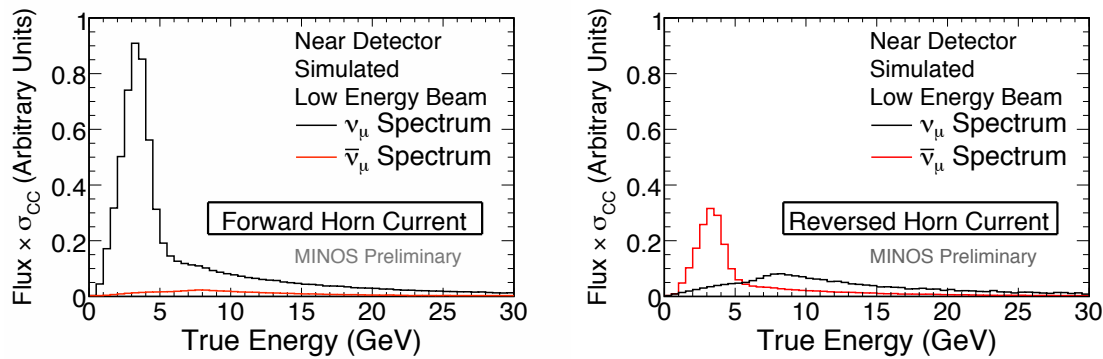


Figure 3.9: The event rate for neutrinos (black) and antineutrinos (red) at the Near Detector in neutrino mode (left) and antineutrino mode (right). While the neutrino-mode beam has high purity (91.7%  $\nu_\mu$ ), the antineutrino-mode beam is only 40%  $\bar{\nu}_\mu$  due to the lower  $\bar{\nu}_\mu$  cross sections.




Figure 3.10: At left, the ratio of  $\pi^+$  ( $\nu_\mu$  parent) to  $\pi^-$  ( $\bar{\nu}_\mu$  parent) production at the NA49 experiment at CERN as well as in several simulations [124]. At right, the  $\nu_\mu$  and  $\bar{\nu}_\mu$  interaction cross sections as a function of energy as measured in many experiments, including MINOS. In both production and interaction antineutrinos have a significantly lower cross section, leading to the differing beam compositions in Figure 3.9. Figure taken from [44].

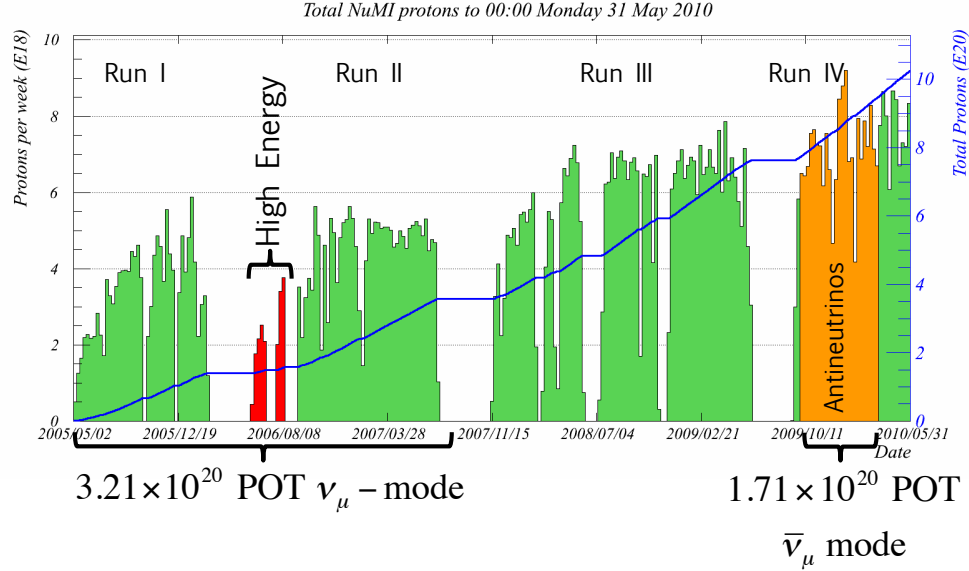


Figure 3.11: The protons-on-target accumulated per week during NuMI operations. The neutrino-mode antineutrino analysis uses the first two runs, excluding the pseudo-high energy running marked in red. The antineutrino-mode analysis uses the data marked in orange.

Figure 3.9 shows a comparison of the neutrino and antineutrino event rate at the Near Detector in neutrino mode and antineutrino mode. The antineutrino-mode beam has significantly more antineutrinos at the low energies necessary for an oscillation measurement, but its overall purity is not as high as the neutrino-mode beam. The antineutrino-mode beam is 40%  $\bar{\nu}_\mu$ , 58%  $\nu_\mu$ , and 2%  $\nu_e/\bar{\nu}_e$ . The large  $\nu_\mu$ -component in antineutrino mode, like the  $\bar{\nu}_\mu$  component in neutrino mode, comes from the neck-to-neck parents that produce the high-energy tail. In fact, the high-energy tail (above approximately 10 GeV) is unchanged between neutrino mode and antineutrino mode.

The reason the high-energy tail is predominantly  $\nu_\mu$ 's and the reason antineutrino-mode focusing peak is about a third lower than the neutrino-mode peak is the same:  $\nu_\mu$ 's have higher cross sections, both at production and when interacting. At the relevant momentum range (4 – 10 GeV parent momentum, which translates to approximately 2 – 5 GeV neutrino energy),  $\pi^+$ 's ( $\nu_\mu$  parents) are produced approximately 30% more often than  $\pi^-$ 's ( $\bar{\nu}_\mu$  parents). This effect can be seen in the ratio of  $\pi^+$  to  $\pi^-$  production from the NA49 experiment at CERN [124] as well as numerous simulations, shown in Figure 3.10 (left).  $\nu_\mu$ 's also have an interaction cross section more than a factor of two larger than  $\bar{\nu}_\mu$ 's, and the effect gets larger at lower energies, as can be seen at numerous experiments (including MINOS's own measurement [44]) in Figure 3.10 (right).

### 3.1.3 MINOS Beam Data

During the last six years of running, MINOS has collected approximately  $12 \times 10^{20}$  protons-on-target (POT) in a variety of beam modes described in Section 3.1.1. Figure 3.11 shows the POT

accumulated per week during its first five years. The data is divided up into run periods separated by accelerator shut-downs and significant beam events (e.g. switching to antineutrino mode). The majority of the data was taken in the low energy neutrino mode. Significant samples were also taken in low energy antineutrino mode ( $1.7 \times 10^{20}$  POT) and pseudo-high energy neutrino mode ( $0.15 \times 10^{20}$  POT). Smaller samples were also taken in a range of target positions and horn currents, including horn-off, for constraining the flux and other specialized applications.

The neutrino-mode antineutrino analysis described in Chapter 5 uses the  $3.2 \times 10^{20}$  POT of low energy running accumulated during Runs I and II. The antineutrino-mode analysis described in Chapter 6 uses the  $1.7 \times 10^{20}$  POT accumulated in antineutrino mode during Run IV.

The addition of helium to the decay volume described in the previous section occurred between Runs II and III. Targets were also changed between Runs I and II and between Runs III and IV. NT02, the target used for three years throughout Runs II and III lasted long enough to show a slow decrease in the number of events in the focusing peak (see Section 3.1.1). The data used in this thesis does not include significant target degradation.

### 3.1.4 Beam Simulation with Flugg

The neutrino flux from NuMI is simulated using the Flugg [125, 126, 127] package to incorporate a GEANT4 [128] geometry into a Fluka [129, 130] simulation of the hadronic production, decay, and transport processes.<sup>6</sup> The technical details are described in Appendix F. The flux simulation is primarily a simulation of the secondary meson beam. It starts with 120 GeV Main Injector protons incident on a graphite target and simulates the production of the secondary mesons and their transport through the horn focusing fields and into the decay pipe, as well as any further downstream interactions. Whenever a particle would decay to produce a neutrino, the properties of the parent are recorded to disk. Then, from that information (position and momentum), the probability of that parent's daughter neutrino reaching one of the detectors, as well as what energy that neutrino would have, can be calculated (these calculations are detailed in Appendix A). Effectively, every neutrino produced is forced to go to the detectors, but with a weight that accounts for its probability of actually doing so, saving significant processing time. This weighted neutrino flux serves as the input to the MINOS detector simulation described in Section 3.5.

The beam simulation uses importance sampling and weighting to reduce the number of low energy particles that need to be simulated. Naturally, the simulation produces many more low energy mesons than high energy ones, which can make it difficult to accumulate enough statistics at higher energies. Relative to the energy distribution of the experiment, the higher energies are under-represented in the beam simulation out of the box. A 1 GeV tracking threshold is included since those mesons will produce approximately 500 MeV neutrinos, which is the lowest energy neutrino

---

<sup>6</sup>I wrote the Flugg beam simulation and produced all the flux Monte Carlo now in use in the MINOS experiment.






Figure 3.12: The ratio of the Near Detector spectrum with helium in the decay pipe to the spectrum with an evacuated decay pipe for neutrinos (left) and antineutrinos (right). The black line represents the older GNuMI simulation, the red line represents the newer Flugg simulation, and the blue points represent the data. As shown, the Flugg-based simulation is significantly better at reproducing the effects of helium seen in the data.

the MINOS detectors are sensitive to, but even then there are still too many low energy particles. In order to speed up the simulation and reduce the space required on disk, pions below 30 GeV are importance sampled and weighted: a fraction of the events are thrown away, but the remaining events are given a weight larger than one so that the total weighted flux remains unchanged. This technique distributes the processing time and statistics more evenly across all energies even though many fewer high energy particles are produced. The weight  $W$  is calculated as,

$$W = W_{\text{parent}} \frac{30 \text{ GeV}}{|P_{\text{total}}|} \quad (3.2)$$

where  $W_{\text{parent}}$  is the importance weight of the particle's parent (initial protons start with weight 1) and  $|P_{\text{total}}|$  is the total momentum of the particle. Weights can never be below 1 and are capped at 100 to prevent a single event from being too important.

The previous beam simulation, called GNuMI, was also part-GEANT and part-Fluka. It simulated the interaction of the primary protons with the target entirely in Fluka and then simulated the transport of the particles leaving the target in GEANT3. GEANT3 uses GFluka, an older less reliable version of the Fluka hadronic interaction model, to model the interactions outside the target. This two-part simulation worked well as long as interactions outside the the target were not important, but that changed when the decay pipe was filled with helium. GFluka significantly overestimated the amount of high energy neutrino production from interactions with the helium. This mis-modeling was particularly troublesome because parents produced in the decay pipe tend to decay very close to the Near Detector, meaning they have large Near-to-Far spectral differences (see Section 5.3). The Flugg simulation, which uses the up-to-date Fluka interaction models both in the target and throughout the beamline, does a much better job of predicting the amount of production




Figure 3.13: The Near Detector  $\nu_\mu$  spectrum, in data (black), raw simulation (blue), and flux-tuned simulation (red), in three different beam configurations. The flux tuning significantly improves the data-simulation agreement in all beam configurations.

from helium interactions, as can be seen in Figure 3.12.

Even with the best available hadronic interaction models, there is still significant uncertainty in simulating the flux, which is reflected in data-simulation discrepancies at the Near Detector. While data-simulation agreement at the Near Detector is not required for the oscillation measurement, the data is much easier to analyze if the flux simulation can be tuned to better describe the observed spectrum. The flux is warped as a function of the properties of the neutrino parents as they leave the target. The warping is determined using a simultaneous fit to multiple beam configurations including standard low energy, horn off (i.e. no focusing), pseudo-medium energy, pseudo-high energy, and runs with varying horn currents at those target positions. Running in these different modes allows the Near Detector data to sample different regions of the  $p_T - p_Z$  space, allowing for a better tuning to the underlying hadron production. Both neutrino and antineutrino samples are used, but the antineutrino sample, coming from unfocused neck-to-neck parents, changes little between configurations. So, to better constrain antineutrino production, the pion charge ratio is constrained to the measurement at NA49 [124]. Data-simulation discrepancies as large as 30% in the high-energy tail are brought into good agreement, as can be seen in the focused neutrino spectra in Figure 3.13 and in the unfocused antineutrino spectrum in Figure 3.14. Even so, the flux tuning has been shown to have little impact on the final oscillation analysis since the uncertainties in the flux cancel in the Near-to-Far extrapolation.




Figure 3.14: The Near Detector  $\bar{\nu}_\mu$  spectrum, in data (black), raw simulation (blue), and flux-tuned simulation (red). The flux tuning, constrained by the NA49 pion charge ratio, significantly improves the data-simulation.

### 3.2 The MINOS Detectors

The MINOS detectors are tracking, sampling calorimeters made up of alternating planes of 1-cm thick plastic scintillator, segmented into 4 cm-wide strips, and 1-in thick steel passive absorbers. The direction of the strips alternates plane-to-plane to allow three-dimensional reconstruction, as shown in Figure 3.15. The detectors are toroidally magnetized (and hence each have a hole cut through them for the coil), enabling track charge-sign identification and momentum measurement for tracks that exit the detector. The two detectors are built with neutrino interaction environments as identical as possible so that uncertainties can cancel out in the analysis.

The Far Detector, which measures the oscillated neutrino spectrum, is made up of 486 octagonal steel-scintillator planes which are 8 m edge-to-edge with a total mass of 5400 metric tons and is divided longitudinally into two “super modules.” It is located in the Soudan iron mine, Minnesota’s oldest, which been inactive since 1962 and now operates as a State park. The detector is 705 m underground (2070 meters-water-equivalent) and 735 km from the neutrino source at Fermilab. Only a few beam neutrinos interact in the Far Detector each day and the cosmic event rate is approximately 0.5 Hz at that depth, so the signal rate at the Far Detector is dominated by spontaneous WLS fiber emission and PMT dark noise which have a combined rate of approximately 4 kHz per PMT. The Far Detector also has a series of scintillator planes placed horizontally over the top of the detector to identify cosmic ray muons whose vertical tracks might otherwise travel far through the gaps between the scintillator before depositing any visible energy, thus mimicking a contained atmospheric neutrino event.




Figure 3.15: The scintillator planes are each attached to a steel plane, arranged so that alternating planes have strips perpendicular to one another to allow 3D event reconstruction. The strips are arranged at  $\pm 45^\circ$  angles, labeled the  $U$  and  $V$  axes. Figure from [131].




Figure 3.16: Pictures of the MINOS Far (left) and Near (right) detectors.

The design of the Near Detector is similar to that of the Far but is smaller and has different electronics to account of its significantly higher event rate. During a single  $10 \mu\text{s}$  beam spill there can be 10 interactions from beam neutrinos alone<sup>7</sup> – a rate four orders of magnitude larger than at the Far Detector. Consequently, the detector does not need to be nearly as large: the Near Detector is made up of 282 planes with a total weight of 980 metric tons. It is 1 km from the NuMI target and is 100 m below ground (225 meters-water-equivalent).

The calibration detector, “CalDet,” was a much smaller, unmagnetized detector. It consisted of 60  $1 \text{ m} \times 1 \text{ m}$  planes and had a total mass of 12 tons. It was exposed to controlled test beams of protons, pions, electrons, and muons at CERN. Both charge signs and momenta ranging from 0.2 GeV to 10 GeV were studied. CalDet was able to provide measurements of the expected topological and calorimetric responses of the Near and Far Detectors for various particles of known momenta.

<sup>7</sup>3.5 reconstructed events per  $10^{13}$  POT.

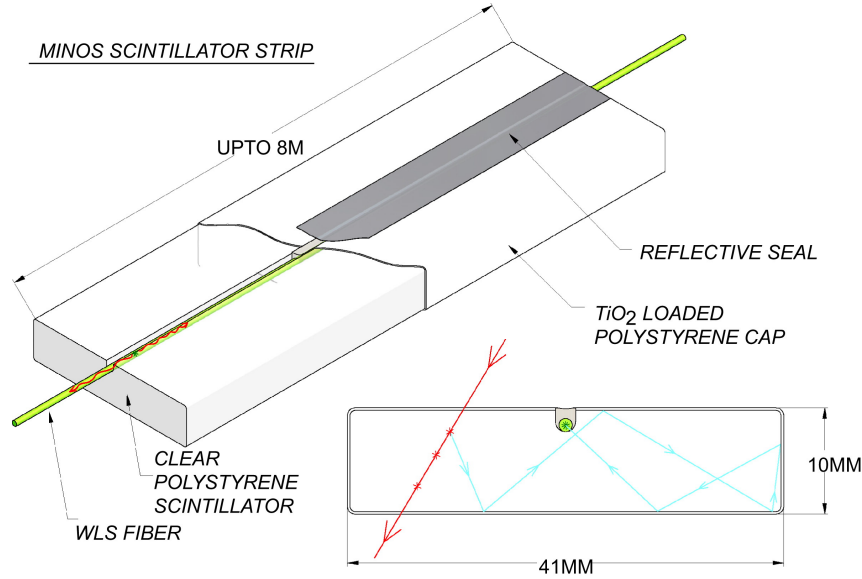


Figure 3.17: Cut-away view of a single MINOS scintillator strip. An ionizing particle passing through the strip produces photons which are multiply reflected by the  $\text{TiO}_2$ -loaded coating until they are picked up by the WLS fiber. The light is then wavelength-shifted and transported to the photodetectors at the edge of the detector.

### 3.2.1 A MINOS Plane

The lowest-level unit of the MINOS scintillator system is the strip. Each scintillator strip is 1 cm thick, 4 cm wide and between 2.5 m (the shortest Near Detector strips) and 8 m (the longest Far Detector strips) in length. Figure 3.17 shows a MINOS scintillator strip. The bulk of the material is clear polystyrene with 1.03% by weight scintillating fluors mixed in at extrusion time. Each strip is co-extruded with a 0.25 mm reflective coating made from polystyrene with 15%  $\text{TiO}_2$  mixed in. This coating traps and reflects the light, keeping each strip optically isolated from its neighbors.

In order to get the light out of the scintillator and to the photodetectors, there is a wavelength-shifting (WLS) fiber running the length of each strip. The strips are extruded with a 2 mm groove in the middle of one of the wide sides into which the fiber is glued and then covered with reflective tape. The WLS fiber, doped with Y11 fluor, absorbs the blue photons ( $\lambda \approx 420$  nm) produced by the scintillator and re-emits them as green photons ( $\lambda \approx 520$  nm), minimizing self-absorption. The green photons are emitted isotropically and any that fall within the fiber's total internal reflection cones are transported out.

Groups of strips, between 16 and 24, are bundled together into modules. Each module is wrapped in aluminum to make it light-tight and to protect it during transportation and installation. It is as modules that the scintillator was transported to the detector sites, and individual modules were attached to the steel plates to build up a steel-scintillator plane. The various modules, and their arrangement into planes, are shown in the Far Detector in Figure 3.18 and in the Near Detector




Figure 3.18: The arrangement of scintillator modules into a  $U$  plane in the Far Detector (the  $V$  plane is the same but  $90^\circ$  rotated). Eight fibers from different modules are routed to a single PMT pixel. The MUX boxes house the PMT's and handle the strip-to-pixel routing. The strips are read out from both ends, but with different strip-to-pixel routing so that strips can be unambiguously “demuxed.” Figure from [131].




Figure 3.19: The arrangement modules into the four types of Near Detector scintillator planes (the steel is the same for all planes). Every fifth plane throughout the detector is an alternating Full  $U$  or  $V$  plane. In the upstream calorimeter region containing the fiducial volume, the intervening planes are alternating Partial  $U$  and  $V$  planes, creating a repeating 10-plane pattern of FU-PV-PU-PV-PU-FV-PU-PV-PU-PV. The beam is centered half-way between the coil hole and the West edge of the detector, so only this region is instrumented in the partial planes. In the spectrometer the intervening planes are left uninstrumented. The Near Detector strips are readout only from the left side. Figure from [131].

in Figure 3.19. The modules were limited in size at the Far Detector by the need to bring them underground through the existing mine elevators.

Each Far Detector plane is identical except for the alternating orientation of the scintillator strips. The geometry of the Near Detector is somewhat more complicated as a response to the significantly higher event rate. The Far Detector fiducial volume fills as much of the detector as possible while still containing all the energy of each event. The first 120 of the Near Detector's 282 planes are called the "calorimeter." Each of the calorimeter's planes is instrumented with scintillator, though only every fifth plane is completely covered with 96 strips (a "Full Plane"). Most of the planes are "Partial Planes" where only the West side of the detector is covered with 65 strips. This cost-saving measure was possible because the Near Detector fiducial volume does not need to be large in order to accumulate enough events – it is a 2 m cylinder centered between the coil hole and the western edge of the detector and is wholly contained within the first 120 planes. The Near Detector's squat shape and asymmetric coil hole placement allow the fiducial volume to be uninterrupted. The remaining 162 planes are called the "spectrometer" and serve only to better measure the momenta of muons by giving them more time to range out, or at least a larger sample of points from which to measure their curvatures.

The steel plates on which the scintillator is mounted make up the majority of the detectors – approximately 95% by weight. The steel provides a relatively inexpensive target material for the neutrino interactions as well as an effective absorber for muon spectroscopy. The plates are made from low-carbon, hot-rolled steel, required to be flat to within 15 mm, and were tested to be low in radioactivity. The steel density was measured to be  $7.85 \pm 0.03$  g/cm<sup>3</sup>, based on samples from various steel batches (heats). In order to be transported down the elevator, the Far Detector planes were assembled from eight individual pieces plug-welded together. The smaller Near Detector plates were fabricated as whole pieces. The quantity of steel needs to be known precisely both for muon range measurements and for calculating the relative neutrino event rate which depends directly on the relative target masses. Each of the component pieces of the Far Detector planes was measured individually, and over the whole detector the plane masses vary with an RMS of 0.62%. The Near Detector planes were not weighed, but ultrasound measurements put the thickness variations at no more 0.3% (in the Far Detector, where both the mass and the thickness were sometimes measured, they were found to correlate strongly with each other).

### 3.2.2 The Magnetic Field

The choice of steel as the bulk target material had an additional advantage beyond its cost – it allowed the detectors to be magnetized. Each detector has a coil that passes through a hole cut through the length of each detector and then around bottom of the Far Detector or the lower East-side corner of the Near Detector. Each of the Far Detector super modules is magnetized separately

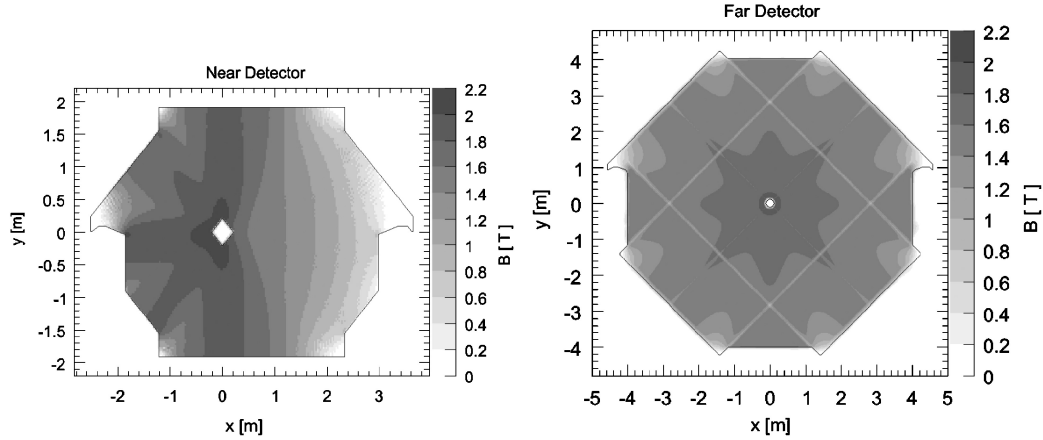


Figure 3.20: Maps of the field strengths in the Near (left) and Far (right) Detectors. The magnetic field, averaged over the fiducial volumes, is 1.42 T in the Far Detector and 1.28 T in the Near Detector. Figure from [131].

using water-cooled, 190-turn coils made from Teflon-insulated stranded copper wire carrying a total current of 15.2 kA-turn producing an average magnetic field strength of 1.42 T in the fiducial volume. The Near Detector, because of its squashed geometry and the off-center placement of the coil hole, requires a current of 40 kA-turn in order to produce a magnetic field comparable to Far Detector field. The Near Detector coil has 8 turns, each of which is made up of six cold-conformed aluminum conductors carrying current in parallel. There are no photodetectors on the coil return side of the Near Detector in order to avoid interference from the fringe fields produced by such high currents. Figure 3.20 shows field maps of the Near and Far Detectors determined using finite-element analysis.

The curvature of a muon track in the toroidally magnetized detector is proportional to the ratio of the track's charge to its momentum,  $q/p$ . The field thus makes it possible to measure the momenta of muons even if they leave the detector before ranging out. It also makes it possible to measure the charge of individual muons. Since neutrinos and antineutrinos produce opposite-sign muons in charged-current interactions, neutrinos and antineutrinos can be separated event-by-event, a unique ability among neutrino oscillation experiments. The magnetic fields in the detectors, like the magnetic fields in the horns, can be reversed to bend either muon sign inwards or outwards.<sup>8</sup> Negative muons from  $\nu_\mu$  interactions are bent inwards while the detector is in neutrino mode, and positive muons from  $\bar{\nu}_\mu$  interactions are bent inwards while the detector is antineutrino mode. The mode of the detectors typically matches the mode of the beam.<sup>9</sup> Without the ability to do charge-sign identification, the measurements made with antineutrinos described in later chapters would not be possible.

<sup>8</sup>Muons focused inwards typically have better-measured momenta since they are more likely to range out in the detector (range-based momentum measurements are more accurate than curvature-based one).

<sup>9</sup>The Far Detector field has always matched the beam field during beam data taking. Sizable samples of Near Detector data were taken with a reversed field, i.e. antineutrino-mode detector with a neutrino-mode beam and vice versa. These samples were used for cross-checks and systematic studies.



### 3.2.3 Detector Readout and Data Acquisition

The end of each module has a light-tight manifold where the WLS fibers from the strips are routed together into an optical connector. From this connector, the light is transported through clear fiber to the photodetectors. In the Near Detector, the photodetector is a 64-anode Hamamatsu (M64) photo-multiplier tube (PMT); in total, 194 PMTs are used. Each PMT pixel receives the light from one scintillator strip. In the upstream calorimeter region each strip is read out individually by the front-end electronics. Downstream, in the spectrometer, the number of front-end channels was reduced (to save cost) by connecting four PMT anodes in parallel, electrically summing their signals. The four strips summed together are separated by about 1 m so that muon tracking can unambiguously assign the hit to the correct strip, which is sufficient since the spectrometer is only used to measure muon tracks.

The Far Detector uses 1452<sup>10</sup> 16-anode Hamamatsu (M16) PMTs. However, each of the 16 pixels receives light from eight scintillator strips. The fiber routing to get the correct optical summing (multiplexing or MUX) on each PMT pixel is handled by 484 MUX boxes, each housing three PMTs. The hits can be unambiguously “demuxed” since each strip is read out from both ends and the two ends have different strip-to-pixel patterns designed to minimize cross-talk.

The front-end electronics integrate, linearize, and digitize the PMT signals before passing them on to the DAQ. At the Far Detector [132], each PMT is read out by a custom-designed, multi-channel integrated circuit called a Viking VA chip. The VA chips are mounted on VA front-end boards (VFB’s) which provide power, biasing controls, and a programmable ASDLite discriminator chip which uses the PMT dynode signal for triggering digitization, usually for light levels above 0.25 PE. VA readout controllers (VARC’s) each handle the digitization, sparsification, pedestal subtraction, and readout<sup>11</sup> for 36 PMTs via 6 VA Mezzanine Modules (VMMs) which each handle 6 PMT’s. Since digitizing a VA chip introduces approximately 5  $\mu$ s of dead time for that chip, a coincidence of two or more dynode triggers within 400 ns from ASDLite chips in the same VARC is required to initiate digitization.

The Near Detector has specialized, fast front-end electronics because of its much higher event rate [133]. Each PMT pixel is attached to a MINOS Electronics for Neutrinos (MENU) circuit board containing a Charge Integrator and Encoder (QIE) chip, a commercial flash analog-to-digital converter (FADC), and a data buffer. The QIE is a custom integrated circuit developed at Fermilab<sup>12</sup> which rapidly integrates charge over a wide dynamic range. It contains four independent copies of the integrating circuits to allow dead-timeless operation when triggered externally. Even when accumulating cosmic ray data, for which there is no external trigger, each PMT has a dead time

<sup>10</sup>There are an additional 64 PMTs for the cosmic Veto Shield.

<sup>11</sup>Since speed is not critical at the Far Detector these process are typically handled by commercially-available components.

<sup>12</sup>QIE chips are also used by the CDF and KTEV experiments.




Figure 3.21: The WLS fibers from multiple strips are grouped together and routed to a multi-pixel photo-multiplier tube. In the Far Detector eight strips are routed to each of 16 PMT pixels. In the Near Detector each of the 64 PMT pixels receives the light from only a single strip. Figures from [131].

of only 0.5%. MINOS Near Detector Electronics Readout (MINDER) boards each readout and timestamp 16 MENU's before passing the digitized data on to data acquisition.

The data acquisition system (DAQ) reads out the front-end electronics continuously at both detectors, but only writes the data to disk if certain trigger conditions are met. The Near Detector electronics are triggered directly by the beam spill signal at Fermilab and all the data during a 100  $\mu$ s window around the beam spill is recorded. The Far Detector cannot be triggered directly since it is not at Fermilab but instead receives the GPS timestamp for the beam spill recorded at the Near Detector via the internet. The Far Detector DAQ has several seconds of buffering capacity so the ‘‘remote spill trigger’’ has time to arrive. Both detectors also have signal-based triggers (for example, the ‘‘plane trigger’’ requires dynode triggers on four out of five consecutive planes) for cosmic ray and atmospheric neutrino interactions that occur outside the beam spill as well as other, more specialized triggers for background monitoring, debugging, and calibration.

More details about the detector design and construction, the readout electronics, and the calibration described in the next section can be found in [131].

### 3.3 Calibration

MINOS uses a multi-step calibration chain to convert the raw detector pulse height to a fully-corrected signal which is consistent across time, position in the detector, and which detector the signal was in. The system uses charge injection and light injection, which measure the photodetectors and readout systems, cosmic ray muons, which measure the scintillator response, and bench-top measurements of the WLS fiber. The full chain can be broken down as follows:

$$Q_{\text{corr}} = Q_{\text{raw}}(s, x, t, d) \times D_{\text{PMT}}(d, t) \times L(d, s, Q_{\text{raw}}) \times D_{\text{scint}}(d, t) \times S(d, s, t) \times A(d, s, x) \times M(d) \quad (3.3)$$

where  $s$  is the strip that was hit,  $x$  is the position along that strip,  $t$  is the time of the hit, and  $d$  is the detector where the hit occurred (Near, Far, CalDet).  $Q_{\text{raw}}$  is the raw pulse height from the strip,  $Q_{\text{corr}}$  is the fully corrected signal that is passed on to reconstruction, and the remaining terms are explained below. The corrected pulse heights are then used to reconstruct physics objects (tracks and showers) in which the pulse heights become physical energies.

#### **PMT and Electronics Drift** $D_{\text{PMT}}(d, t)$

The ‘drift’ is the change in PMT and electronics gain over time. Short-term variations are generally related to changes in local temperature and can be eliminated with good environmental control. There are also long-term changes that cannot be removed so easily, related to seasonal environmental changes and component aging,

amounting to no more than a 4% annual change in the overall detector gain.

The PMT and electronics gain is calibrated using a Light Injection (LI) system that pulses UV LEDs to illuminate the each WLS fiber 300 times per hour in the Far Detector and 1000 times per hour in the Near Detector. The UV LEDs simulate the light from the scintillator, but in a controlled and reproducible fashion. This LI data is collated at each detector on a monthly basis to produce an average response per photoelectron per channel, which enters the signal correction as  $D_{\text{PMT}}(d, t)$ .

#### **Linearity** $L(d, s, Q_{\text{raw}})$

The goal of the linearity calibration is to correct the output of the electronics so that it is linearly related to the true illumination. In the Near Detector, each integrator/digitizer is calibrated by injecting known amounts of charge and observing the output signals. This electronics-only linearity correction is applied directly in the front-end electronics.

The Far Detector electronics, as well as the PMTs at both detectors, also show 5% – 10% non-linearity starting above about 100 photoelectrons. This non-linearity is also corrected using the LI system. Monthly, an extend light injection run is performed where each strip end is pulsed 1000 times at a wide range of light levels, allowing the PMT/electronics response as function of true illumination to be measured. This gain is then linearized offline.

#### **Scintillator Drift** $D_{\text{scint}}(d, t)$

Like the electronics and photodetectors, the scintillator response changes as a function of temperature and degrades as it ages. Aging reduces the light level by about 2% annually and larger short-term changes have been shown to correlate well with temperature variations. The calibration of the scintillator drift is performed using the large sample of through-going muons accumulated at each detector. The total pulse-height per plane deposited by these muons is calculated each day. Though the total amount of energy deposited by these muons varies with zenith angle and is different between the detectors, the amount of energy deposited at any one detector is expected to be constant over time. The correction, then, is calculated by comparing the median response each day to the median response at  $t_0$ , the beginning of the experiment:

$$D_{\text{scint}}(d, t) = \frac{\text{Median response}(d, t_0)}{\text{Median response}(d, t)}. \quad (3.4)$$

#### **Strip-to-Strip Uniformity** $S(d, s, t)$

The detector response also varies strip-by-strip versus time. Individual strip ends

can vary from the detector average by as much as 30%. These effects are also calibrated out using cosmic ray muon tracks. The mean light level at each strip end due to cosmic ray muons is measured after linearization. Hits are corrected for known effects from attenuation (as the light travels through the WLS fiber) and the path-length of the muon through the scintillator so that each strip is calibrated for an idealized muon travelling at normal incidence through the center of the strip. The strip-to-strip calibration is then calculated relative to the mean detector response

$$S(d, s, t) = \frac{\text{Mean detector response}(d, t)}{\text{Mean strip end response}(d, s, t)}. \quad (3.5)$$

#### **Attenuation** $A(d, s, x)$

Hits that occur close to the edges of the detector will have a higher light level than hits deep in the interior due to attenuation along the WLS fiber. In principle, cosmic ray muons could be used for the attenuation correction, as well, but it is more accurate to use the measurements taken at various position along each scintillator strip at the production factory using a known  $\gamma$  source. These ‘module mapper’ data were then fit to a double exponential in each strip

$$A(d, s, x) = A_1^{d,s} e^{-x/L_1^{d,s}} + A_2^{d,s} e^{-x/L_2^{d,s}} \quad (3.6)$$

where  $x$  is the length along the strip and  $L_1^{d,s}, L_2^{d,s}$  stand for the two attenuation lengths. The attenuation correction was subsequently checked with through-going muons and found to be consistent to 4%.

#### **Inter-detector** $M(d)$

All the previous calibrations have served to produce temporally and spatially consistent data at each of the three detectors. In order to normalize the energy scales of the Near and Far Detectors and CalDet to each other, the stopping cosmic ray muon sample was used. Only a small fraction of all cosmic ray muons stop in the detectors (0.3% in the Far Detector), but there are enough cosmic rays that this is a sufficient sample. Stopping muons are used since they produce well-understood energy depositions, especially between 0.5 GeV and 1.1 GeV where  $dE/dx$  is basically constant. The method of using only the portion of the muon track when it was in that energy range is called the ‘‘track window technique’’ and reduces a 2% uncertainty in the track’s end position to a 0.2% uncertainty in the energy deposition. Using the stopping muon sample the inter-detector calorimetric responses were calibrated to a consistency of better than 2%.

### Absolute Track and Shower Energy

Once the detectors are all calibrated to give the same hit-to-hit response, the overall response to physics objects needs to be evaluated. These absolute track and shower energy scales were determined by measurements with CalDet while it was exposed to various test beams at CERN. Data was taken with electrons, pions, muons, and protons of both charge signs with momenta ranging from 200 MeV to 10 GeV, with other detectors in the beamline to identify the particles. Thus, the particle species<sup>13</sup>, charge, and momenta were all known a priori before the event was observed in the detector. The data were then compared to the GEANT3 detector simulation.

The predicted range of stopping muons was modeled to better than 3%, thus benchmarking the accuracy of the GEANT3 simulation and the GEANT3 swimmer used to reconstruct track energies. The detector response to electrons agreed to better than 2%. The response to pion and proton showers guided the choice of hadronic interaction model for the detector (GCALOR [134]) and showed that the Monte Carlo could reproduce the measured hadronic showers to better than 6% at all energies. The reconstruction of track and shower energy is discussed in greater detail in Section 3.4.

## 3.4 Reconstruction

The MINOS reconstruction is a C++ based framework that takes a snapshot of digitized (or simulated) detector readouts and determines the visible energy and other properties of the products of the underlying neutrino interaction.

The input snapshot is called a ‘snarl’ and, in the Near Detector, typically contains several neutrino interactions. Each snarl corresponds to a beam-spill window or other trigger. The first step is to separate the individual neutrino interactions both spatially and temporally, allowing overlapping or simultaneous events to be distinguished. Figure 3.22 shows an example Near Detector snarl with several events distributed in space and time. In principle, slicing is applied to the Far Detector data, as well, but multiple neutrino interactions during a single beam spill are rare.

### 3.4.1 Event topologies

MINOS beam events fall into four categories: charged current muon neutrinos ( $CC\nu_\mu$ ), charged current muon antineutrinos ( $CC\bar{\nu}_\mu$ ), charged current electron (anti)neutrinos ( $CC\nu_e$ ), and neutral currents (NC). There are also  $\nu_\tau$  and  $\bar{\nu}_\tau$  events, but very few and they are indistinguishable from

---

<sup>13</sup>The upstream detectors could not distinguish pions and muons, but the particles produce very topologically distinct events in the detector so a posteriori separation was not a problem.




Figure 3.22: An example of a Near Detector snarl with several neutrino event distributed in space (left) and time (right). Figure taken from [135].

neutral currents in the MINOS detectors. As described in Section 2.2.4,  $CC\nu_\mu$  and  $CC\bar{\nu}_\mu$  events are characterized by long, curving muon tracks ( $\mu^+$  for  $\bar{\nu}_\mu$  and  $\mu^-$  for  $\nu_\mu$ ) with a hadronic shower at the interaction vertex.  $CC\nu_\mu$  and  $CC\bar{\nu}_\mu$  events can be distinguished by the direction the track curves in the detector magnetic field. While running with neutrino-mode beam, both detector fields are typically tuned to bend  $\mu^-$  inwards and  $\mu^+$  outwards, and vice versa during antineutrino-mode beam. The detector fields are chosen this way because tracks bent inwards are less likely to exit the detector, allowing the more accurate range momentum measurement to be used.

$CC\nu_e$  and NC interactions produce events without muon tracks, though they sometimes still have short reconstructed tracks.  $CC\nu_e$  events are characterized by compact electromagnetic showers and NC events are characterized by more diffuse hadronic showers, but distinguishing between them is not necessary for either antineutrino analysis. Specific analyses have been performed for both the  $CC\nu_e$  sample [136] and the NC sample [137], but in the analyses presented in this thesis the events without muon tracks are background. Given the predominance of NC events, the  $CC\nu_e$  events are wrapped into the NC sample. Figure 3.23 shows event displays from a  $CC\nu_\mu$ , a  $CC\bar{\nu}_\mu$ , and an NC interaction.

### 3.4.2 Tracks

Once the snarl has been split up into events, a track finder searches for small track-like segments – several hits in an approximate line across several planes. The track finder then joins these segments together to produce a ‘seed track.’ The seed track is then fit using a multi-pass Kalman filter [138]. The filter moves forwards and backwards along the track, attempting to estimate the state (e.g. momentum) of the underlying muon at each point along the track, including the effects of noise and multiple scattering in addition to the expected curvature in the magnetic field. It makes the final decision on whether or not a particular hit is part of the track.

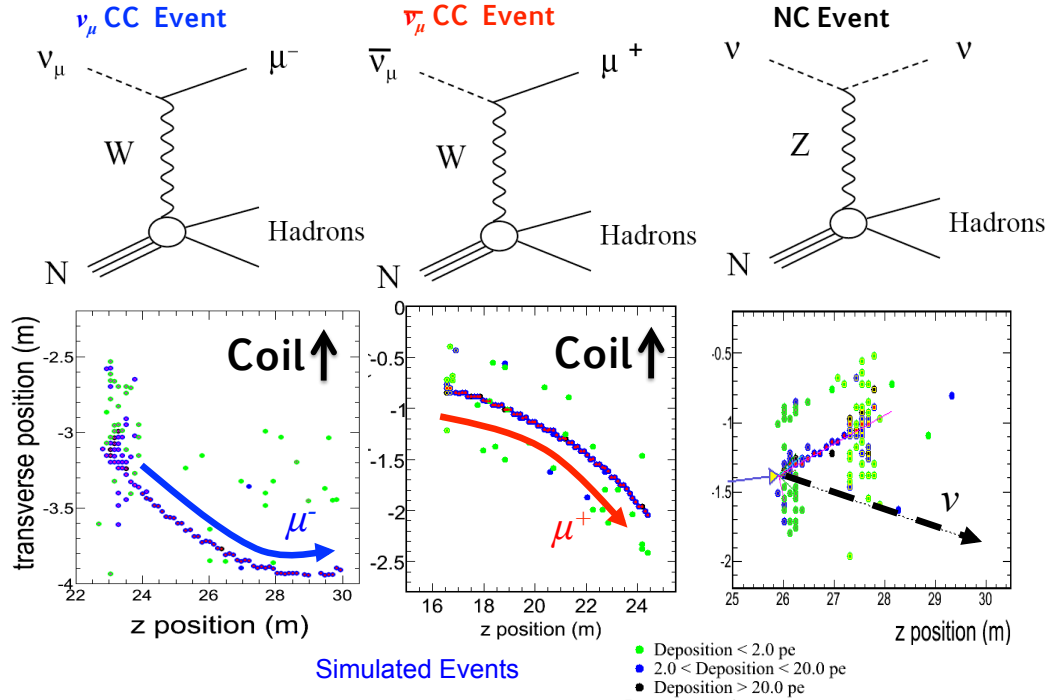


Figure 3.23: The three event topologies relevant to the antineutrino analysis:  $CC\nu_\mu$  (left),  $CC\bar{\nu}_\mu$  (center), and NC (right). The top row shows the Feynman diagram and the bottom row shows a representative simulated event in one view (i.e. only U planes). CC events are characterized by long muon tracks which curve in opposite directions for  $CC\nu_\mu$  and  $CC\bar{\nu}_\mu$ . NC events do not have true muon tracks, but can have fake tracks which make them a background at low energy. The green points are hits with light levels below two photo-electrons and are not included in the analysis.

After two passes with the filter, the fitted momentum state (or more precisely, charge-to-momentum ratio) of the particle's first hit is curvature-based estimate of the track's initial momentum when created in the detector. Other properties of the track fitting, for example how well the track's curvature (momentum) was measured, are also recorded. The curvature of the track is proportional to the ratio of its charge to its momentum,  $q/p$ . At 3 GeV the resolution of this measurement is 11%.

If the track ends in the detector, a second, more accurate measurement of the momentum is made using the range of the track through the plastic and steel. At 3 GeV, the momentum measurement by range has a resolution of 4.6%. It is used for all tracks that do not exit the detector or end in the uninstrumented coil hole. Muons with energies between 10 MeV and 10 GeV lose energy almost exclusively through ionization (see Figure 27.1 on page 286 in [35]) and are well described by the Bethe-Bloch equation [35, 139]. This is precisely the energy regime relevant to oscillations in MINOS. In practice, the energy is measured by swimming a muon backwards along the track using the GEANT3 simulation package [140] and progressively adding back in the energy that would have been lost in each steel and scintillator plane.

CalDet confirmed the accuracy of the Bethe-Bloch equation as tabulated by Groom [139] with



material-specific density effects tabulated by Sternheimer [141, 142]. Modifying the GEANT3 simulation to use the Groom tabulation produced data-MC agreement at better than 2% [131, 143]. The remaining 2% is taken as a systematic uncertainty in the analysis. Since 95% of the energy loss occurs in the steel planes, each of which is nominally 1.46 radiation lengths thick, the amount of steel the muon passes through must also be known precisely. The density of the steel was measured to an accuracy of 0.3% and the Near and Far Detector plane thicknesses were measured to 0.1%<sup>14</sup> and 0.2%<sup>15</sup> respectively.

The curvature-based measurement of the track momentum was calibrated by comparison with the range-based measurement. The two energy measurements were compared for stopping tracks, and it was found that the curvature-based measurement generally agreed with the range-based measurement to within 1%. This 1% is conservatively added linearly to the 2% uncertainty from the range measurement, leading to a total uncertainty of 3%.

### 3.4.3 Showers

Once the tracks have been identified, the remaining hits in proximity to one another are grouped together into showers. Hits that are part of a track, but with more energy than the muon would have deposited, have the track portion of the energy subtracted before being included in a shower. Unlike the muon, whose energy is measured topologically, the shower energy is measured calorimetrically. The MINOS detectors are too coarse (each ‘pixel’ is 4 cm across and separated longitudinally by 5 cm of steel and air) to reliably distinguish the component particles in the shower. Instead, the energy of the shower is reconstructed calorimetrically: it is estimated based on the total energy deposited by all of its constituent hits.

The absolute shower energy scale is also calibrated using the test beams at the calibration detector. This detector provided the opportunity to measure the detector response to hadronic particles of known energy. The electron shower data agreed with the GEANT3 simulation to less than 2%. The hadronic shower measurements showed data-MC agreement of 6% and helped guide the choice of GCALEOR [134] as the hadronic interaction model in the simulation.

The hadronic and electromagnetic shower energy resolutions can be adequately modeled by the simulation, and they are parameterized as  $56\%/\sqrt{E} \oplus 2\%$  for hadrons and  $21.4\%/\sqrt{E} \oplus 4\%$  for electrons, where  $E$  is the particle’s energy in GeV.

---

<sup>14</sup>Measured by ultrasound

<sup>15</sup>Measured by weight

### 3.5 Simulation

The two-detector design reduces, but does not eliminate, the need for simulation. The MINOS data is simulated in a multi-stage process that begins with the Flugg beam simulation described in Section 3.1.4. Neutrinos are sampled from the flux simulation, using rejection-sampling to account for the importance weighting. They are traced through models of the Near and Far detector halls, allowing the neutrinos to interact both in the detector and in the surrounding material.

The neutrino interactions are generated by the NEUGEN [144] program. NEUGEN simulates quasi-elastic and inelastic neutrino interactions at a wide range of energies, from 100 MeV to 100 GeV. It simulates the hadronization process (described in Section 2.2.4) with the AGKY model [145], which uses a combination of PYTHIA/JETSET [146] at high hadronic invariant mass and the KNO phenomenological model [147] at low invariant mass, with a smooth transition between the two models. NEUGEN also includes the INTRANUKE [148] model of intranuclear rescattering to account for the interactions of the hadronic particles as they leave the nucleus.

Once the products of the interaction leave the nucleus, their simulation is taken over by GMINOS which includes a detailed geometric model of the detector written in GEANT3 [140]. It also includes a detailed model of the magnetic field created using finite element analysis and measured B-H curves from steel samples. GMINOS transports the particles through the detector geometry, recording their strip-by-strip energy depositions as the particles lose energy into the steel and scintillator. At this stage, multiple neutrino interactions, both from in the detector and in the surrounding material, are overlaid into a single simulated snapshot in time to reflect the high Near Detector event rate.

Once the energies have been deposited on scintillator strips, the C++ based PhotonTransport program takes over. PhotonTransport generates photons in the scintillator based on the GMINOS energy depositions, transports those photons into the WLS fiber and onto the PMT cathode where they are converted into photoelectrons. It includes the detailed behavior of the PMTs and electronics, including non-linearity, noise, cross-talk, and triggering. In order to model the real-world detector as well as possible, the PhotonTransport simulation is ‘decalibrated’ by applying the measured calibration constants in reverse. Thus the simulation includes the best knowledge available on light-levels, attenuation, non-linearity and gains. Each simulated run is given a fictitious date from some time during actual data taking and calibration constants from that time are used. Later, when calibrations are re-applied, each Monte Carlo run is re-calibrated using the same date that was used to produce it.

At this stage, the simulation is as similar to the real data as possible and both are handled in the same way during reconstruction.

## Chapter 4

# An Oscillation Analysis in Brief

All oscillation analyses in MINOS follow the same basic structure. Everything begins with a signal in the detector: light in the scintillator strips. The light is amplified by photomultiplier tubes, digitized by the front-end electronics, and becomes a single ‘hit.’ The reconstruction algorithm then groups many hits together in space and time to produce ‘tracks’ (typically muons) and ‘showers.’ Tracks and showers that share a vertex are then grouped together into a neutrino ‘event.’

There are a number of classes of beam-induced events that occur in the detectors: charged-current (CC)  $\nu_\mu$  and  $\bar{\nu}_\mu$ , neutral current (NC), CC  $\nu_e$  and  $\bar{\nu}_e$  (both inherent in the beam and a possible appearance signal), and CC  $\nu_\tau$  and  $\bar{\nu}_\tau$  (only at the Far Detector).<sup>1</sup> Each oscillation analysis is focused on looking for changes between the event samples in the the Near and Far Detectors. The first step in any analysis, then, is to select as pure a sample as possible of the event class of interest. Here, the signal sample is CC  $\bar{\nu}_\mu$ ’s, whose main backgrounds are CC  $\nu_\mu$ ’s and NC’s.

Once a sample has been selected (many more details on that process are given in Sections 5.1 and 6.1), the event energies are required. The energy of a CC neutrino event is estimated by summing the shower energy, measured via calorimetry, and the track energy, measured via the range the muon travels or the amount it curves in the detector (more details are given in Section 3.4). These energy measurements are collected together to form Near and Far Detector energy spectra. While closely related, these two spectra are not identical. In addition to having lower statistics because of being further away from the neutrino source, the Far Detector also has spectral differences related to a combination of the geometry of the beamline and the kinematics of the meson decays that produce the neutrinos (see Section 5.3).

The Monte Carlo simulation is used to account for these Near-to-Far differences, converting the measured Near Detector spectrum to a ‘prediction’ of what the Far Detector spectrum would look like with any arbitrary choice of oscillation parameters. This process is called extrapolation.

---

<sup>1</sup>All these event categories are also produced by neutrinos produced by cosmic rays in the atmosphere, but these can be effectively eliminated by selecting only events in-time with the beam ‘spill triggers,’ which account for less than 0.01% of the Far Detector live time. The average rate of atmospheric neutrino interactions is less than half that of beam neutrinos.

Since these Far Detector predictions are based on the Near Detector spectrum, cross section and flux systematics largely cancel.<sup>2</sup> Take, for example, a systematic error in the simulation's neutrino interaction cross section; say it is high by 10%. With only one detector there would appear to be a 10% deficit of neutrinos. With two detectors, however, the simulation is only required to convert the measured Near Detector data into a Far Detector prediction. Since neutrinos interact in both detectors using the same cross section, they will both be 10% high but the *relative* number of events expected at the Far Detector for a given number of events in the Near Detector remains unchanged. The systematic has been 'cancelled out' in the extrapolation process – this is the power of the two-detector design.

Since the prediction can be produced for any arbitrary set of oscillation parameters, those parameters can be varied to find the values that best fit the data. In practice, a search is performed over the possible oscillation parameters to find the ones that maximize the likelihood of the observed data spectrum given the prediction. Then, starting from the best fit parameters, a two-dimensional confidence interval (contour) can be drawn showing how the data constrain the values of the parameters. For more details, see Sections 5.6, 5.7, and 6.5.

---

<sup>2</sup>Given the spectral differences mentioned above, the systematics do not cancel completely, but their effects can be estimated and are small. See Sections 5.4 and 6.3.

## Chapter 5

# Antineutrinos in a Neutrino Beam

The first  $7.2 \times 10^{20}$  POT of data was taken with the NuMI beamline running in neutrino mode as described in Section 3.1.1. However, the neutrino-mode beam has a small (approximately 7%) component of antineutrinos. Two measurements were made with the antineutrinos from the first  $3.2 \times 10^{20}$  POT of running: one looking for antineutrino oscillations via disappearance, and one looking for neutrino-to-antineutrino transitions via antineutrino appearance.<sup>1</sup> The disappearance analysis measures the  $|\Delta\bar{m}_{\text{atm}}^2|$  and  $\sin^2(2\bar{\theta}_{23})$  parameters from the oscillation survival probability:

$$P(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu) = 1 - \sin^2(2\bar{\theta}_{23}) \sin^2\left(\Delta\bar{m}_{\text{atm}}^2 \frac{L}{4E}\right) \quad (5.1)$$




Figure 5.1: Simulated Far Detector reconstructed energy spectrum showing the effect of transitions with  $\alpha = 0.12$  and oscillations with  $|\Delta\bar{m}_{\text{atm}}^2| = 5.65 \times 10^{-3} \text{eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 1$ . The values chosen correspond to the lowest parameter values with measurable effects at 99% CL.

<sup>1</sup>I performed the neutrino-to-antineutrino transition analysis.




Figure 5.2: The  $p_T$  vs.  $p_Z$  distribution of the  $\pi^\pm$  parents that produce neutrinos (left) and antineutrinos (right) at the Near Detector when the beam is in low-energy neutrino-mode mode. The unfocused component has a broad range of  $p_Z$ , and hence total momentum, producing the diffuse high-energy tail.

where  $L$  is the baseline over which the oscillations occur and  $E$  is the energy of the neutrino. The appearance analysis constrains the possibility that some fraction,  $\alpha$ , of the neutrinos that MINOS has observed to disappear are actually transitioning to antineutrinos:

$$P(\nu_\mu \rightarrow \bar{\nu}_\mu) = \alpha \sin^2(2\theta_{23}) \sin^2\left(\Delta m_{\text{atm}}^2 \frac{L}{4E}\right) \quad (5.2)$$

which would be visible as an anomalous low-energy peak in the antineutrino spectrum. A simulated example of each of these signals is shown in Figure 5.1.

Antineutrinos are produced in the decays of pions produced by colliding primary protons with the graphite target.



The antineutrinos in the neutrino-mode beam come primarily from low- $p_T$  pions leaving the target headed directly down the axis of the beamline and thus avoid being defocused. These “neck-to-neck” pions, so-called since they pass directly through the necks of both horns, are unaffected by the horn magnetic fields. Without the momentum-selecting benefit of focusing, the antineutrinos are left with a broader spectrum with a higher peak energy than the neutrinos (7 GeV instead of 3 GeV). A comparison of the neutrino and antineutrino parents’  $p_T - p_Z$  distributions can be seen in Figure 5.2.

## 5.1 Selection

The first step in the analysis of the antineutrinos is to select as large and pure a sample of them as possible. Of course, the antineutrino itself cannot be directly observed, so its properties must be inferred from the outgoing remnants of the antineutrino’s collision with a nucleus in the detector. Only charged current (CC) interactions are desired in the analysis since the neutrino flavor can be determined from the flavor of the outgoing lepton,

$$\bar{\nu}_\mu + N \rightarrow \mu^+ + \text{hadrons.} \quad (5.5)$$

Selecting the antineutrino component of the neutrino-mode beam would be impossible if not for MINOS’s magnetized detectors. The separation of neutrinos and antineutrinos relies on the fact that the two charged current interactions produce opposite sign outgoing leptons:  $\bar{\nu}_\mu \rightarrow \mu^+$  and  $\nu_\mu \rightarrow \mu^-$  (the samples are sometimes referred to as positive charge, PQ, and negative charge, NQ, respectively). The toroidal magnetic fields (approximately 1.4 T) in the two detectors cause the outgoing muons to curve in different directions depending on their charge. When the detectors are in neutrino mode,<sup>2</sup> negative muons curve inwards and positive muons curve outwards. The reconstruction algorithm measures the track’s curvature, which is proportional to the ratio of the track’s charge/momentum ( $q/p$ ).

There are two main backgrounds to the CC  $\bar{\nu}_\mu$  sample: CC  $\nu_\mu$  interactions whose charge has been mis-reconstructed (sometimes called ‘wrong sign’ or WS),

$$\nu_\mu + N \rightarrow \mu^- (\text{mis-id as } \mu^+) + \text{hadrons} \quad (5.6)$$

or neutral current (NC) interactions of any neutrino species which have a shower element that fakes a muon track,

$$\nu_x + N \rightarrow \nu_x + \text{fake muon} + \text{hadrons.} \quad (5.7)$$

The first selection step is to apply basic “preselection” cuts:

- The beam and detector must have been in good operating condition.
- The event must occur during the beam spill, greatly suppressing atmospheric neutrinos and cosmic rays.
- The event vertex must occur in a fiducial volume separated from the edges of the detector, ensuring that all the energy of the event is contained in the detector and can be measured.

---

<sup>2</sup>The detector mode always corresponds to the beam mode for the data in this thesis, though some data were taken with the Near Detector field reversed relative to the horn fields for systematics studies.




Figure 5.3: Contamination in the positive sample after preselection only. The black line represents the total contamination ( $1 - \text{purity}$ ), the red line represents the fraction of events that are mis-identified  $\nu_\mu$  CC (wrong sign) and the blue line represents the fraction of events that are neutral currents. After preselection the sample is still mostly background, especially at low energies.

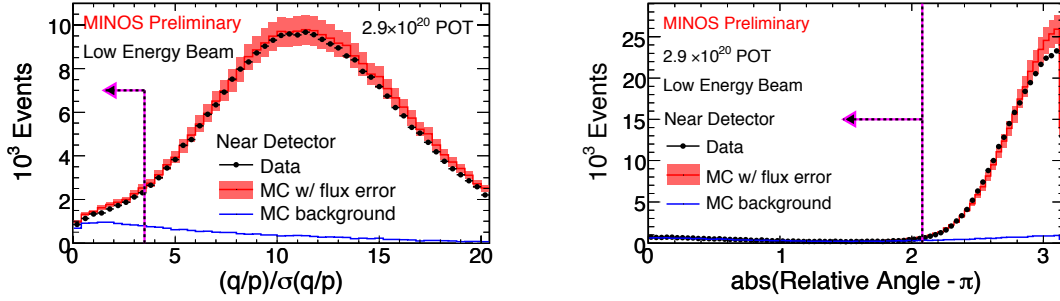


Figure 5.4: Charge-sign selection variables  $(q/p) / \sigma_{q/p}$  and  $|\text{Relative Angle} - \pi|$  are shown for the Near Detector in data (black points) and simulation (total in red, background in cyan). The flux uncertainty on the Monte Carlo is represented by the shaded red bars. In each plot, all other selection cuts have been applied. These selectors keep only well-measured positive tracks.

- The event must have a reconstructed track, eliminating most neutral current events.

The next step in selecting antineutrinos is to keep only events with positive reconstructed charge (the PQ sample). However, the antineutrino component is so small that even this sample is dominated by backgrounds (see Figure 5.3). The backgrounds are split evenly between neutral currents and wrong-sign neutrinos. The positive sample is at least half background at all energies, and the contamination gets worse below 7 GeV.

An additional selection step is applied for each of these backgrounds. Two cuts are made to address the charge sign, and a third cut is made on a likelihood-based CC/NC separator. The two



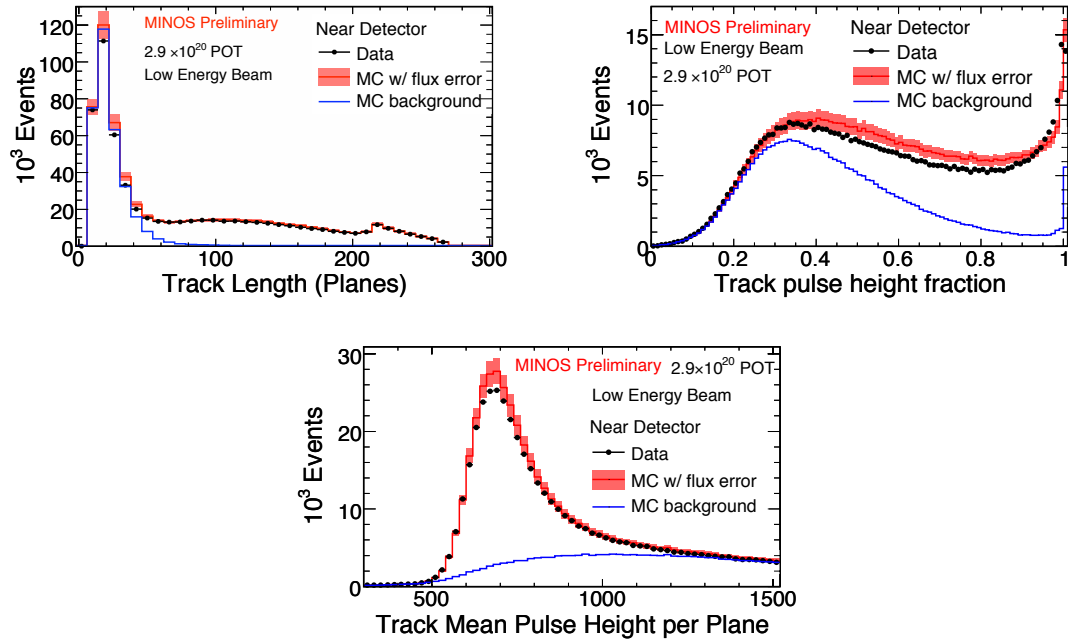


Figure 5.5: The three variables that form the PDFs in the  $DpID$  CC/NC separator: the track length, the fraction of the event energy in the track, and the mean energy deposited per plane are shown for the Near Detector in data (black points) and simulation (total in red, background in cyan). The flux uncertainty on the Monte Carlo is represented by the shaded red bars. Each shows some separation between the background and the bulk of the sample.

charge-sign selection variables are the ratio of the track's curvature ( $q/p$ ) to the uncertainty on that curvature ( $\sigma_{q/p}$ ) and the relative angle between the straight-line projections of the first few hits and the last few hits of the track. The distributions of these two variables are shown in Figure 5.4.

The CC/NC separation parameter, called  $DpID$ , is built up from 1-dimensional PDFs of three variables that each have some power to distinguish charged current events from neutral currents.

### Track length

True muon tracks tend to be longer (i.e. cross more planes) than the tracks of particles coming from the hadronic shower.

### Track energy fraction

The fraction of the event energy that is in the track (lepton) as opposed to the shower (hadrons). It is related to the kinematic  $y$  (inelasticity).

### Track energy per plane

The amount of energy deposited per plane of the track. It is related to  $dE/dx$  which can distinguish true muons (typically minimum-ionizing) from the tracks formed by hadronic shower components.

These three variables are shown in Near Detector data and simulation in Figure 5.5. The distribution




Figure 5.6: CC/NC separation parameter ( $DpID$ ) is shown for the Near Detector in data (black points) and total Monte Carlo with flux uncertainty (red line and shaded bars). Also shown is the neutral current background (magenta), mis-identified  $\nu_\mu$  background (green), and total background (blue). The CC/NC separator has some power to reject the wrong-sign background in addition to the neutral current background due to the higher average inelasticity ( $y$ ) of neutrinos compared to antineutrinos, which is one of the input variables. All other cuts have been applied.




Figure 5.7: The contamination in the sample of events passing preselection (dashed blue, same as the black in Figure 5.3) as well as the contamination after the full selection has been applied (solid blue). This contamination is made up of approximately equal parts neutral current interactions and mis-identified negative charged current interactions. Also shown is the selection efficiency (red). By comparing the dashed and solid blue lines, it is clear that the selection dramatically increased the purity of the antineutrino sample.

of the separation parameter, again in Near Detector data and simulation, is shown in Figure 5.6. This figure also shows the two individual background components. In addition to removing neutral current events,  $DpID$  is also effective at removing many wrong-sign events because one of its input PDFs is related to inelasticity (kinematic  $y$ ) and neutrino interactions typically have a higher  $y$  than antineutrinos.

Figure 5.7 shows the performance of the antineutrino selection in purity and efficiency<sup>3</sup> as well as the purity in the positive sample before selection (as shown in Figure 5.3). The CC/NC separator and the two selection cuts on charge-sign dramatically improve the purity of the sample from 34% to 97% while keeping the overall efficiency at 82% in the Far Detector.

## 5.2 Near Detector Spectrum

Now that a pure sample of antineutrinos has been obtained, the next step is to build up an energy spectrum from the selected events. The energy of each event is reconstructed in two parts: the muon track and the hadronic shower. For tracks that end within the detector (contained tracks), the range of the track is the best measure of its energy. If the track leaves the detector its energy is instead measured by the amount the track curves in the magnetic field. The shower energy is measured calorimetrically: it is the sum of the energy deposited into the scintillator strips for all non-track hits in the event. For more details see Section 3.4. The reconstructed track and shower energy distributions in the Near Detector are shown in Figure 5.8.

Combining the track and shower energy for each event, the full energy spectrum can be built up at the Near and Far Detectors. The selected antineutrino energy spectrum at the Near Detector in the neutrino-mode beam, in data and Monte Carlo, can be seen in Figure 5.9, along with the flux uncertainty on the Monte Carlo. Thanks to MINOS's two-detector design, this systematic




Figure 5.8: Reconstructed momenta of  $\mu^+$  tracks (left) and reconstructed shower energy (right) in the Near Detector. The red histogram represents the Monte Carlo expectation with the flux error, the blue histogram represents the total (charged and neutral current) background with the background uncertainty. Black points represent data.

<sup>3</sup>Calculated from simulation and taken relative to all true interactions in the fiducial volume.




Figure 5.9: The selected  $\bar{\nu}_\mu$  energy spectrum is shown in the Near Detector in data (black points) and simulation (total in red, background in cyan). The flux uncertainty on the Monte Carlo is represented by the shaded red bars.

uncertainty on the flux as well as the uncertainty from antineutrino cross-sections cancel in the extrapolation to the Far Detector (see Section 5.3 for more details on extrapolation). The figure illustrates that the amount of background in the final spectrum is relatively small, approximately 3%.

Note that the simulation shown has been reweighted to account for the mismodeling of hadronic interactions in the NuMI target using a fit to multiple beam running conditions. While necessary to understand the Near Detector distributions, it has little bearing on the oscillation result since flux uncertainties cancel in the extrapolation. See Section 3.1.4 for more details on the reweighting.

### 5.3 Extrapolation (the Beam Matrix)

The MINOS experiment was designed with two detectors in order to reduce the effect of systematic uncertainties. Many systematics, such as the neutrino flux, the neutrino cross section, and the modeling of the hadronic energy, affect both detectors in the same way, allowing them to effectively “cancel out.” For example, imagine the actual neutrino flux were 10% higher than simulated. With only one detector those 10% more events at the Far Detector would change the apparent oscillation probability. With two detectors, however, that flux increase would increase the number of events at the Near Detector as well. So, if we predict the Far Detector based on the Near Detector that 10% increase would be expected and thus not affect the measured oscillation probability.

There is an important caveat though. The example above relied on the change in the flux having




Figure 5.10: Diagram of the neutrino parents in the NuMI decay pipe. A parent will typically have a wide range of neutrino decay angles that reach the Near Detector and a very narrow range that will reach the Far Detector.

the same effect at both detectors. This comes from an implicit assumption that both detectors see an identical flux. However, for MINOS this is not quite true due to the kinematics of the decays that produce the neutrinos. For a given parent, the energy of the daughter neutrino in the parent's rest frame is fixed:

$$E_\nu^* = \frac{m_p^2 - m_\mu^2}{2m_p} \quad (5.8)$$

where  $m_p$  is the mass of the parent ( $p = \pi^\pm, K^\pm$ ),  $m_\mu$  is the mass of the muon and the neutrino mass is negligible. However, the energy of the neutrino in the lab frame depends on the relative angle,  $\theta$ , between the parent's direction of travel and the neutrino's:

$$E_\nu = \frac{E_\nu^*}{\gamma_p (1 - \beta_p \cos \theta)} \quad (5.9)$$

where  $\gamma_p$  is the parent's Lorentz factor and  $\beta_p$  is its velocity. The flux is also a function of angle. While the parent emits the neutrino isotropically in its rest frame, once boosted to the lab frame the flux becomes angle-dependent:

$$\frac{dN}{d \cos \theta} = \frac{1}{2\gamma_p^2 (1 - \beta_p \cos \theta)^2} \quad (5.10)$$

where again  $\gamma_p$  is the parent's Lorentz factor and  $\beta_p$  is its velocity. The derivations of these formulas can be found in Appendix A.

The Far Detector is sufficiently distant so that for a given parent, there is only one narrow range of angles that will produce a neutrino that will reach the Far Detector, uniquely determining the neutrino energy from that parent. The Near Detector, however, covers a much wider solid angle since it is significantly closer to the end of the decay pipe. Consequently, this same parent can produce neutrinos at a range of energies in the Near Detector (see Figure 5.10). The effect on the spectrum can be seen in Figure 5.11: neutrinos of a particular energy in the Near Detector (shaded regions on the left) correspond to a range of parent energies with different decay angles which will produce a range of energies at the Far Detector (shaded regions on the right). The Far Detector distribution

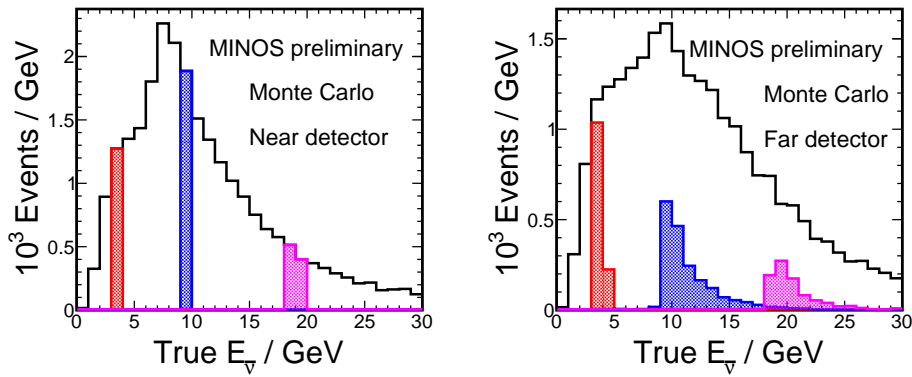


Figure 5.11: The relationship between the energies of  $\bar{\nu}_\mu$  events observed in the Near Detector with those observed in the Far Detector. The colored regions on the left and right show the differing neutrino energy distributions in the detectors for neutrinos coming from the same parents.

gets wider at higher energies since higher energy parents tend to get further down the decay pipe (i.e. closer to the Near Detector) before decaying, enhancing the solid-angle effect (and outweighing the increased Lorentz boost which tends to narrow the outgoing neutrino energy distribution). Since the larger angles tend to have lower energies, events at the Near Detector tend to shift downward into the peak, leading to a more peaked spectrum at the Near Detector than at the Far Detector.

In order to get a correct prediction of the Far Detector given the Near Detector spectrum, the Monte Carlo is used to create a “beam matrix” that relates the Near Detector spectrum to the Far Detector spectrum.<sup>4</sup> The rows each correspond to a bin of Far Detector true neutrino energy and the columns each correspond to a bin of Near Detector true neutrino energy. Each whole column is effectively the Far Detector neutrino energy spectrum that would be produced by the collection of parents that produced neutrinos at the Near Detector energy to which that column corresponds. Each column is normalized to one Near-Detector neutrino so that by matrix multiplying with the Near Detector spectrum the Far Detector spectrum is obtained. The matrix for antineutrinos in the neutrino-mode beam is illustrated in Figure 5.12.

The matrix is populated by taking many simulated neutrino parents and forcing them to decay towards both the Near Detector<sup>5</sup> and the Far Detector. The corresponding neutrino energies and the probabilities of those decay directions are calculated. The matrix is then filled in the element defined by the two detector energies with a weight defined by the probability of that parent producing a neutrino at each detector. Once the full matrix is filled each column is normalized to a single Near Detector neutrino as described above.

Since the matrix is based on simulation, it assumes that the Near and Far Detector spectra are in

<sup>4</sup>I played a key role in adapting the extrapolation, designed for the neutrino analysis, to antineutrinos.

<sup>5</sup>Locations within the Near Detector are selected randomly in three dimensions, the same procedure used in generating the Near Detector simulation. This more accurately samples the range of decay angles that reach the Near Detector.




Figure 5.12: The beam matrix for  $\bar{\nu}_\mu$ 's in the neutrino-mode beam. Each cell relates a Far Detector energy bin to a Near Detector one. The content of each cell represents the mean number of  $\bar{\nu}_\mu$  events expected in the Far Detector for one event in the Near Detector. This distribution is treated as a matrix to relate the energies measured in the Near Detector to those expected in the Far Detector.

true neutrino energy, are perfectly pure, and that events are selected perfectly. In reality, of course, there are backgrounds and selection efficiencies, and only the reconstructed visible energy of each event is available. Each of these effects (purity, efficiency, visible energy) is corrected using the full detector Monte Carlo, both at the Near Detector and at the Far Detector. Backgrounds and efficiency are corrected by multiplying or dividing the appropriate histogram. The reconstructed visible energy is converted to corrected ('true') neutrino energy by matrix-multiplying a two-dimensional histogram which relates reconstructed and true energies in the Monte Carlo. This way the energy spectrum of selected antineutrinos in the Near Detector can be transformed into the neutrino flux  $\times$  cross section in neutrino energy at the Near Detector. This Near Detector flux  $\times$  cross section is multiplied by the beam matrix to produce a corresponding Far Detector flux  $\times$  cross section in neutrino energy. The corrections applied for the Near Detector, but now based on the Far Detector simulation, are then applied in reverse: the spectrum is converted back to reconstructed energy by multiplying with a true-to-visible matrix, the spectrum is reduced by multiplying the selection efficiency, and data-based background predictions are added, finally resulting in a *prediction* of the Far Detector spectrum given the spectrum observed in the Near Detector without any oscillations. The prediction can be produced for any choice of oscillation parameters by applying the oscillation probability to the unoscillated Far Detector flux  $\times$  cross section before the other corrections are applied.

## 5.4 Systematics

There are a number of systematic uncertainties with the potential to effect the oscillation and transition measurements. They fall into a few general categories: energy reconstruction, backgrounds, and extrapolation.

### 5.4.1 Energy Reconstruction Systematics

#### Track energy scale

Track energy as measured by range has a systematic uncertainty of 2% determined using CalDet, CalDet (see Section 3.4). Track energy as measured by curvature has an additional 1% uncertainty, determined by comparing range and curvature momentum measurements for stopping tracks. These uncertainties are taken as fully correlated between the two detectors.

#### Relative shower energy scale

The relative shower energy systematics come from uncertainties in the energy calibration procedure using cosmic ray muons (see Section 3.3). Data-simulation differences in the various calibration steps are added in quadrature and give an uncertainty of 2.4% in the Near Detector and 2.3% in the Far [149]. The errors are uncorrelated between the detectors.

#### Absolute shower energy scale

The absolute shower energy systematic uncertainty is taken as fully correlated between the two detectors and has two major components. The first component stems from uncertainties in the detector response to single hadrons as measured in the calibration detector at the CERN test beam and is 5.7% at all energies (see Section 3.4). The second component is energy-dependent and encapsulates uncertainties in hadron production and intranuclear effects. It is 8.2% at the lowest energies, dropping off to 3% above 10 GeV [150]. The final systematic has the energy-dependent form

$$\sigma_{\text{shw}} = 6.6\% + (3.5\%) \times e^{\frac{-E_{\text{shw}}}{1.44 \text{ GeV}}} \quad (5.11)$$

which is taken as fully correlated bin-to-bin.

### 5.4.2 Background Systematics

#### Neutral current and charged current $\nu_\mu$ backgrounds

The systematic uncertainty on the neutral current and charged current  $\nu_\mu$  backgrounds was evaluated by comparing data and Monte Carlo in a background-



Systematic	Uncertainty
Steel Thickness	0.2%
Scintillator Thickness	0.2%
FD Live Time	1.0%
ND Fiducial Bias (z)	1.9%
ND Fiducial Bias (y)	0.7%
ND Fiducial Bias (x)	0.7%
N/F Selection Bias	3.0%

Table 5.1: Components of the Near-to-Far normalization systematic uncertainty.

dominated sample of events with  $0 < DpID < 0.25$ , giving a systematic uncertainty of 50% on the amount of NC background. Due to the higher  $y$ -distribution of  $\nu_\mu$ 's relative to  $\bar{\nu}_\mu$ 's, this region had relatively equal amounts of both background components (see Figure 5.6). The systematic was evaluated assuming that one, the other, or both backgrounds were responsible for the data-simulation discrepancies in this region, and it showed that a 50% systematic uncertainty covers all cases.

### 5.4.3 Extrapolation Systematics

#### Near-to-Far normalization

The 4% normalization systematic incorporates several systematic uncertainties, all of which change the relative number of events expected at the two detectors per POT. It is dominated by a 3% uncertainty on the difference in reconstruction and selection efficiency in the two detectors. Table 5.1 shows all the contributions. The ‘fiducial bias’ uncertainties refer to data-simulation differences in the non-uniformity of the vertex distributions coming from acceptance effects due to the geometry of the Near Detector.

#### Downstream events

A +50% and –100% uncertainty is used on the number of neutrinos coming from interactions outside of the target as described in Section 5.4.4.

#### Cross-sections

Several systematic effects are evaluated as part of the cross section systematic, both on the overall cross section and on various NEUGEN [144] interaction model parameters. Some affect both neutrinos and antineutrinos and others are specific to antineutrinos. While there is some residual systematic effect because of spectral differences between the detectors, the effect is negligibly small in the oscillation analysis.<sup>6</sup>

<sup>6</sup>The cross section uncertainty does become important for the transition analysis, but only the case of a large signal where size of  $\alpha$  is anti-correlated with the  $\nu_\mu/\bar{\nu}_\mu$  cross section ratio.

### Flux modeling

The flux modeling uncertainty encapsulates a number of sources of error, including hadron production, beam optics (horn positions, currents, etc.), the position of the target, and the amount of material in the beamline. The flux errors are evaluated by moving around the fit parameters in the beam tuning fit within their uncertainties and observing the effect on the flux. Again, the majority of the errors cancel between the two detectors, but some residual uncertainty remains because the two detectors do not see identical fluxes.

### 5.4.4 Downstream Parents

The antineutrinos in the neutrino-mode beam, making up only 7% of the event rate, are sensitive to details of the flux that are negligible in other analyses which use the focused events. The vast majority of the focused events come from mesons that were produced in the target. These mesons then travel along the decay pipe where they decay into neutrinos. However, mesons are also produced due to collisions with other pieces of material downstream of the target. These outside-the-target collisions involve both the primary protons and the secondaries produced in the target. These downstream parents produce both neutrinos and antineutrinos, but are only a significant contribution to the smaller antineutrino flux. The relative contributions to the Near and Far Detector antineutrino event rates from parents produced in different regions of the beamline can be seen in Figure 5.13.<sup>7</sup>




Figure 5.13: The simulated Near Detector spectrum and Far/Near ratio for the various antineutrino parents:  $\pi^-$  and  $K^-$  produced upstream and parents produced in the decay pipe. Note that the decay pipe component is primarily at low energies and has a distinctly different Far/Near ratio from the other flux components.

<sup>7</sup>I studied downstream production of antineutrinos in detail.




Figure 5.14: The Near (color) and Far (black) Detector spectra for parents produced in the decay pipe (left) and upstream (right) regions. The upstream spectrum consists mostly  $\bar{\nu}_\mu$ 's from parents produced in the target, but it also includes parents produced in the horns and other material in the target region. The plots are normalized to summed area, that is the summed area of the two Near Detector plots is equal to that of the two summed Far Detector plots. This shows that the decay pipe component looks very different to the Near and Far Detectors while the upstream region looks relatively similar.

Of particular concern are parents produced in the decay pipe – they tend to decay very close to the Near Detector, leading to a large difference in their flux between the Near and Far Detectors. They also appear primarily at low energies where the oscillation and transition signals would be most prominent. This can be seen dramatically in Figure 5.14. The colored curves represent the spectrum in the Near Detector and the black represent the spectrum in the Far Detector. While neutrinos from upstream parents produce similar spectra at both detectors, the neutrinos from decay pipe parents have a much higher rate at the Near Detector than they do at the Far Detector (relative to the total rate at each respective detector). If the number of neutrinos coming from parents produced in the decay pipe were mis-modeled, it would make a relative change in the low-energy event rate between the detectors – precisely the signal being searched for in these analyses. The production of particles outside the target is thus a major systematic uncertainty for the antineutrino analysis.

A method was proposed to determine a systematic uncertainty on the simulation's downstream hadronic interaction model using the difference between the data taken with an evacuated decay pipe and data taken with helium in the decay pipe (leading to an enhanced rate of downstream interactions). However, the method ended up being impossible due to concerns about the applicability of a constraint on interactions with helium to interactions with steel (the majority of parents produced in the decay pipe are produced in interactions with the iron walls) as well as the simulation's poor modeling of the effect of the helium (see Section 3.1.4 for a more complete discussion of the modeling of helium in the decay pipe).

In the end, a simpler, if cruder, constraint on the decay pipe uncertainty was used. For the purpose of evaluating the systematic, it was assumed that all data-Monte Carlo discrepancies at the Near Detector were due to the decay pipe, both nominally and with all systematic uncertainties applied in both directions. The systematic uncertainty on the decay was taken as the largest shift

in each direction required to make the integral number of low energy events agree between data and Monte Carlo. The systematic uncertainty was evaluated as  $^{+50\%}_{-100\%}$ . It was more a worst-case scenario than a true  $1\sigma$  uncertainty, but it was acceptable since it certainly covered our uncertainty, and it had a similar sized effect on the result as the other systematic uncertainties (that is, tiny compared to the statistical uncertainty). It is described in detail in Appendix B.

### 5.4.5 Effect on the Analyses

The effect of each systematic uncertainty on the oscillation and transition results is estimated using the simulation. Systematic shifts are applied to Monte Carlo events to produce shifted high-statistics Near Detector and Far Detector spectra. The shifts are applied both positively and negatively, producing two sets of spectra. The total systematic uncertainty can then be examined several ways.

Figure 5.15 shows the Far Detector systematic error band constructed from all the systematic uncertainties summed in quadrature. The correlation in the systematics between the two detectors, which generally leads to cancellation, needs to be taken into account. The systematically shifted Near Detector spectrum is extrapolated to the Far Detector, producing a systematically shifted prediction. The shift in the systematically shifted Far Detector spectrum is then divided out of the shifted prediction, approximating the cancellation that occurs when fitting.

An oscillation analysis is also performed (see Section 5.6) for each systematic shift using the systematically shifted Near and Far Detector spectra as fake data. The amount the best fit moves compared to using the nominal Monte Carlo shows the size that systematic effect on the oscillation




Figure 5.15: Total systematic error band on the Far Detector prediction. The band is obtained by adding the effect of each individual systematic shift on the FD predicted energy spectrum in quadrature.




Figure 5.16: The shifts in the best fit oscillation parameters induced by the application of systematic shifts to the fake data. The cross section systematic is the sum in quadrature of all the component cross section systematics.




Figure 5.17: The shifts in the best fit transition probability as a function of true transition probability. The colored bars represent the individual  $1\sigma$  errors (including statistics in red) and the white error bar represents the quadrature sum of all the errors. Note that the total error consists almost entirely of the statistical error, especially at low transition probability.

result. The sizes of these shifts can be seen in Figure 5.16.

Finally, a transition analysis is performed (see Section 5.7) for each systematic shift using the systematically shifted Near and Far Detector spectra as fake data at a range of possible transition probabilities. The amount the best fit moves compared to using the nominal Monte Carlo shows the size that systematic effect on the transitions result. The size of each individual uncertainty is represented by the colored bars in Figure 5.17 and the sum of them all in quadrature is represented by the white error bars. Note that the total error consists almost entirely of the statistical error, especially at low transition probabilities.

## 5.5 Far Detector Data

Figure 5.18 shows the Far Detector data spectrum along with the prediction based on Near Detector data without oscillations, the prediction with the same oscillation parameters measured for neutrinos in [77], and the predicted background spectrum. 42 events were observed with an unoscillated expectation of  $66.6 \pm 8.2(\text{stat}) \pm 5.0(\text{syst})$  and an oscillated expectation of  $60.3 \pm 7.8(\text{stat}) \pm 4.6(\text{syst})$  assuming the neutrino oscillation parameters as above. The two-sigma deficit is at high energies so it is not associated with previously observed oscillations [63, 65, 66, 77]. In order to be sure there is not a problem with selection, comparisons between Far Detector data and simulation in the selection variables are shown in Figures 5.19-5.21. The deficit is limited to the signal region in the CC/NC separation parameter and appears at low  $(q/p) / \sigma_{q/p}$  since the deficit is at higher energies (higher energy tracks have poorer charge-determination). There is no evidence of an excess of events being inadvertently removed by a problem with the selection. Also shown are the track vertex and end positions in Figure 5.22. The vertex plots show an apparent asymmetry with no event vertices in the right-most section of the detector, but, as described in the next section, there is no evidence of a detector problem in this region and fake data studies show that such an asymmetry is reasonably likely with so few events.




Figure 5.18: Far Detector  $\bar{\nu}_\mu$  data spectrum (black points) compared to predictions with no oscillations (solid red histogram) and with the  $\nu_\mu$  oscillation parameters from [77] (dashed blue histogram). The background is also displayed (gray shaded histogram).




Figure 5.19: CC/NC separation parameter ( $DpID$ ) is shown in the Far Detector in data (black points), oscillated simulation (solid red line) and unoscillated simulation (dashed red line). All other selection cuts have been applied. The line and arrow mark the region removed by the cut. The deficit appears only in the signal region and does not appear to be due to a mismodeling of the CC/NC separator.




Figure 5.20: Charge-sign selection variable  $(q/p) / \sigma_{q/p}$  is shown in the Far Detector in data (black points), oscillated simulation (solid red line) and unoscillated simulation (dashed red line). All other selection cuts have been applied. The deficit appears largest at low values since this variable is correlated with energy: higher energy tracks have poorer charge-determination. The line and arrow mark the region removed by the cut.




Figure 5.21: Charge-sign selection variable  $|\text{Relative Angle} - \pi|$  is shown in the Far Detector in data (black points), oscillated simulation (solid red line) and unoscillated simulation (dashed red line). All other selection cuts have been applied. The line and arrow mark the region removed by the cut. There appear to be no missing events at lower relative angles, as might be expected if there were a problem with this selector.






Figure 5.22: Selected antineutrino event vertex (left) and end (right) positions in the Far Detector as a function of  $x$ - and  $y$ -coordinates. The apparent asymmetry in vertex  $x$  positions was shown to be consistent with statistical fluctuations ( $p > 0.05$  compared to distributions from pseudo-experiments), given the number of events observed.

### 5.5.1 Cross-checks

After the box was opened, several cross-check studies were undertaken. The first step was to understand how unlikely it was to observe so large a deficit. Figure 5.23 shows the distribution of the number of selected events in 10,000 fake experiments, including systematic shifts. The probability of seeing 42 events or below was 2.4%, but the probability of being as far from the mean (high or low) as observed was 5.1%. The likelihood of the left-right vertex asymmetry was also examined and found to be consistent with statistical fluctuations: comparisons with pseudo-experiment distributions gave  $p$ -values  $> 0.05$  for the values measured in the data.<sup>8</sup>

A study was also undertaken to see what effect a mismodeling of precisely which events the selection cuts remove in the Near Detector might produce. This effect was modeled by changing the selection cuts up and down in only the Near Detector and seeing the size of this effect on the Far Detector prediction. The ratios of the shifted to nominal predictions are shown in Figure 5.24 and the changes are extremely small between 5 GeV and 30 GeV where the majority of the antineutrino spectrum is.

The event rate over time was also examined. If the deficit were concentrated in a single time period it would suggest a beam or detector problem. Figure 5.25 shows the number of Far Detector antineutrino events per proton-on-target in 4-month blocks. There is no evidence of a single time period being responsible for the deficit – it appears constant over time. Figure 5.26 shows the Near Detector antineutrino event rate. An excess of Near Detector events could cause an apparent deficit at the Far Detector, but the Near Detector event rate appears stable over time, as well.

<sup>8</sup>I performed the pseudo-experiment studies of the event counts and vertex asymmetry.




Figure 5.23: The distribution of the number of selected events in 10,000 fake experiments. In addition to randomly selected events, each experiment also has a random set of systematic shifts so the effects of systematic errors are included. Our mean experiment had 58.3 events and our actual experiment had 42 events. The probability of getting 42 events or below is 2.4% and the probability of being as far from the mean as observed is 5.1%.




Figure 5.24: Ratio of the nominal Far Detector prediction to the modified predictions with higher and lower cuts on the individual selection variables. The high and low cut values are chosen such that the efficiency  $\times$  purity of each modified selection is 3% lower than the maximum. None of these changes can reproduce the deficit seen in the Far Detector.

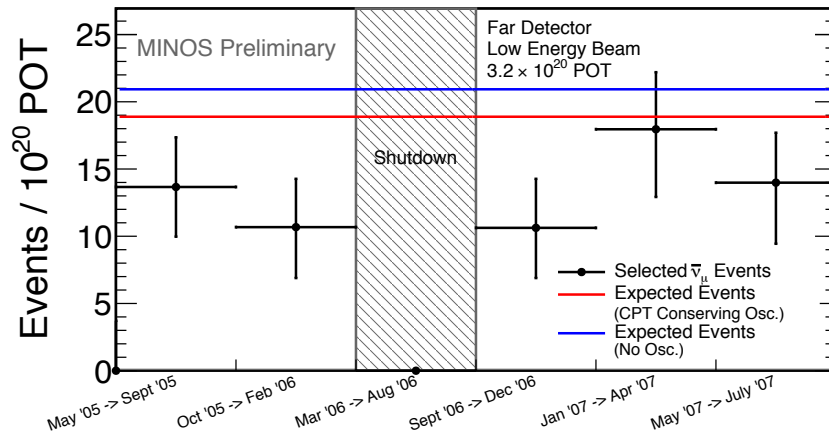


Figure 5.25: Number of  $\bar{\nu}_\mu$  events in the Far Detector per  $10^{20}$  POT per period of 4 months. This plot includes Run I and II data represented by black points with Poisson errors. The red line represents the expected event rate assuming the same oscillation parameters as measured for neutrinos [77]. The blue line represents the expected event rate in the absence of oscillations. The deficit appears to be in all time periods.




Figure 5.26:  $\bar{\nu}_\mu$  event rate per  $10^{16}$  POT in the Near Detector over time. A fit to a constant is provided. There is no evidence of a period of abnormally high Near Detector data that might explain the relative deficit.

## 5.6 Oscillation Analysis

The oscillation parameters are measured using a two-parameter fit for  $|\Delta\bar{m}_{\text{atm}}^2|$  and  $\sin^2(2\bar{\theta}_{23})$  from Equation 5.1. The fit finds the oscillation parameters that maximize the likelihood (minimize the quantity  $-2\ln L$ ) of the data given the prediction at those oscillation parameters:

$$-2\Delta\ln L(\mathbf{a}) = 2 \sum_i \left[ p_i(\mathbf{a}) - d_i + d_i \ln \frac{d_i}{p_i(\mathbf{a})} \right] \quad (5.12)$$

where  $\mathbf{a} = [\Delta\bar{m}_{\text{atm}}^2 \sin^2(2\bar{\theta}_{23})]$  contains the fit parameters,  $i$  counts over energy bins,  $d_i$  is the number of data events in bin  $i$ , and  $p_i(\mathbf{a})$  is the predicted number of events in bin  $i$  given the parameters in  $\mathbf{a}$ . Note that the prediction,  $p_i$ , is generated based on the Near Detector data as described in Section 5.3. The oscillation formula (Equation 5.1) is applied to the predicted spectrum in true energy before it is transformed into reconstructed energy for comparison with the data. Once a likelihood surface has been produced, the contours are drawn using the Feldman-Cousins method as described in Section 5.8.

The allowed oscillation contours are shown in Figure 5.27. The MINOS neutrino oscillation parameters from [77] are contained within the 90% contour. The best fit is at a relatively high mass splitting since the spectral deficit is at high energy, and even higher mass splitting cannot be excluded. Much of the power of the analysis comes from counting events, rather than shape information, due to its low statistics. Consequently, the results cannot distinguish the best fit point with a moderate mixing angle from fast oscillations with a large mixing angle. The one-dimensional exclusion curve, assuming maximal mixing, is shown in Figure 5.28. This data improves the exclusion from the global fit for a range of low mass-splitting values.

## 5.7 Transition Analysis

The transitions fit is performed in the same way as the oscillation fit described above, except it is a one-parameter fit: the vector of fit parameters in Equation 5.12 becomes  $\mathbf{a} = [\alpha]$  where  $\alpha$  is the probability that a neutrino transitions into an antineutrino from Equation 5.2. Again, once a likelihood curve has been produced, the confidence intervals are determined using the Feldman-Cousins method as described in Section 5.8.

The best fit was to no transitions or  $\alpha = 0$ . The 68%, 90%, and 99% upper limits on  $\alpha$  are 0.0067, 0.026, 0.067, respectively. The  $-2\Delta\ln L$  surface and exclusion confidence as a function of transition probability can be seen in Figure 5.29. The final results are better than our sensitivity since the low number of events in the Far Detector data makes an appearance signal even more unlikely.




Figure 5.27: The 68% (red), 90% (blue), and 99.7% (black) results contours are shown. The best fit point is marked as a triangle. The white area is excluded at 99.7% confidence. The MINOS best fit  $\nu_\mu$  parameters from [77] are included in the 90% contour.




Figure 5.28: The blue line represents the level at which we can exclude oscillations at a given mass splitting assuming maximal mixing. The black line represents the equivalent, but based on the global fit from Maltoni *et al.* [76]




Figure 5.29: At left is the  $-\Delta \ln L$  curve from the fit for Transitions. It shows both the original curve (black) and the Feldman-Cousins corrected curve for 90% confidence with systematics. At right is a plot showing the Feldman-Cousins corrected exclusion confidence as a function of transition probability.

## 5.8 Feldman-Cousins Method

The Feldman-Cousins method [151] was used in the antineutrino analysis in order to draw correct confidence limit contours for a given level of significance despite non-Gaussian errors. It also includes the effects of systematic uncertainties.<sup>9</sup>

The idea behind Feldman-Cousins is to use fake experiments to build up empirical  $-2\Delta \ln L$  distributions for the fit as a function of the parameter values. For each choice of parameter values, fake experiments are generated assuming those parameter values. For each experiment, the  $-2\Delta \ln L$  between the true parameters chosen for that experiment and the best fit are recorded. Over many fake experiments this builds up an empirical  $-2\Delta \ln L$  distribution which can be used to determine what value of  $-2\Delta \ln L$  is required to cover, for example, 90% of fake experiments at that choice of true parameters. This procedure is repeated at many parameter values building up a line or surface out of the  $-2\Delta \ln L$ 's required to cover 90% of experiments. The actual contours are drawn by comparing the likelihood surface at a given set of parameters to the value on the Feldman-Cousins surface. If the likelihood at a particular point is below the Feldman-Cousins surface for those parameters, that point is inside the contour.

The classic case for needing a Feldman-Cousins correction is when drawing confidence intervals near a physical boundary. The boundary gives extra information so the gaussian confidence interval will be too conservative. Put another way, a lower  $-2\Delta \ln L$  value is required to cover a given percentage of fake experiments. In practice, this comes about because fake experiments that would have been unphysical “pile-up” at the physical boundary, meaning they have a smaller  $-2\Delta \ln L$  from the truth than they would have had otherwise. An example of this effect, in the transition




Figure 5.30: Example  $-2\Delta \ln L$  distributions for two transition probabilities. On the left is the distribution for  $\alpha = 6.67 \times 10^{-3}$  (very close to the physical boundary at 0). On the right is the distribution for  $\alpha = 0.6$ , well away from any boundaries. The black vertical lines are at  $-2\Delta \ln L = 2.71$  in both plots, the canonical one-dimensional 90% coverage value. The blue lines show the actual  $-2\Delta \ln L$  required to encompass 90% of fake experiments. On the right, away from the boundary, this value is very close to the canonical value. On the left, however, many experiments are piled up at lower  $-2\Delta \ln L$ 's so a  $-2\Delta \ln L = 1.41$  is sufficient to cover 90% of experiments.

<sup>9</sup>I developed and implemented the Feldman-Cousins procedure used in the antineutrino analysis.




Figure 5.31: Feldman-Cousins 90% coverage grid with systematic shifts included in the individual experiments. The color on the coverage grid represents the  $-2\Delta \ln L$  required to cover 90% of experiments. If the experiment was perfectly gaussian (and without systematics), the grid would be a uniform 4.61 (light blue) everywhere.

analysis, can be seen in Figure 5.30.

The two-dimensional Feldman-Cousins surface for oscillations can be seen in Figure 5.31. The oscillation surface is shown for 90% confidence so the canonical two-dimensional value is 4.61. Areas near the physical boundary are typically below this value due to the effect described above. The regions around  $\sin^2(2\bar{\theta}_{23}) \approx 0.5$  and at high  $|\Delta\bar{m}_{\text{atm}}^2|$  tend to be above this canonical value since there is a degeneracy between them: fast oscillations (high  $|\Delta\bar{m}_{\text{atm}}^2|$ ) and  $\sin^2(2\bar{\theta}_{23}) \approx 0.5$  both reduce the overall event rate by about a half. While these two hypotheses do produce different low energy spectral shapes, there are not enough statistics in the antineutrino data to resolve the






Figure 5.32: The red, blue and green lines show the  $-2\Delta \ln L$  value needed to cover 68%, 90% and 99% of mock experiments as a function of the true transition probability. The dashed lines show the statistical-only case whereas the solid lines show the case where systematics were also included in each of the mock experiments. The horizontal dot-dash lines at 1.0, 2.71, 6.67 are the one-dimensional Gaussian expectation. The effect of the physical boundaries is clear: the curves drop below the gaussian line near  $\alpha = 0$  and  $\alpha = 1$  and without systematics the curves reproduce the gaussian expectation everywhere else. Adding systematics always increases  $-2\Delta \ln L$  value required, leading to widened confidence intervals as expected.

difference.

The one-dimensional Feldman-Cousins curves for transitions can be seen in Figure 5.32. Shown are the 68%, 90%, and 99% curves for gaussian, statistics-only, and statistics+systematics. Again, the effect of the physical boundaries is clear: the curves drop below the gaussian line near  $\alpha = 0$  and  $\alpha = 1$ . Well away from the physical boundaries without systematics the curves reproduce the gaussian expectation. Adding systematics always increases  $-2\Delta \ln L$  value required. This increase leads to widened confidence intervals as expected. The effect of systematics increases with transition probability since the larger the appearance signal, the more significant the neutrino cross section uncertainties become.



## Chapter 6

# Antineutrinos in an Antineutrino Beam

The sensitivity of the antineutrino oscillation analysis with the neutrino-mode beam, presented in the previous chapter, is fundamentally limited for two reasons: the small size of the antineutrino component of the neutrino-mode beam, and the high average energy of the antineutrino events. These two effects limit the sensitivity to oscillations in the parameter region where oscillations have already been measured for neutrinos. The solution to both of these shortcomings is the same: a focused antineutrino beam. Such a beam can be produced by reversing the direction of the current in the NuMI horns, focusing negative mesons instead of positive ones, giving an antineutrino focusing peak. MINOS accumulated  $1.7 \times 10^{20}$  POT with NuMI in antineutrino mode thanks to a proposal from the Caltech neutrino group.

### 6.1 Selection

The first step in investigating the antineutrino-mode data is to select a pure sample of signal interactions, in this case charged-current anti-muon neutrino ( $CC\bar{\nu}_\mu$ ) interactions. Of course, the antineutrinos cannot be observed directly, so instead the goal is to identify the remnants of the antineutrino hitting a nucleus,

$$\bar{\nu}_\mu + N \rightarrow \mu^+ + \text{hadrons.} \quad (6.1)$$

As with the analysis of antineutrinos in the neutrino beam (see Section 5.1), the primary backgrounds are  $CC\nu_\mu$  interactions whose charge has been mis-reconstructed (sometimes called ‘wrong sign’ or WS),

$$\nu_\mu + N \rightarrow \mu^- (\text{mis-id as } \mu^+) + \text{hadrons} \quad (6.2)$$




Figure 6.1:  $q/p$  distribution of selected events before charge sign selection in the Near Detector. The red curve represents MC expectation with the flux uncertainty and black dots represent data. Antineutrino-like events are on the right (positive) and neutrino-like events are on the left (negative).

or neutral current interactions of any neutrino species which have a shower element that fakes a muon track (NC's),

$$\nu_x + N \rightarrow \nu_x + \text{fake muon} + \text{hadrons}. \quad (6.3)$$

The first selection step is to apply the same preselection cuts used in the neutrino-mode analysis:

- The beam and detector must have been in good operating condition.
- The event must occur during the beam spill, eliminating atmospheric neutrinos and cosmic rays.
- The event vertex must occur in a fiducial volume separated from the edges of the detector, ensuring that all the energy of the event is contained in the detector and can be measured.
- The event must have a reconstructed track, eliminating most neutral current events.

The first piece of the selection proper is a cut on the reconstructed charge of the muon track to eliminate the bulk of the  $CC\nu_\mu$ 's in the sample. In practice, the quantity measured in the detector is the track's curvature, which is proportional to the ratio of charge to momentum ( $q/p$ ). The distribution of this quantity in the Near Detector is shown in Figure 6.1. Only tracks with a positive reconstructed charge ( $q/p > 0$ ) are accepted.

The second piece of the selection is a cut on a multivariate CC/NC separation parameter called *RoID*. It is the output of a 4-parameter k nearest neighbors (kNN) algorithm [152]. A kNN uses a multi-dimensional space populated with simulated events, both signal and background. A metric

based on these 4 variables defines the distance between two events. The separation parameter for a given input event is then constructed based on what fraction of the  $k$  nearest neighbors ( $k$  events with the smallest distance from the input event) are CC events.

*RoID* uses 4 variables (seen in Figure 6.2) which can distinguish between charged current and neutral current interactions. The first two variables described below are topological and the second two describe the energy deposition along the track.

#### **Number of active planes in the track**

True muon tracks tend to be longer (i.e. cross more planes) than the tracks of particles from the hadronic shower.

#### **Transverse profile parameter**

Also called fraction of pulse height in the track, this variable looks at the amount of energy deposited in the transverse vicinity of the track, away from the shower at the vertex. True muon tracks typically leave only a single hit on a scintillator plane while non-muon tracks typically have other unconnected hits sitting near the reconstructed track. So, the larger the fraction of energy in the track as opposed to surrounding it, the more likely that the track is a muon.

#### **Average pulse height per plane in the track**

The average deposited energy (pulse height) per plane of the track away from the event vertex. It is related to  $dE/dx$  and distinguishes muons, which are typically minimum-ionizing, from the more energetic interactions of hadronic shower particles.

#### **Ratio of mean low pulse height to mean high pulse height**

This variable compares the mean energies of the lowest and highest energy strips in the track. Muon energy loss is relatively uniform, typically occurring through ionization where large energy losses are rare. Hadronic shower energy loss happens through completely different physical processes with much larger fluctuations in deposited energy. The closer in energy the low mean and the high mean are, the more likely the event is a muon.

The final output parameter from *RoID* in the Near Detector is shown in Figure 6.3. While this variable primarily distinguishes charged current interactions from neutral current interactions, it also eliminates some wrong sign background:  $CC\nu_\mu$ 's have a higher average  $y$ -distribution than  $CC\bar{\nu}_\mu$ 's, meaning  $\nu_\mu$ 's tend to look a little more NC-like.

The efficiency and purity performance of the selection in the Far Detector without oscillations are shown in Figure 6.4. The overall reconstruction and selection efficiency, with a CC/NC cut at

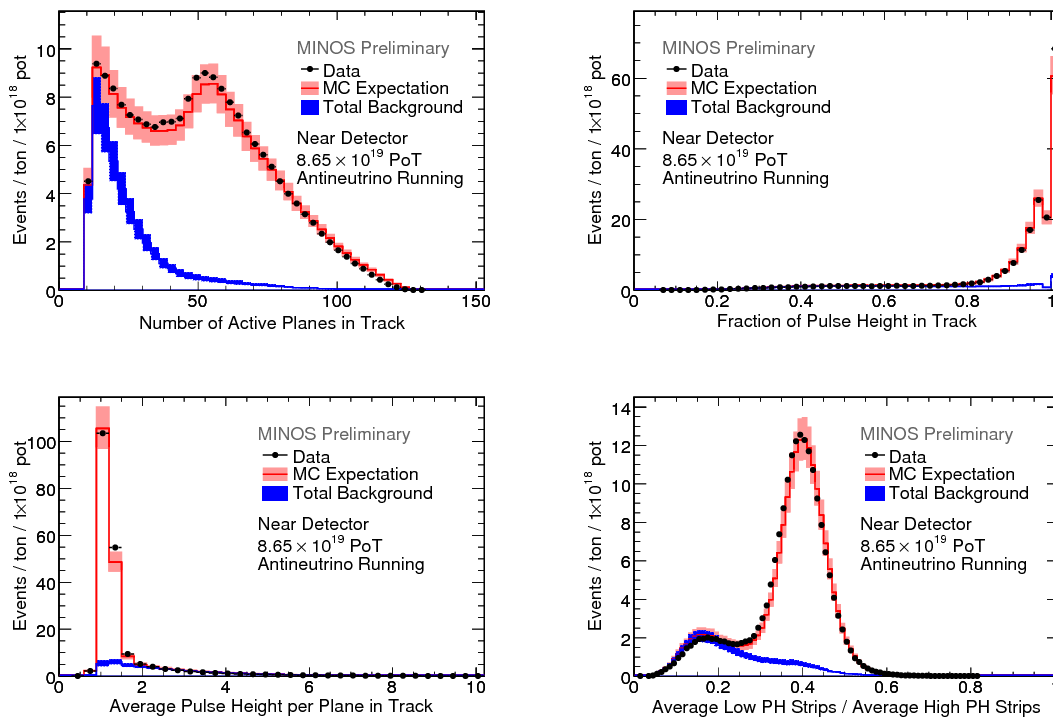


Figure 6.2: Distribution of the 4 kNN input variables before the CC/NC selection cut is applied. The red histogram represents the Monte Carlo expectation with the flux error, the blue histogram represents the total (charged and neutral current) background with the background uncertainty. Black points represent data. Each shows some separation between the background and the bulk of the sample.




Figure 6.3: CC/NC separation parameter on a semi-log scale. The red histogram represents the simulation with the flux error, the blue histogram represents the total background with the background uncertainty. Black points represent data. Events with a CC/NC separation parameter less than 0.3 (30% CC's in the nearest neighbors) are cut. The peak in the background at high-PID contains  $CC\nu_\mu$ 's whose charge-sign has been mis-identified.




Figure 6.4: Performance of the antineutrino selection ( $RoID > 0.3$ ,  $q/p > 0$ ) in the Far Detector. The dashed lines show the contamination before selection and the solid show efficiency and contamination after selection. The  $CC\nu_\mu$  contamination rises at higher energies because these tracks do not curve as much and so are more difficult to reliably assign a charge to. The background is reduced to nearly nothing in the 1 – 6 GeV region most sensitive to oscillations.

0.3, is 93%<sup>1</sup> and the sample is 92% pure antineutrinos. Below 6 GeV, the region most sensitive to oscillations, the purity is 98% with the same efficiency. The wrong sign background is largest where the curvature of the muon track is difficult to measure: at low energies ( $< 2$  GeV) where the track is too short and at high energies ( $> 10$  GeV) where the track is too straight. NC's are only significant at low energies where a shower element might be mistaken for a short muon track.

## 6.2 Near Detector Spectrum

Now that a pure sample of antineutrinos has been obtained, the next step is to build up an energy spectrum from the selected events. The energy of each event is reconstructed in two parts: the muon track and the hadronic shower. For tracks that end within the detector (contained tracks), the range of the track is the best measure of its energy. If the track leaves the detector its energy is instead measured by the amount the track curves in the magnetic field. The shower energy is measured calorimetrically: it is the sum of the energy deposited into the scintillator strips for all non-track hits in the event. For more details see Section 3.4. The reconstructed track and shower energy distributions in the Near Detector are shown in Figure 6.5.

Combining the track and shower energy for each event, the full energy spectrum can be built up at the Near and Far Detectors. The selected antineutrino energy spectrum at the Near Detector in

<sup>1</sup>Calculated from simulation and taken relative to all true interactions in the fiducial volume.

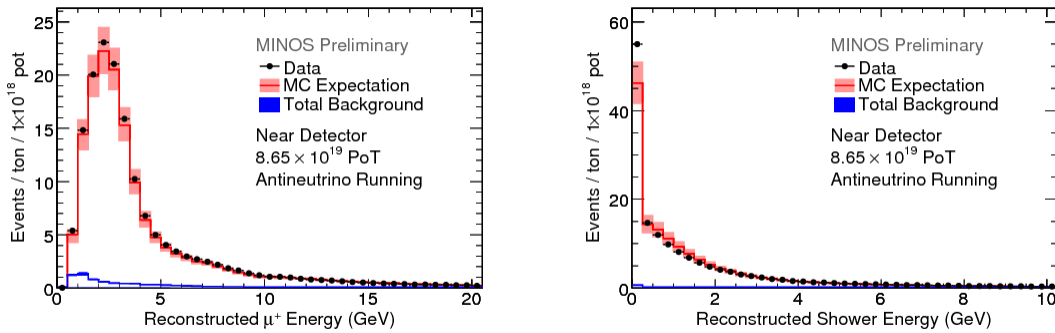


Figure 6.5: Reconstructed  $\mu^+$  track momentum (left) and shower energy (right) in the Near Detector. The red histogram represents the Monte Carlo expectation with the flux error, the blue histogram represents the total (charged and neutral current) background with the background uncertainty. Black points represent data. The track and shower energy of each event are summed to estimate the  $\bar{\nu}_\mu$  energy.




Figure 6.6: Reconstructed energy distribution of events selected as antineutrinos in the Near Detector. The red histogram represents the Monte Carlo expectation with the flux error, the blue histogram represents the total (charged and neutral current) background with the background uncertainty. Black points represent data. Thanks to MINOS's two-detector design, this flux uncertainty as well as systematic uncertainty from the antineutrino cross section will cancel in the extrapolation to the Far Detector. This figure only includes the first half of the Near Detector data – initial studies were done with a partial Near Detector data sample and it was demonstrated that the Near Detector data did not change between the first and second halves of the run period.



the antineutrino-mode beam, in data and Monte Carlo, can be seen in Figure 6.6, along with the flux uncertainty on the Monte Carlo. Thanks to MINOS's two-detector design, this flux uncertainty as well as systematic uncertainty from the antineutrino cross section will cancel in the extrapolation to the Far Detector (see Section 5.3 for more details on extrapolation).<sup>2</sup>

Note that the simulation shown has been reweighted to account for the mismodeling of hadronic interactions in the NuMI target using a fit to multiple beam running conditions. While necessary to understand the Near Detector distributions, it has little bearing on the oscillation result since flux uncertainties cancel in the extrapolation. See Section 3.1.4 for more details on the reweighting.

## 6.3 Systematics

There are a number of systematic uncertainties with the potential to affect the oscillation measurement. They fall into a few general categories: energy reconstruction, backgrounds, and extrapolation. Thanks to the two-detector design of the MINOS experiment, none of these systematic uncertainties is significant in comparison to the statistical uncertainty from the size of the antineutrino data set.

### 6.3.1 Energy Reconstruction Systematics

#### Track energy scale

Track energy as measured by range has a systematic uncertainty of 2% determined using CalDet (see Section 3.4). Track energy as measured by curvature has an additional 1% uncertainty, determined by comparing range and curvature momentum measurements for stopping tracks. These uncertainties are taken as fully correlated between the two detectors.

#### Relative shower energy scale

The relative shower energy systematics come from uncertainties in the energy calibration procedure using cosmic ray muons. Data-simulation differences in the various calibration steps are added in quadrature and give an uncertainty of 1.85% in the Near Detector and 1.05% in the Far Detector [153]. The errors are uncorrelated between the detectors.

#### Absolute shower energy scale

The absolute shower energy systematic uncertainty is taken as fully correlated between the two detectors and has two major components. The first component stems from uncertainties in the detector response to single hadrons as measured in CalDet at the CERN test beam and is 5.7% at all energies (see Section 3.4). The second

---

<sup>2</sup>I had a central role in adapting the extrapolation procedure for the antineutrino beam.

Systematic	Uncertainty
Steel Thickness	0.2%
Scintillator Thickness	0.2%
FD Live Time	0.32%
ND Fiducial Bias (z)	0.43%
ND Fiducial Bias (y)	0.14%
ND Fiducial Bias (x)	0.53%
N/F Selection Bias	1.3%

Table 6.1: Components of the Near-to-Far normalization systematic uncertainty.

component is energy-dependent and encapsulates uncertainties in hadron production and intranuclear effects. It is 8.2% at the lowest energies, dropping off to 3% above 10 GeV [154]. The final systematic has the energy-dependent form

$$\sigma_{\text{shw}} = 6.6\% + (3.5\%) \times e^{\frac{-E_{\text{shw}}}{1.44 \text{ GeV}}} \quad (6.4)$$

which is taken as fully correlated bin-to-bin.

### 6.3.2 Background Systematics

#### Neutral current background

The neutral current background systematic was evaluated by comparing data and Monte Carlo in an NC-dominated sample of events below a cut value of 0.3 in  $RoID$ , giving a systematic uncertainty of 20% on the amount of NC background.

#### Charged current $\nu_\mu$ background

As with the NC background, the  $CC\nu_\mu$  background was evaluated by comparing data and Monte Carlo in a wrong sign-enhanced sample with  $(q/p)/(\sigma_{q/p}) > 2.3$ ,<sup>3</sup> giving a systematic uncertainty of 30% on the amount of wrong-sign background.

### 6.3.3 Extrapolation Systematics

#### Near-to-Far normalization

The 1.54% normalization systematic incorporates several systematic uncertainties, all of which change the number of events expected at the two detectors per POT. It is dominated by a 1.3% uncertainty on the difference in selection efficiency in the two detectors, evaluated by hand-scanning events in both detectors in data and Monte Carlo. Also included are uncertainties in the fiducial mass of the detectors,

<sup>3</sup>This variable is used in the neutrino-mode selection and is described in Section 5.1.

the spatial uniformity of the detector acceptances, and the live time of the Far Detector. All the components are tabulated in Table 6.1.

### Cross-sections

A number of uncertainties are evaluated, both on the overall cross section and on various NEUGEN interaction model parameters. Some affect both neutrinos and antineutrinos and others are specific to antineutrinos. While the majority of the cross section uncertainty cancels between the two detectors, some residual uncertainty remains because of the spectral differences between the detectors.

### Flux modeling

The flux modeling uncertainty encapsulates a number of sources of error, including hadron production, beam optics (horn positions, currents, etc.), the position of the target, and the amount of material in the beamline. The flux errors are evaluated by moving around the fit parameters in the beam tuning fit within their uncertainties and observing the effect on the flux. Again, the majority of the errors cancel between the two detectors, but some residual uncertainty remains because the two detectors do not see identical fluxes.

## 6.3.4 Effect on the Analysis

As with the neutrino-mode analysis (see Section 5.4), the effect of each systematic uncertainty on the oscillation results is estimated using the Monte Carlo simulation. Systematic shifts are applied to Monte Carlo events to produce shifted Near and Far Detector spectra. The shifts are applied both positively and negatively, producing two sets of spectra. The total systematic uncertainty can then be examined several ways.

Figure 6.7 shows the Far Detector systematic error band constructed from all the systematic uncertainties summed in quadrature. The correlation in the systematics between the two detectors, which generally leads to cancellation, needs to be taken into account. The systematically shifted Near Detector spectrum is extrapolated to the Far Detector, producing a systematically shifted prediction. The shift in the systematically shifted Far Detector spectrum is then divided out of the shifted prediction, approximating the cancellation that occurs when fitting.

An oscillation analysis is also performed (see Section 6.5) for each systematic shift using the systematically shifted Near and Far Detector spectra as fake data. The amount the best fit moves compared to using the nominal Monte Carlo shows the size of that systematic effect on the oscillation result. The sizes of these shifts, which are approximately an order of magnitude smaller than the statistical uncertainty, can be seen in the colored lines in Figure 6.8.




Figure 6.7: Total systematic error band on the Far Detector prediction. The band is obtained by adding the effect of each individual systematic shift on the FD predicted energy spectrum in quadrature, taking into account the cancellation from extrapolation.

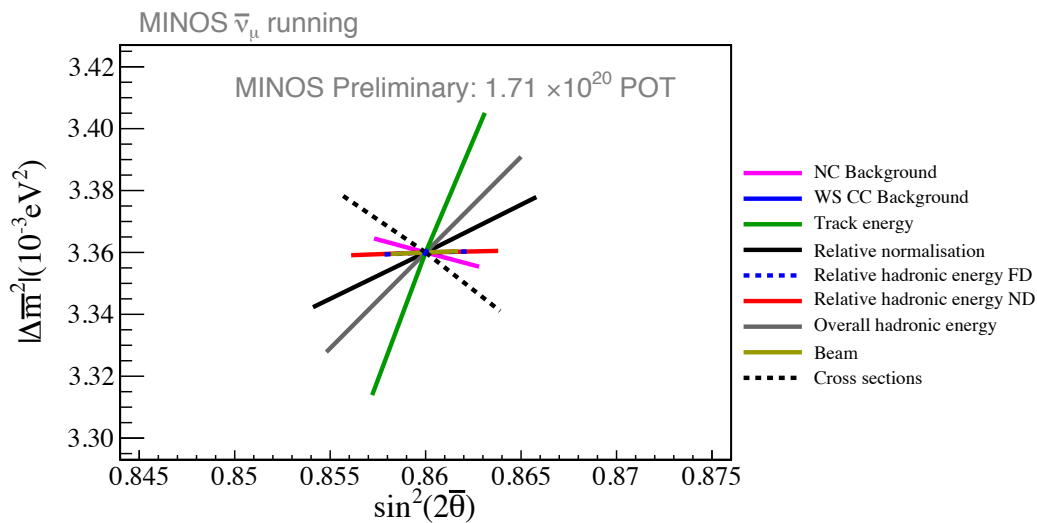


Figure 6.8: The shifts to the best fit oscillation parameters induced by the application of systematic shifts to the fake data. The cross section systematic is the sum in quadrature of all the component cross section systematics. The systematic uncertainties are approximately an order of magnitude smaller than the statistical uncertainty.

## 6.4 Far Detector Data

Figure 6.9 shows the measured Far Detector energy spectrum along with the total spectrum and background predicted without oscillations based on the Near Detector data. 156 events are predicted in the Far Detector without oscillations and 97 events were actually observed. Before fitting for oscillations, however, cross-checks are required to be sure that the Far Detector events being used are actually antineutrinos.

Figure 6.10 shows the  $q/p$  (curvature) distribution for selected Far Detector tracks of both charge signs. Antineutrinos have positive  $q/p$  and neutrinos have negative  $q/p$ . Both neutrinos and antineutrinos agree well with the (oscillated) Monte Carlo simulation, lending confidence that the antineutrinos are actually antineutrinos and the neutrinos are actually neutrinos.

There are also kinematic differences between charged current neutrinos and antineutrinos that can provide an independent check on the contents of each selected sample. Particularly, the  $y$ -distribution (fraction of event energy going into the hadronic component instead of the muon) is skewed strongly towards low- $y$  for antineutrinos while it is more uniform for neutrinos. We can see this expected behavior in the Monte Carlo and how it is fulfilled in the data in Figure 6.11.

Finally, we can also take advantage of the choice of detector magnetic field direction to distinguish neutrinos from antineutrinos. While the beam was in antineutrino mode, the two detectors had their magnetic fields set to focus positive (antineutrino) tracks inwards and defocus negative (neutrino) tracks outwards. While track vertices are distributed uniformly for both neutrinos and antineutrinos




Figure 6.9: The selected  $\bar{\nu}_\mu$  energy spectrum at the Far Detector. The red histogram represents the prediction from the Near Detector without oscillations, the gray histogram represents the expected background, and the black points represent data.

(see Figure 6.12), the track ends for antineutrinos should cluster towards the center of the detector while the neutrino track ends should cluster at the outer edges of the detector. Precisely this behavior is seen in Figure 6.13.




Figure 6.10:  $q/p$  distribution of selected Far Detector events before charge sign selection. The solid red curve represents the oscillated simulation, the dashed red curve represents the unoscillated simulation, and the black points represent data. The data agrees well with the (oscillated) simulation.

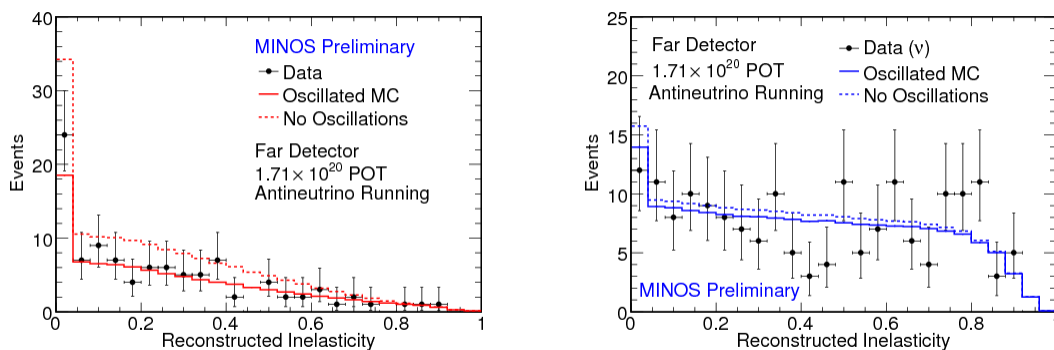


Figure 6.11: Inelasticity distribution of selected  $\bar{\nu}_\mu$  (red, left) and  $\nu_\mu$  (blue, right) interactions in the Far Detector. The solid colored histograms represents the simulation assuming the best fit oscillation parameters, the dashed histograms represent the unoscillated simulation, and the black points represent data. The data is consistent with the differing  $y$ -distributions of antineutrinos and neutrinos.

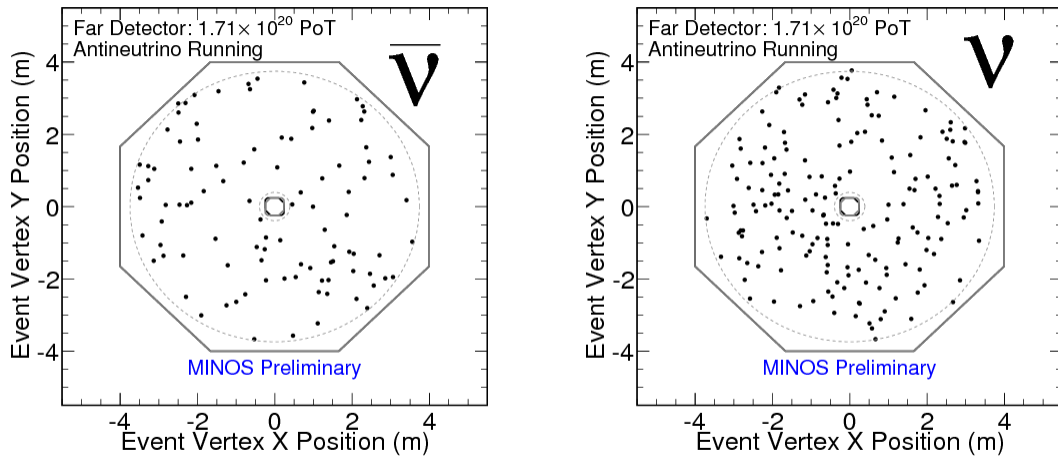


Figure 6.12: The event vertices of selected  $\bar{\nu}_\mu$ 's (left) and  $\nu_\mu$ 's (right). They are distributed uniformly throughout the fiducial volume for both  $\bar{\nu}_\mu$ 's and  $\nu_\mu$ 's.

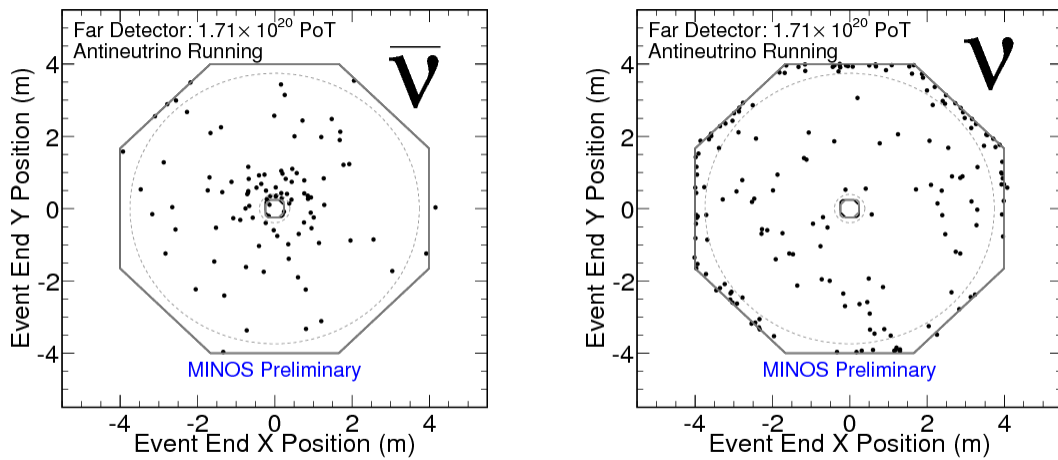


Figure 6.13: The track end positions of selected  $\bar{\nu}_\mu$ 's (left) and  $\nu_\mu$ 's (right). The  $\bar{\nu}_\mu$ 's are being focused inward by the detector's magnetic field while the  $\nu_\mu$ 's are being defocused towards the edges.

## 6.5 Oscillation Results

A fit to the oscillation parameters,  $|\Delta\bar{m}_{\text{atm}}^2|$  and  $\sin^2(2\bar{\theta}_{23})$ , was performed, assuming the two-neutrino oscillation probability

$$P(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu) = 1 - \sin^2(2\bar{\theta}_{23}) \sin^2\left(1.27\Delta\bar{m}_{\text{atm}}^2 \frac{L}{E}\right) \quad (6.5)$$

where  $L = 735$  km and  $E$  is the neutrino energy in GeV. The fit was performed by minimizing the negative log likelihood between the data and an oscillated prediction assuming Poisson statistics:

$$-2 \ln L(\mathbf{a}) = 2 \sum_{i=1}^{100} \left[ p_i(\mathbf{a}) - d_i + d_i \ln \frac{d_i}{p_i(\mathbf{a})} \right] \quad (6.6)$$

where  $p_i(\mathbf{a})$  is the predicted content in bin  $i$  as a function of the oscillation parameters  $\mathbf{a} = [\Delta\bar{m}_{\text{atm}}^2 \sin^2(2\bar{\theta}_{23})]$ ,  $d_i$  is the number of data events observed in bin  $i$ , and the index  $i$  runs over the following 100 bins: one bin from 0 to 0.5 GeV, then 0.25 GeV bins up to 20 GeV, 1 GeV bins up to 30 GeV, and 2 GeV bins up to 50 GeV.<sup>4</sup>

The oscillation parameters which best fit the data (i.e. maximize the above likelihood) are  $|\Delta\bar{m}_{\text{atm}}^2| = (3.36_{-0.40}^{+0.46}(\text{stat}) \pm 0.06(\text{syst})) \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860_{-0.12}^{+0.11}(\text{stat}) \pm 0.01(\text{syst})$ . Note that the sign of  $|\Delta\bar{m}_{\text{atm}}^2|$  cannot be determined since the probability depends on the squared sine of this quantity. The data is shown with the best fit oscillation prediction in Figure 6.14. The

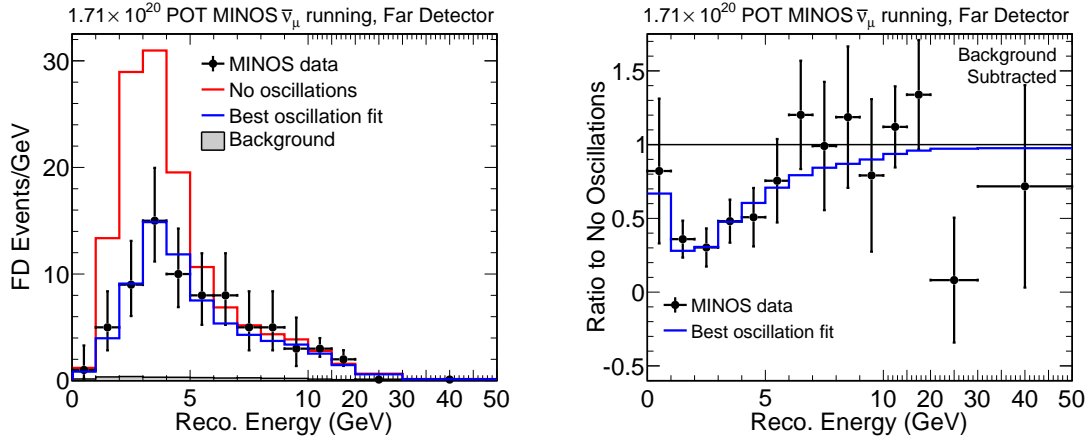


Figure 6.14: On the left is the selected  $\bar{\nu}_\mu$  energy spectrum at the Far Detector. The histograms represent the predictions from the Near Detector data without oscillations (red), with the best fit oscillation parameters  $|\Delta\bar{m}_{\text{atm}}^2| = 3.36 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860$  (blue), and for the background (filled gray). The black points represent the data. On the right are the ratios of data and best fit to the no-oscillation hypothesis with the background subtracted. The no-oscillation hypothesis is excluded at  $6.3\sigma$ .

<sup>4</sup>I performed the fit to the antineutrino data and produced the results and contours. I adapted the Feldman-Cousins procedure used here from the procedure I created for the analysis with the neutrino-mode beam.






Figure 6.15: The 68% and 90% antineutrino oscillation contours from antineutrino-mode running. The best fit is at  $|\Delta\bar{m}_{\text{atm}}^2| = 3.36 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860$ . They are compared with the 90% confidence global fit without MINOS data [76].

best fit has a  $\chi_{\text{Fit}}^2 = 94$  with 98 degrees of freedom, which gives a  $p$ -value of 0.59, but that assumes that  $\chi_{\text{Fit}}^2$  is actually distributed as a  $\chi^2$ , which is an unreliable assumption given the large number of bins relative to the number of data points. The no-oscillation scenario (shown in red in Figure 6.14) is excluded at  $6.3\sigma$  ( $-2\Delta \ln L = 43$ ). That number is calculated before systematics and Feldman-Cousins corrections have been applied, but the systematics are always small and Feldman-Cousins has little effect at large  $-2\Delta \ln L$  values.

Figure 6.15 shows the two-dimensional confidence intervals (contours) at the 68% and 90% confidence levels for the oscillation parameters. The contours were determined using the Feldman-Cousins method, using the same technique as in the neutrino-mode analysis (see Section 5.8). While the best fit of the mixing angle is 0.860, we do not exclude maximal mixing at 90% confidence. The contour is compared with a global fit for atmospheric antineutrino mixing to all neutrino experiments (except MINOS)[76]. MINOS provides a significant new constraint on the mass-splitting, but is still consistent with the global fit. Note that none of the experiments in the global fit can distinguish neutrinos from antineutrinos on an event-by-event basis – they are limited to statistical separation in combined samples.

Figure 6.16 shows the same MINOS and global fit contours compared with the neutrino-mode oscillation measurement (see Chapter 5). Though the best fit points are dramatically different, the contours do have significant overlap: the antineutrino mode contour lies directly over the lowest mass-splitting region of neutrino mode result.




Figure 6.16: The 68% and 90% antineutrino oscillation contours from antineutrino-mode running compared with the 90% antineutrino contour from neutrino-mode running. The best fit for antineutrino mode is at  $|\Delta\bar{m}_{\text{atm}}^2| = 3.36 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860$  and the best fit for neutrino mode is at  $|\Delta\bar{m}_{\text{atm}}^2| = 17.7 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.55$ . They are compared with the 90% confidence global fit without MINOS data [76]. The 90% contours are in good agreement.

### 6.5.1 Comparison to Neutrinos

A natural question is how consistent the antineutrino oscillation parameters are with the neutrino oscillation parameters already measured by MINOS:  $|\Delta m_{\text{atm}}^2| = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\theta_{23}) > 0.90$ .

Figure 6.17 shows the Far Detector spectrum and ratio, but in addition to the best fit prediction it shows, represented by the dashed line, the prediction at the oscillation parameters measured by MINOS for the neutrinos,  $|\Delta m_{\text{atm}}^2| = 2.32 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\theta_{23}) = 1$  [67]. Figure 6.18 shows the contours from the two analyses overlaid. There is some tension between the two measurements as the 90% contours are only barely overlapping. To get a combined likelihood consistent with the antineutrino analysis, we must rephrase the problem and use a Feldman-Cousins proce-

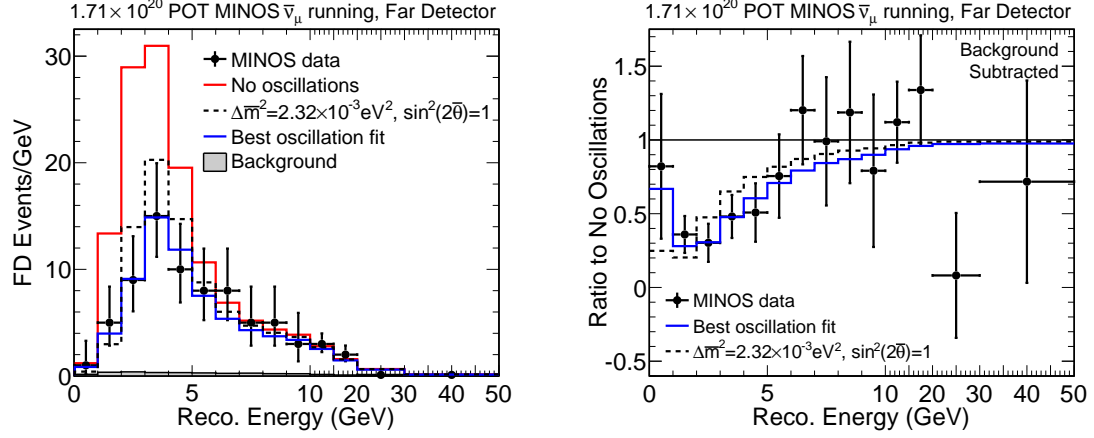


Figure 6.17: As in Figure 6.14 with the additional dashed histogram representing the prediction with the neutrino best fit parameters  $|\Delta\bar{m}_{\text{atm}}^2| = 2.32 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}) = 1.0$ .




Figure 6.18: The 68% and 90% antineutrino oscillation contours from antineutrino-mode running and the neutrino oscillation contours from Runs I-III of neutrino-mode running. The antineutrino best fit is at  $|\Delta\bar{m}_{\text{atm}}^2| = 3.36 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860$  while the neutrino best fit is at  $|\Delta m_{\text{atm}}^2| = 2.32 \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\theta_{23}) = 1$ . There is some tension between the two measurements as the 90% contours are only barely overlapping.

dure.<sup>5</sup> Instead of 2 experiments each with 2 parameters, it must be treated as 1 experiment with 4 parameters. Then, the question is how much better an oscillation fit with all 4 parameters free is than one where the  $|\Delta m_{\text{atm}}^2|$ 's and  $\sin^2(2\theta_{23})$ 's are constrained to be the same. It can be expressed more clearly with a simple variable transformation. The 4 original parameters,  $|\Delta m_{\text{atm}}^2|$ ,  $|\Delta\bar{m}_{\text{atm}}^2|$ ,  $\sin^2(2\theta_{23})$ , and  $\sin^2(2\bar{\theta}_{23})$  become only  $|\Delta m_{\text{atm}}^2|$ ,  $\sin^2(2\theta_{23})$  with two new variables defined

<sup>5</sup>I developed the Feldman-Cousins procedure for determining the significance of the difference between the neutrino and antineutrino data.




Figure 6.19: One example fake experiment. The lower right is antineutrinos and the other three are neutrinos in Runs I, II, and III. The blue histogram is the 4-parameter best fit while the red is the 2-parameter best fit. The  $-2\Delta \ln L$  between the data and each fit is calculated, and it is the difference in these  $-2\Delta \ln L$ 's that is recorded.

as  $\alpha = (|\Delta \bar{m}_{\text{atm}}^2| - |\Delta m_{\text{atm}}^2|)$  and  $\beta = (\sin^2(2\bar{\theta}_{23}) - \sin^2(2\theta_{23}))$ . Only  $\alpha$  and  $\beta$  are relevant to the consistency of the two results, so the 4-parameter likelihood can be reduced to a 2-parameter profile-likelihood,<sup>6</sup> treating  $|\Delta m_{\text{atm}}^2|$  and  $\sin^2(2\theta_{23})$  as nuisance parameters. Then, the goal is to evaluate the difference between the best fit point and  $(0, 0)$  on the  $\alpha - \beta$  profile-likelihood surface.

In the Feldman-Cousins prescription, the significance of the difference between the likelihood given the best fit parameters and the likelihood given any other set of parameters is evaluated using fake experiments produced assuming the alternate parameters. In this case, the best fit point is the 4-parameter best fit (or the two data sets fit separately) and the alternate point is the 2-parameter best fit where neutrino and antineutrino parameters are assumed to be equal ( $\alpha = 0, \beta = 0$ ). However, to generate the fake experiments, the values for the nuisance parameters,  $|\Delta m_{\text{atm}}^2|$  and  $\sin^2(2\theta_{23})$ , are needed as well. These are taken from the profile-likelihood of the data, i.e. the best fit  $|\Delta m_{\text{atm}}^2|$  and  $\sin^2(2\theta_{23})$  values of a 2-parameter fit to both data sets.

Fake experiments now need to be produced using these ‘true’ oscillation parameters. The method for producing the experiments is the same as that described in Section 5.8; it has just been extended

<sup>6</sup>The profile-likelihood is a well-known statistical technique for reducing the number of parameters in a surface. The value of the profile-likelihood at a given subset of the parameters is the minimum of a fit to the other (nuisance) parameters while holding the given parameters fixed [35].




Figure 6.20: The empirical (blue) and theoretical (red) distribution of  $-2\Delta \ln L$ 's between the 4-parameter and 2-parameter fits. Our data has a  $-2\Delta \ln L = 7.07$ , which is smaller than 1.99% of fake experiments. The uncertainty is the binomial error on the fraction of experiments above the data value.

to produce neutrino fake experiments as well as antineutrino fake experiments. The neutrino experiments are produced in three separate parts, one for each neutrino mode run period.<sup>7</sup>

Systematic uncertainties for both neutrinos and antineutrinos are incorporated, but in different ways consistent with the two analyses. The neutrino systematics are incorporated as nuisance parameters in the fitting routine for that analysis [67], so their effect is already included in the neutrino likelihood surface from which we get the neutrino part of the data  $-2\Delta \ln L$ . The antineutrino systematics, on the other hand, are incorporated into the antineutrino fake experiments.

The Feldman-Cousins prescription tells us to build up an empirical  $-2\Delta \ln L$  distribution from the  $-2\Delta \ln L$ 's between the input parameters and the best fit parameters in each of our pseudo-experiments. In this case, “best fit” refers to the 4-parameter best fit and the “input parameters” refers to the 2-parameter best fit (the input parameters being  $\alpha = 0, \beta = 0$ ). An example fake experiment with these two fits is shown in Figure 6.19.

Evaluating the significance of our data is then as simple as finding what percentage of fake experiments (which assume the input parameters) have a larger  $-2\Delta \ln L$  than the data. This distribution is shown in Figure 6.20. For our data, the measured  $-2\Delta \ln L = 7.07$ , which is smaller than 1.99% of fake experiments, so there is a 1.99% probability of observing neutrino and antineutrino data sets at least this discrepant if they are oscillating with the same underlying parameters.

<sup>7</sup>Note, the neutrino data sets are produced and fit in Runs I, II, and III, but not in separate resolution bins as done in [67]. Of course, in principle it is best if the Feldman-Cousins experiments exactly reproduce the procedure taken on the real data. In this case, it is not so crucial since the issue at hand is how “gaussian” the experiment is, rather than its sensitivity. The neutrino analysis is performed under the reasonable (for that analysis) assumption of gaussian statistics. Any variation on the method, for example resolution binning, will change the sensitivity but should not change how “gaussian” the likelihood surface is.



## Chapter 7

# Conclusion

Throughout its history, the neutrino has been at the forefront of new discoveries in physics, particularly in the breaking of symmetries that were previously believed to be absolute. According to  $\mathcal{CPT}$  symmetry,  $\nu_\mu$ 's and  $\bar{\nu}_\mu$ 's should show the same disappearance behavior due to oscillations. MINOS has made measurements of atmospheric-sector antineutrino oscillations with greater precision than was possible before, thanks to its unique ability to identify individual neutrinos and antineutrinos in a long-baseline oscillation experiment. While studying the antineutrino component of the neutrino-mode beam, we have observed no evidence of neutrino-to-antineutrino transitions and limit the fraction of oscillating  $\nu_\mu$ 's that might be transitioning to  $\bar{\nu}_\mu$ 's to less than 0.026, which others have used to set an improved limit on the effective Majorana  $\nu_\mu$  mass,  $|\langle m_{\mu\mu} \rangle|$  [117].

We have measured the  $\bar{\nu}_\mu$  oscillation parameters to be  $|\Delta\bar{m}_{\text{atm}}^2| = (3.36_{-0.40}^{+0.46}(\text{stat}) \pm 0.06(\text{syst})) \times 10^{-3} \text{ eV}^2$  and  $\sin^2(2\bar{\theta}_{23}) = 0.860_{-0.12}^{+0.11}(\text{stat}) \pm 0.01(\text{syst})$ , giving a significantly tighter constraint on the mass-splitting than the previous world combined fit. When compared with the MINOS measurement of  $\nu_\mu$  disappearance, there is a 2% probability of having two data sets that are at




Figure 7.1: The  $\bar{\nu}_\mu$  contours with different amounts of additional antineutrino-mode data. The data already accumulated is used and the additional data is simulated assuming the oscillation parameters measured in the first  $1.7 \times 10^{20}$  POT. With double the current data the  $\nu_\mu$  and  $\bar{\nu}_\mu$  contours are well separated.




Figure 7.2: The  $\bar{\nu}_\mu$  contours with different amounts of additional antineutrino-mode data. The data already accumulated is used and the additional data is simulated assuming the oscillation parameters measured for neutrinos in [67]. With double the current data the  $\nu_\mu$  and  $\bar{\nu}_\mu$  contours overlap significantly.

least as discrepant as these, assuming they have the same underlying oscillation parameters. As of now, this measurement cannot exclude the possibility of having the same neutrino and antineutrino atmospheric-sector oscillation parameters at better than 2%, which is insufficient evidence to warrant discarding  $\mathcal{CPT}$  conservation or introducing a new particle or force.

However, the measurement of  $\bar{\nu}_\mu$  disappearance was made with only six months of antineutrino-mode beam data, giving fewer than 100 events at the Far Detector. The analysis is still highly-limited by statistics and a small amount of additional running can lead to a dramatic improvement in the precision of the antineutrino measurement. Figure 7.1 shows the contours that might be possible if additional data taking in antineutrino mode gives the same best-fit oscillation parameters. With double the current data set ( $3.5 \times 10^{20}$  POT), the likelihood of seeing such discrepant distributions while assuming equal underlying oscillation parameters drops to approximately 0.1%.<sup>1</sup> If, on the other hand, the antineutrinos and neutrinos do have the same underlying oscillation parameters, the new data is likely to significantly increase the overlap of the contours, as can be seen in Figure 7.2. MINOS has accumulated an additional  $1 \times 10^{20}$  POT of antineutrino-mode data which is currently being analyzed.

Neutrino masses, and the flavor oscillations that derive from them, are the first concrete evidence of physics beyond the Standard Model. If the difference between  $\nu_\mu$  and  $\bar{\nu}_\mu$  oscillations becomes more significant with more data, it could be evidence of further physics beyond the Standard Model.  $\mathcal{CPT}$  invariance, a symmetry at the heart of quantum field theory, requires that particles and their antiparticles have the same mass. If the apparent difference between  $|\Delta m_{\text{atm}}^2|$  and  $|\Delta \bar{m}_{\text{atm}}^2|$  is taken at face value, then it would imply that the neutrino and antineutrino masses are different - a violation of  $\mathcal{CPT}$  invariance. The difference observed might also be due to as yet unobserved interactions of

<sup>1</sup>As with Figures 7.1-7.2, the already accumulated data plus simulated data with the same oscillation parameters is used to calculate this number.



---

the neutrinos and antineutrinos in the matter they travel through, which can modify the oscillation probability. Several possible models of both  $\mathcal{CPT}$  violation and new interactions have already been examined in light of the data presented in this thesis [78, 79, 80, 83, 93, 95].

If the  $|\Delta m_{\text{atm}}^2| - |\Delta \bar{m}_{\text{atm}}^2|$  difference does not become more significant after the new MINOS data has been analyzed, there will likely be no more MINOS antineutrino data. However, the study of antineutrinos will not end with MINOS. After the measurement of  $\theta_{13}$ , the next goals of the neutrino physics community are to measure the mass hierarchy and  $\mathcal{CP}$  violation. Both of these measurements require observing differences between the  $\nu_e$  and  $\bar{\nu}_e$  appearance probabilities in  $\nu_\mu$  and  $\bar{\nu}_\mu$  beams. Consequently, future long baseline experiments which hope to tackle these questions will need to accumulate significant data sets with both  $\nu_\mu$  and  $\bar{\nu}_\mu$  beams.  $\text{NO}\nu\text{A}$  and  $\text{LBNE}$  are already planning to accumulate both sets of data. Thus, in order to further the goals of the field, the study of  $\nu_\mu$  and  $\bar{\nu}_\mu$  oscillations must continue. If the apparent neutrino-antineutrino discrepancy persists, it is likely that one of these future experiments will be able to determine whether there is more new physics in the neutrino sector.



## Appendix A

# Meson and muon decay kinematics

The majority of antineutrinos in the MINOS detectors are produced from three processes:<sup>1</sup>

$$\pi^- \rightarrow \bar{\nu}_\mu + \mu^- \quad (\text{A.1})$$

$$K^- \rightarrow \bar{\nu}_\mu + \mu^- \quad (\text{A.2})$$

$$\mu^+ \rightarrow \bar{\nu}_\mu + \nu_e + e^+, \quad (\text{A.3})$$

the two-body decays of negative pions and kaons, and the decay of muons. The mesons decay primarily into muons and muon neutrinos because of the suppression by a factor of  $10^4$  of the  $\bar{\nu}_e/e^+$  channel due to the weak interaction's chiral selectivity. This is demonstrated at the end of this appendix since the calculation relies on the kinematics derived below.

### A.1 Probability that the Neutrino Reaches the Detectors

Given a particular parent  $p$  with mass  $m_p$  and velocity  $\beta_p$ , and the relative angle between the parent's direction of travel and the path to the detectors, we can calculate the probability that its neutrino daughter will reach either MINOS detector, and what energy it would have if it did so. Take  $\theta$  as the angle between the direction of motion of the parent and that of the outgoing neutrino in the lab frame and  $\theta^*$  as this same angle in the parent's rest frame.

The simplest case is two-body meson decay (Equations A.2 and A.3). Since pions and kaons are spinless, the neutrino is emitted isotropically in the parent's rest frame,<sup>2</sup>

$$\frac{dN}{d\cos\theta^*} = \frac{1}{2}. \quad (\text{A.4})$$

However, this expression must be boosted to the lab frame in order to see the distribution of neutrinos

<sup>1</sup>98.5% in neutrino mode at the Near Detector, the remaining 1.5% coming from three-body decays of  $K^-$  and  $K_L^0$ .

<sup>2</sup>Normalized so that  $\int_{-1}^1 \frac{dN}{d\cos\theta^*} d\cos\theta^* = 1$ .

seen by the detectors,  $dN/d\cos\theta$ . The transformation takes the form,

$$\frac{dN}{d\cos\theta} = \frac{dN}{d\cos\theta^*} \frac{d\cos\theta^*}{d\cos\theta}. \quad (\text{A.5})$$

Let  $p_{\parallel}$  be the component of the neutrino's momentum,  $\mathbf{p}_{\nu}$ , that is parallel to the parent's momentum  $\mathbf{p}_p$ . We will neglect the mass of the neutrino, meaning  $|p_{\nu}| = E_{\nu}$ . We can then write,

$$\cos\theta = \frac{p_{\parallel}}{E_{\nu}} \quad \cos\theta^* = \frac{p_{\parallel}^*}{E_{\nu}^*}. \quad (\text{A.6})$$

We can then perform a Lorentz boost from the parent's rest frame to the lab frame in one dimension, boosting only the parallel portion of the neutrino's momentum:

$$\begin{bmatrix} E_{\nu}^* \\ p_{\parallel}^* \end{bmatrix} = \begin{bmatrix} \gamma_p & -\gamma_p\beta_p \\ -\gamma_p\beta_p & \gamma_p \end{bmatrix} \begin{bmatrix} E_{\nu} \\ p_{\parallel} \end{bmatrix} \quad (\text{A.7})$$

$$\begin{bmatrix} E_{\nu}^* \\ p_{\parallel}^* \end{bmatrix} = \begin{bmatrix} \gamma_p(E_{\nu} - \beta_p p_{\parallel}) \\ \gamma_p(p_{\parallel} - \beta_p E_{\nu}) \end{bmatrix}. \quad (\text{A.8})$$

Then, we can substitute into the angle in the parent's rest frame to get an expression for it in terms only of lab frame variable:

$$\cos\theta^* = \frac{p_{\parallel}^*}{E_{\nu}^*} \quad (\text{A.9})$$

$$= \frac{\gamma_p(p_{\parallel} - \beta_p E_{\nu})}{\gamma_p(E_{\nu} - \beta_p p_{\parallel})} \quad (\text{A.10})$$

$$= \frac{\frac{p_{\parallel}}{E_{\nu}} - \beta_p}{1 - \beta_p \frac{p_{\parallel}}{E_{\nu}}} \quad (\text{A.11})$$

$$= \frac{\cos\theta - \beta_p}{1 - \beta_p \cos\theta}. \quad (\text{A.12})$$

Then, by taking the derivative with respect to  $\cos\theta$  we can get the transformation factor required for Equation A.5,

$$\frac{d\cos\theta^*}{d\cos\theta} = \frac{d}{d\cos\theta} \left( \frac{\cos\theta - \beta_p}{1 - \beta_p \cos\theta} \right) \quad (\text{A.13})$$

$$= \frac{\beta_p(\cos\theta - \beta_p)}{(1 - \beta_p \cos\theta)^2} + \frac{1}{1 - \beta_p \cos\theta} \quad (\text{A.14})$$

$$= \frac{1 - \beta_p^2}{(1 - \beta_p \cos\theta)^2} \quad (\text{A.15})$$

$$= \frac{1}{\gamma_p^2(1 - \beta_p \cos\theta)^2}. \quad (\text{A.16})$$

We can now calculate the probability per unit area of a neutrino reaching one of the detectors

at a distance  $z_D$  and angle  $\theta_D$  from the parent's decay point. Take some small area element on the face of the detector of area  $A$ . The probability that the parent emits a neutrino passing through this area element can be found by integrating  $dN/d\cos\theta$  over the solid angle subtended by  $A$ :

$$P = \int_A \frac{1}{2\pi} \frac{dN}{d\cos\theta} d\Omega, \quad (\text{A.17})$$

where the factor of  $\frac{1}{2\pi}$  is required to normalize out the integral over the azimuthal angle. Given the long distance between the decay point and the detector and the small size of the area  $A$ , we can take the angular decay distribution as constant over solid angle subtended by our area,

$$P = \frac{1}{2\pi} \frac{dN}{d\cos\theta} \Big|_{\cos\theta=\cos\theta_D} \Omega_A \quad (\text{A.18})$$

$$= \frac{1}{2\pi} \frac{1}{\gamma_p^2 (1 - \beta_p \cos\theta_D)^2} \frac{A}{z_D^2}. \quad (\text{A.19})$$

Thus we get a probability per unit area of reaching one of the detectors of

$$p = \frac{1}{2\pi z_D^2 \gamma_p^2 (1 - \beta_p \cos\theta_D)^2}. \quad (\text{A.20})$$

## A.2 Energy of the Neutrino that Reaches the Detectors

The next step is to calculate what energy the neutrino would have if it reached the detector. The first step is to calculate the neutrino's energy in the parent's rest frame,  $E_\nu^*$ . Let  $q_p^*$ ,  $q_\mu^*$ , and  $q_\nu^*$  be the 4-momentum of the parent, muon, and neutrino in the parent's rest frame. Conservation of energy and momentum give us

$$q_\mu^* + q_\nu^* = q_p^*. \quad (\text{A.21})$$

Rearranging and taking the dot product of each side with itself,

$$q_\mu^* \cdot q_\mu^* = (q_p^* - q_\nu^*) \cdot (q_p^* - q_\nu^*) \quad (\text{A.22})$$

$$|q_\mu^*|^2 = |q_p^*|^2 + |q_\nu^*|^2 - 2q_p^* \cdot q_\nu^* \quad (\text{A.23})$$

$$m_\mu^2 = m_p^2 + 0 - 2(E_p^* E_\nu^* - \mathbf{p}_p^* \cdot \mathbf{p}_\nu^*). \quad (\text{A.24})$$

Since we are in the parent's rest frame, its 4-momentum is

$$q_p = \begin{bmatrix} E_p^* \\ \mathbf{p}_p^* \end{bmatrix} = \begin{bmatrix} m_p \\ \mathbf{0} \end{bmatrix}. \quad (\text{A.25})$$

Using this, our expression becomes

$$m_\mu^2 = m_p^2 - 2(m_p E_\nu^* - \mathbf{0} \cdot \mathbf{p}_\nu) \quad (\text{A.26})$$

$$m_\mu^2 = m_p^2 - 2m_p E_\nu^* \quad (\text{A.27})$$

$$E_\nu^* = \frac{m_p^2 - m_\mu^2}{2m_p}. \quad (\text{A.28})$$

In order to find the energy of the neutrino at the detectors, we need to boost this energy into the lab frame. We will take advantage of the fact that the inner product of two 4-vectors is a Lorentz invariant. Taking the unstarred  $q$ 's as the lab frame 4-momenta and  $\theta$  as the angle between the parent's direction of travel and the path to the detectors,

$$q_p^* \cdot q_\nu^* = q_p \cdot q_\nu \quad (\text{A.29})$$

$$E_p^* E_\nu^* - \mathbf{p}_p^* \cdot \mathbf{p}_\nu^* = E_p E_\nu - \mathbf{p}_p \cdot \mathbf{p}_\nu \quad (\text{A.30})$$

$$m_p E_\nu^* = E_p E_\nu - |\mathbf{p}_p| E_\nu \cos \theta \quad (\text{A.31})$$

$$E_\nu = \frac{m_p E_\nu^*}{E_p - |\mathbf{p}_p| \cos \theta} \quad (\text{A.32})$$

$$E_\nu = \frac{E_\nu^*}{\frac{E_p}{m_p} - \frac{|\mathbf{p}_p|}{m_p} \cos \theta} \quad (\text{A.33})$$

$$E_\nu = \frac{E_\nu^*}{\gamma_p (1 - \beta_p \cos \theta)} \quad (\text{A.34})$$

$$E_\nu = \frac{(m_p^2 - m_\mu^2)/2m_p}{\gamma_p (1 - \beta_p \cos \theta)}. \quad (\text{A.35})$$

Thus we have an expression for the energy of the neutrino as a function of the parent's mass and velocity and the relative angle at which the neutrino must travel to reach the detectors.

### A.3 Helicity Suppression in Meson Decays

Take the pion decay shown in Equation A.2. In principle,  $\bar{\nu}_e$ 's and  $e^+$ 's could also have been produced but that channel is suppressed due to the weak interaction's chiral selectivity – only left-handed particles and right-handed antiparticles interact.

Let us look at the decay in the rest frame of the  $\pi^-$ . The outgoing antineutrino must have right-handed chirality and, since the neutrino mass is negligible, right-handed helicity. Since the pion has spin-0, conservation of angular momentum requires that the outgoing lepton and antineutrino must have opposite spins:

$$\overleftarrow{\bar{\nu}_l} \leftarrow \pi^- \rightarrow \overrightarrow{l^-}, \quad (\text{A.36})$$

meaning both the  $\bar{\nu}_l$  and  $l^-$  come out with right-handed helicity. If the lepton  $l^-$  were massless, this

process would be completely forbidden since then its helicity and chirality would always be equal and the right-handed  $l^-$  would couple to the weak force.

However, since the  $l^-$  does have mass, its right-handed helicity state has a component with left-handed chirality.<sup>3</sup> Let us take a Dirac spinor with right-handed helicity and look at its chiral projections:

$$u_{\uparrow} = P_R u_{\uparrow} + P_L u_{\uparrow} = \frac{1}{2} \left( 1 + \frac{|\mathbf{p}|}{E+m} \right) u_R + \frac{1}{2} \left( 1 - \frac{|\mathbf{p}|}{E+m} \right) u_L. \quad (\text{A.37})$$

Since only the left-handed chiral component will contribute, the matrix element for the decay must be proportional to the coefficient of the left-handed projection:

$$\mathcal{M} \propto \frac{1}{2} \left( 1 - \frac{|\mathbf{p}|}{E+m} \right). \quad (\text{A.38})$$

Let us calculate this coefficient in the pion rest frame,

$$\mathcal{M} \propto \frac{1}{2} \left( 1 - \frac{|\mathbf{p}_l^*|}{E_l^* + m_l} \right). \quad (\text{A.39})$$

From conservation of momentum and taking the neutrino mass as negligible,

$$|\mathbf{p}_l^*| = |-\mathbf{p}_\nu^*| = |\mathbf{p}_\nu^*| = E_\nu^*, \quad (\text{A.40})$$

and from conservation of energy,

$$E_l^* = E_\pi^* - E_\nu^* = m_\pi - E_\nu^*, \quad (\text{A.41})$$

and now substituting back into the matrix element:

$$\mathcal{M} \propto \frac{1}{2} \left( 1 - \frac{E_\nu^*}{m_\pi + m_l - E_\nu^*} \right). \quad (\text{A.42})$$

Substituting in the expression for  $E_\nu^*$  from Equation A.28 and taking parent  $p = \pi$ ,

$$\mathcal{M} \propto \frac{1}{2} \left( 1 - \frac{\frac{m_\pi^2 - m_l^2}{2m_\pi}}{m_\pi + m_l - \frac{m_\pi^2 - m_l^2}{2m_\pi}} \right) \quad (\text{A.43})$$

$$\propto \frac{1}{2} \left( 1 - \frac{m_\pi^2 - m_l^2}{m_\pi^2 + 2m_l m_\pi + m_l^2} \right) \quad (\text{A.44})$$

$$\propto \frac{1}{2} \left( \frac{2m_l m_\pi + 2m_l^2}{(m_\pi + m_l)^2} \right) \quad (\text{A.45})$$

$$\propto \frac{m_l}{m_\pi + m_l}. \quad (\text{A.46})$$

---

<sup>3</sup>It must since the particle moves at less than the speed of light, meaning a Lorentz-transformation can change the helicity.

So, the ratio of the branching ratios for producing a  $\bar{\nu}_e$  and a  $\bar{\nu}_\mu$  is approximately:

$$\frac{P(\pi^- \rightarrow e^- \bar{\nu}_e)}{P(\pi^- \rightarrow \mu^- \bar{\nu}_\mu)} \approx \frac{m_e^2 / (m_\pi + m_e)^2}{m_\mu^2 / (m_\pi + m_\mu)^2} \approx 1 \times 10^{-4}. \quad (\text{A.47})$$

The detailed calculation of the ratio of matrix elements and phase-space factors can be found in [155] and comes to

$$\frac{P(\pi^- \rightarrow e^- \bar{\nu}_e)}{P(\pi^- \rightarrow \mu^- \bar{\nu}_\mu)} = \frac{m_e^2 (m_\pi^2 - m_e^2)^2}{m_\mu^2 (m_\pi^2 - m_\mu^2)^2} = 1.28 \times 10^{-4}, \quad (\text{A.48})$$

which is quite close to the experimental value  $1.23 \pm 0.02 \times 10^{-4}$ .



## Appendix B

# The Decay Pipe Systematic

Decay pipe production, or downstream production, refers to the neutrinos that come from the decay of hadrons produced in the decay pipe rather than in the target. The interest in studying it is that 14% of the Near Detector antineutrino events and 6% of the Far Detector antineutrino events come from parents produced in the decay pipe. Figure B.1 shows the vertex distribution (weighted to the number of events expected in the Near Detector) color coded for the different production regions. The decay pipe is in red. While the downstream production also exists for neutrinos, it is negligible in comparison to the much larger focused sample. The trouble with the sample is two-fold. First, it does not produce the same spectrum in the Near and Far Detector (see Figure B.2) so the uncertainties in production cross-sections do not fully cancel between the two detectors. Second, the downstream production is modeled with GFluka [156], not Fluka05 [129, 130], our preferred




Figure B.1: The vertex of the production (immediate) parent of antineutrinos interacting in the Near Detector in the neutrino-mode beam. Note the log scales.




Figure B.2: The Near (color) and Far (black) Detector charged-current antineutrino spectra broken down by parent production region. The left shows the spectra from neutrino parents produced upstream and the right shows the spectra from neutrino parents produced in the decay pipe. The Near and Far spectra are scaled to the same number of events summed across the two histograms. Note that the decay pipe spectrum has much larger Near/Far differences.

and tuned hadroproduction model.<sup>1</sup> As a consequence, this so-called downstream production is a significant systematic error for the analysis of neutrino-mode antineutrinos.

In order to constrain the uncertainty on decay pipe production, the Near Detector data is used. The Near Detector Monte Carlo is first reweighted using a special set of beam tuning weights generated using the most recent version of the beam fits but without using the  $\bar{\nu}_\mu$  Near Detector spectrum (NA49 data is used as a constraint on the  $\pi^+/\pi^-$  ratio). Then, the decay pipe component is scaled up or down in order to make the total number of Monte Carlo events match the data. This scale will then be taken as the systematic uncertainty on the decay pipe. This gives a ‘worst case’ systematic uncertainty – the true systematic error will almost certainly be less than this value. The range of the integral to get the total number of events starts at 0 and ends at 13 GeV. This value was chosen because it gives the largest systematic error band, consistent with the idea of bracketing the decay pipe systematic, rather than trying to estimate an exact  $1\sigma$  uncertainty.

However, there are systematic uncertainties on the Near Detector spectrum. In order to avoid having these uncertainties mask a discrepancy that might be due to the decay pipe, the above evaluation is done with all other systematics applied (see Section 5.4), where their effects are summed in quadrature. By looking at the two extremes of the systematic shifts, the possible range of decay pipe scalings allowed by the Near Detector data plus its systematics can be determined. This gives us a systematic uncertainty on the decay pipe events of  $-100\%$  and  $+50\%$ . The effect of this systematic on the Far-over-Near ratio can be seen in Figure B.6.

<sup>1</sup>For the antineutrino-mode analysis, the beam simulation does use a consistent hadroproduction model (see F), but since antineutrino-mode antineutrinos are focused, decay pipe production is not an important systematic in that sample.

Systematic	Shift	
SKZP	+6.4%	-6.4%
Combined MaRes	+5.0%	-4.3%
NuBar DIS 2	+4.4%	-3.8%
NuBar Overall	+3.8%	-3.8%
Combined Overall	+3.5%	-3.5%
Backgrounds	+2.7%	-2.7%
Combined DIS 2	+2.3%	-2.3%
Combined MaQE	+2.0%	-1.7%
NuBar Res	+1.9%	-1.9%
NuBar QEL	+1.1%	-1.1%
Combined DIS 3	+0.8%	-0.0%

Table B.1: The systematic errors on the Near Detector, in order of their effect on the total number of Near Detector events. The systematics that shift events between bins but do not change the total number of events (e.g. energy shifts) all have very small effects (less than 1%) and are excluded from the table.




Figure B.3: The Near Detector spectrum with no systematics applied. The data is the black points, the original spectrum, with only the SKZP weights applied is the dashed line, the non-decay pipe component is in blue, the scaled decay pipe is in red, and the scaled total spectrum is in black. The shift needed in order to match the Monte Carlo to the data in this case is  $-27\%$ .




Figure B.4: The Near Detector systematic error bands, obtained from adding the effects of all systematics in quadrature. The significant systematics are broken down in Table B.1.




Figure B.5: The Near Detector spectrum with all systematics applied. The left figure has all positive systematics and the right has all negative systematics. As in Figure B.3, the data is the black points, the original spectrum, with only the beam tuning weights applied is the dashed line, the non-decay pipe component is in blue, the scaled decay pipe is in red, and the scaled total spectrum is in black. The shifts needed in order to match the Monte Carlo to the data in these cases are  $-100\%$  and  $+50\%$ .




Figure B.6: The Far-over-Near ratios for the two calculated systematic errors from Figure B.5. The effects on the Far-over-Near ratio give a sense of how significant this systematic will be for the final analysis.



## Appendix C

# Feldman-Cousins Method for the Antineutrino Analyses

### C.1 Introduction

This appendix describes how the Feldman-Cousins method [151] for determining confidence intervals was implemented for the antineutrino analysis. Using Feldman-Cousins is necessary for several reasons. First, the effects being studied exist near physical boundaries where gaussian confidence intervals are too conservative (cover more than the expected fraction of experiments for a given confidence level). For oscillations there are boundaries at  $\sin^2(2\bar{\theta}_{23}) = 0, 1$  and  $|\Delta\bar{m}_{\text{atm}}^2| = 0$  and for transitions there are boundaries at transition probabilities  $\alpha = 0, 1$ . Second, when the statistics in the spectrum are low, the oscillation formula has a degeneracy between fast oscillations with a high  $|\Delta\bar{m}_{\text{atm}}^2|$  and a low  $|\Delta\bar{m}_{\text{atm}}^2|$  with  $\sin^2(2\bar{\theta}_{23}) \approx 0.5$ . This degeneracy creates the opposite effect as the physical boundary: the gaussian confidence intervals are too optimistic (do not cover enough experiments). Third, the Feldman-Cousins prescription allows us to introduce the effects of systematics into our confidence intervals and contours.

### C.2 The Feldman-Cousins Method

The basic principle of the Feldman-Cousins method is to determine, using a Monte Carlo simulation, what  $-2\Delta \ln L$  gives the correct coverage (68%, 90%, etc.) at different values of the true parameters. For example, a  $-2\Delta \ln L = 2.71$  usually gives 90% coverage in a one-dimensional likelihood, but near a physical boundary, the fake experiments that would have had unphysical best fits “pile up” at the boundary, meaning the  $-2\Delta \ln L = 2.71$  will cover more than 90% of the experiments. Thus, to get the right coverage on the confidence interval at this point, a  $-2\Delta \ln L < 2.71$  is needed. The solution is to start with the right coverage (in this case, 90%) and work backwards to see what value of  $-2\Delta \ln L$  this corresponds to. The effect can be interpreted as the physical boundary providing

extra information, meaning the confidence interval can be drawn a little smaller. An example of this effect in the transition analysis can be seen in Figure C.1.




Figure C.1: Example  $-2\Delta \ln L$  distributions for two transition probabilities. On the left is the distribution for  $\alpha = 6.67 \times 10^{-3}$  very close to the physical boundary at 0. On the right is the distribution for  $\alpha = 0.6$ , well away from any boundaries. The black vertical lines are at  $-2\Delta \ln L = 2.71$  in both plots, the canonical 90% coverage value. The blue lines show the actual  $-2\Delta \ln L$  required to encompass 90% of fake experiments. On the right, away from the boundary this value is very close to the canonical value. On the left, however, many experiments are piled up at lower  $-2\Delta \ln L$ 's so a  $-2\Delta \ln L = 1.41$  is sufficient to cover 90%.

The prescription for using these empirical coverages is simple. Once a grid of  $-2\Delta \ln L$ 's has been produced, we can use it to determine the correct contour for our actual experiment. For the real data, we generate a  $-2\Delta \ln L$  surface from the best fit point. For each point on the surface of this experiment, if the value of  $-2\Delta \ln L$  is below that on the corresponding point on the Feldman-Cousins surface, then that point is inside the contour. Otherwise, if the point on the experimental grid is higher than the Feldman-Cousins point, then that point is outside the contour. In practice, this means that the Feldman-Cousins surface is simply subtracted from the measured surface, and all the zeros determine the contour.

### C.3 Generating Pseudo-Experiments

The crucial step in the method is how the pseudo-experiments are generated. There are two parts of the process: the first is to generate a Far Detector prediction based on the Near Detector Monte Carlo. The second is to generate the individual Far Detector fake experiments based on this prediction. Two different methods were tried for both steps: a faster, simpler method and a method that allowed systematics to be added (discussed further in Section C.4).

The first method tried for step one is to use a static reconstructed energy spectrum from Near Detector Monte Carlo as the input to the matrix method (see section 5.3). The more complex method is to loop over a reduced library of the Near Detector Monte Carlo to generate a Near Detector fake data histogram for each fake experiment. Since both methods use the same set of selected events, the histogram is produced using the same set of energies and weights. So, if systematics are not



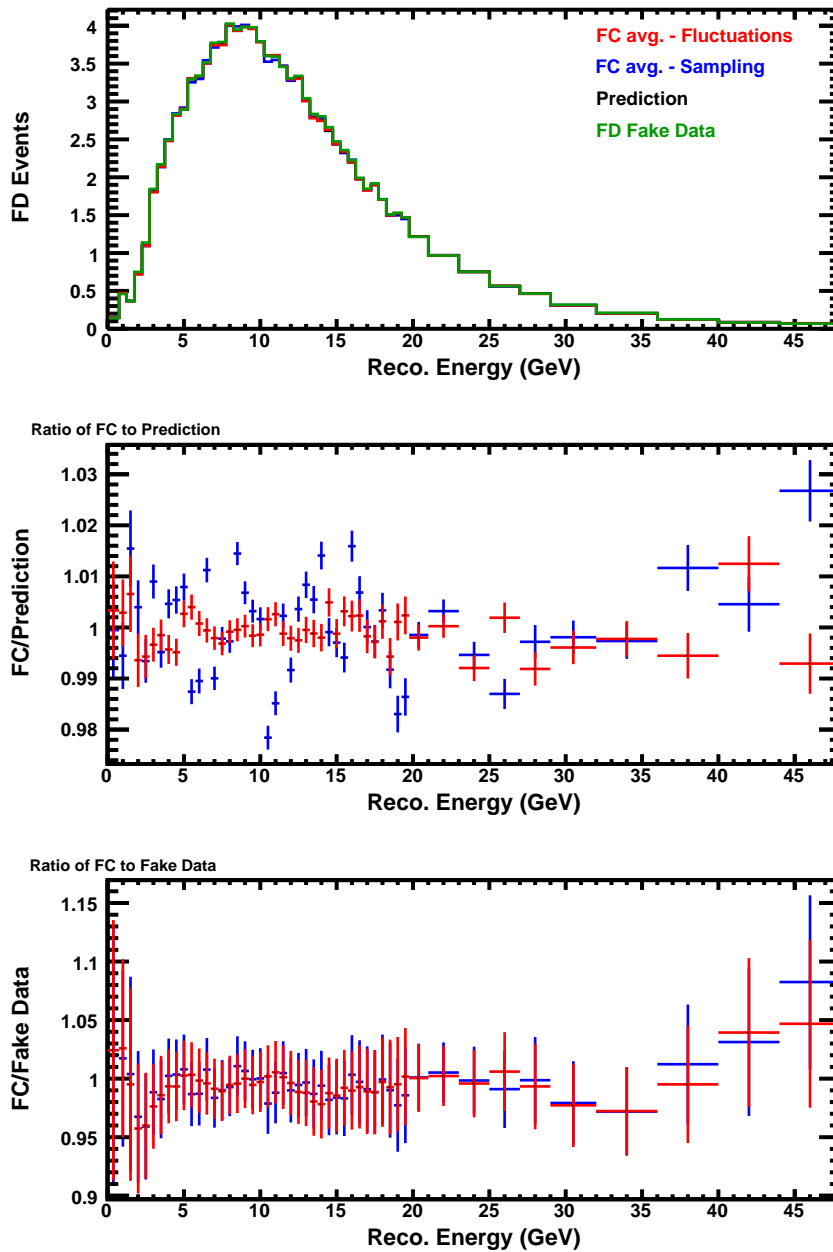


Figure C.2: The two Feldman-Cousins methods compared with each other, the Matrix Method prediction and with oscillated Far Detector fake data. Everything is consistent to within a few percent. The error bars in the middle plot (FC/Prediction) are too small because the error in the shared denominator (Prediction) are not included. This is because the statistical error on the prediction is difficult to calculate.

applied, the methods give identical results. This newly generated histogram is used as the Near Detector input data for the fitter.

Two different methods were tried for the second step and show consistent results. The ‘Fluctuations’ method is to apply Poisson-fluctuations to each bin of the predicted spectrum (that is, set each bin’s content to a random number drawn from a Poisson distribution with a mean of the bin’s contents). This fluctuated histogram would then be the pseudo-experiment the fit was performed on. The advantage of this method is that it is simple and fast to implement, and is thus not prone to errors. The disadvantage is that it is difficult to apply systematic shifts (this will be described in greater detail in C.4).

The ‘Sampling’ method is to sample actual, selected  $\bar{\nu}_\mu$  events from the Far Detector Monte Carlo. The events are separated into several different samples: signal  $\bar{\nu}_\mu$ ’s, mis-identified  $\nu_\mu$ ’s, neutral currents, and  $\nu_\tau/\bar{\nu}_\tau$ ’s. The true energy spectra<sup>1</sup> for each of these samples is predicted using the matrix method (the separated spectra are taken before being combined into a final prediction as the method usually does). Then, for each sample, the total number of events to be drawn is determined by Poisson-fluctuating the integrated number of events in the prediction. The events are drawn at random, and to ensure that the drawn events have the right energy distribution, rejection sampling [35], a well-known statistical technique, is used. The final pseudo-experiment is constructed using the reconstructed energies of the events which were chosen based on their true energies.

Figure C.2 shows the two Feldman-Cousins methods compared with both the Matrix Method prediction and the oscillated fake data. The Feldman-Cousins histograms are the average of 50,000 pseudo-experiments. Both methods are consistent with each other and are reasonably close to the prediction and fake data (a few percent variation).

## C.4 Applying Systematics

The principle for applying systematics for Feldman-Cousins is simple: in each fake experiment, pick a random value for each systematic uncertainty and incorporate these shifts into the pseudo-experiment. For systematics like normalization, this is very simple no matter how the pseudo-experiments are produced – simply scale the prediction before fluctuating or sampling (essentially, scaling the number of events drawn). For others it can be more complicated – they are described in Table C.1.

Since the normalization systematic was easy to implement for both methods, it was used as a cross-check to make sure the systematics were being applied properly. The resulting surfaces (four in all, new and old method with and without normalization), are shown in Figure C.3 (Transitions were used as a test since the 1-dimensional surfaces are easier to compare). As you can see, the

---

<sup>1</sup>Reconstructed energy is used for the neutral currents since this defines their behavior in the detectors and the true energy can vary wildly depending on how much energy is taken away by the exiting neutrino.

Systematic	Size	ND Method	FD Method
Normalization	4%	-	Scale the number of events drawn for all samples.
Backgrounds	50%	Scale the weight of each background event.	Scale the number of events drawn for the background samples.
Track Energy Range	2%	Modify the energy component before adding to the reconstructed energy histograms.	FD Only FD Only
Overall Shower Energy	10%		
Track Energy Curvature	4%		
Relative Shower Energy	3.3%		
NuMuBar Cross-Section	$1\sigma$	Scale the weight of each event in true energy as shown in Figure C.4.	Use the weights in true energy shown in Figure C.4 to scale acceptance.
SKZP	$1\sigma$	Scale the weight of each event using its associated error.	Use each event's error to scale its acceptance (described below).
Decay Pipe	37%	Scale the weight of decay pipe events.	Scale the acceptance of decay pipe events.

Table C.1: Table of systematics and how they are implemented.

systematics had a clear affect, and the new and old methods are consistent with each other with and without systematics.

Systematics that would be applied by changing the weights of events in a high-statistics histogram (like the Near Detector) need some care in application to the Far Detector where only a small number of unweighted events are used. It is easier when the events being weighted can be treated as a group (i.e. the background scales); the total number of those events can be scaled. When the scaling needs to happen on an event-by-event basis, care needs to be taken to account for the two separate effects of the systematic: changing the likelihood of having that particular event *and* changing the total number of events. If an event is drawn that needs to be weighted, two rejection probabilities are calculated: one unmodified, and one systematically shifted up or down. The event is rejected based on the shifted probability. If the event would have been accepted otherwise but is rejected because of the systematic, the total number of events drawn for this pseudo-experiment is reduced by one. Conversely, if an event that would have been rejected is accepted because of this systematic, the number of events drawn is increased by one.

For the decay pipe, there was an option to either do the acceptance scaling event-by-event or treat it as a separate sample like the backgrounds. While the ‘Separate Sample’ method was not chosen for production because it required modification of much of the matrix method chain and thus introduced numerous possibilities for small bugs, it was implemented temporarily as a cross check. Figure C.5 shows the predictions from both decay pipe methods compared to each other. It is clear that they give basically identical spectra.

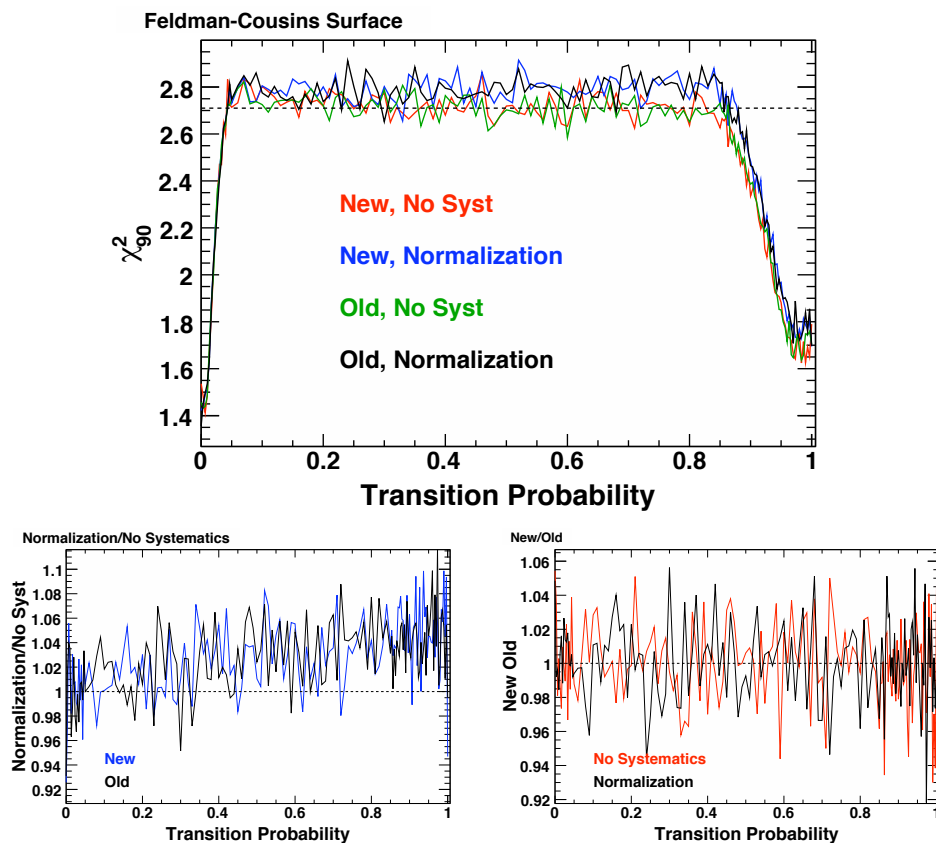


Figure C.3: A comparison of the two Feldman-Cousins methods (old - fluctuated histogram, new - event drawing) both with and without a normalization systematic applied. As you can see in the lower left ratio plot (blue), applying the normalization systematic has an observable effect. And, looking at both lower ratio plots, it is clear that the new and old Feldman-Cousins methods are giving consistent results.




Figure C.4: The curve used to parameterize the antineutrino cross section uncertainty as a function of energy. It is compared to the  $1\sigma$  errors on the combined neutrino and antineutrino cross-sections components as well as the antineutrino cross section ratio components.




Figure C.5: The two methods for applying the decay pipe systematic compared to each other. Here, the decay pipe systematic has been forced to be exactly +37%, not varied, so its effects can be investigated. They clearly give results consistent to within 2%.

## C.5 Feldman-Cousins Surfaces

### C.5.1 Oscillations

The 90% Feldman-Cousins surface for oscillations is shown in Figure C.6. Near the physical boundaries the empirically-corrected sensitivity is better than the sensitivity assuming gaussian errors, as expected. The loss of power from the degeneracy between fast oscillations and  $\sin^2(2\bar{\theta}_{23}) \approx 0.5$  is also apparent in both areas of parameter space.

### C.5.2 Transitions

The final transitions Feldman-Cousins surface with systematics and 50,000 experiments per point is shown in Figure C.7. We see the expected drop-off near the physical boundaries. You can also see that across most of the range, the FC curves sit a little bit above their nominal values (dotted black lines). This increase is from the inclusion of systematics and will lead to slightly wider confidence intervals.

Figure C.8 shows what range of confidence intervals we might expect using these surfaces on experiments with no transitions happening. Both the lower and upper limits are shown. 90.6% of the lower limits are zero meaning the coverage is nearly perfect (we would expect 90% of the intervals to contain the true value). The upper limits of the 10,000 experiments range from  $\alpha = 1.1\%$  to 27%, though the typical values (middle quartiles) range from 3.9% to 9.5% with a median of 6.5%.




Figure C.6: Feldman-Cousins 90% coverage grid with systematic shifts included in the individual experiments. The color on the coverage grid represents the  $-2\Delta \ln L$  required to cover 90% of experiments. If the experiment were perfect gaussian (and without systematics), the grid would be a uniform 4.61 (light blue) everywhere.




Figure C.7: The red, blue and green lines show the  $-2\Delta \ln L$  value needed to cover 68%, 90% and 99% of mock experiments as a function of the true transition probability. The dashed lines show the statistical-only case whereas the solid lines show the case where systematics were also included in each of the mock experiments. The horizontal dot-dash lines at 1.0, 2.71, 6.67 are the Gaussian expectation. The effect of the physical boundaries is clear: the curves drop below the gaussian line near  $\alpha = 0$  and  $\alpha = 1$  and without systematics the curves reproduce the gaussian expectation everywhere else. Adding systematics always increases  $-2\Delta \ln L$  value required, leading to widened confidence intervals as expected.




Figure C.8: Above is pictured the upper (blue) and lower (red) limits of the 90% confidence intervals of 10,000 fake experiments with no transitions. The true value ( $\alpha = 0$ ) is contained 90.2% of the time (the large red peak at zero). The box and whisker plot of the Upper Limits gives, at the vertical bars from from left to right, the minimum, 1<sup>st</sup> quartile, median, 3<sup>rd</sup> quartile, and maximum.



## Appendix D

# Detector Edge Study

### D.1 Introduction

A discrepancy was observed between the data and the simulation in where the west edge of the Near Detector appeared to be. The first evidence of the problem can be seen in Figure D.1. In it, the end points of positive tracks are plotted in  $x$ . Positive ( $\bar{\nu}_\mu$ ) tracks are used because they are defocused towards the outside edges of the detector. By placing a selection  $-0.82 < y < 0.82$ , the vertical west edge of the detector was selected (see inset in Figure D.1). When data (black) and Monte Carlo (red) are compared, a large discrepancy is apparent: the data extends further out in  $x$  than the




Figure D.1: The track end point in  $x$  in meters for positive tracks in the Near Detector. Only end points falling in the colored region of the inset are included, selecting the vertical west edge of the detector. Simulation is shown in red and data in black. The data extends several centimeters past the end of the simulation.

Monte Carlo.

## D.2 Plane Types

The first question to answer is what planes are being affected by this problem. For reference, Figure D.2 shows a drawing of how the Near Detector planes are constructed from modules. Figure D.3 contains the same events as Figure D.1, but separated by plane type: partial and full, U and V. It is clear from the figure that the discrepancy is only in the full planes, not the partial planes, though it is not isolated to either orientation. Figure D.4 shows the track end positions separately for the calorimeter and the spectrometer, but now looking at the two diagonal segments on the west side. The discrepancy is apparent in these angled sections, but it is not isolated to either the calorimeter or the spectrometer.

## D.3 U and V

$x$  is not a fundamental property of the geometry of the detector: it is naturally oriented in U and V. Again, it is simpler to look at flat edges, so a U cut was applied to the V plots and a V cut applied to the U plots so only the flat edges are selected. The analyzed regions are shown in insets in Figure D.5. This figure also shows that the effect can be isolated by looking only at track end positions along the length of the strips. Note, the plane only measures position based on which strip a hit occurs in and cannot measure position along a strip. This means the effect only occurs in the

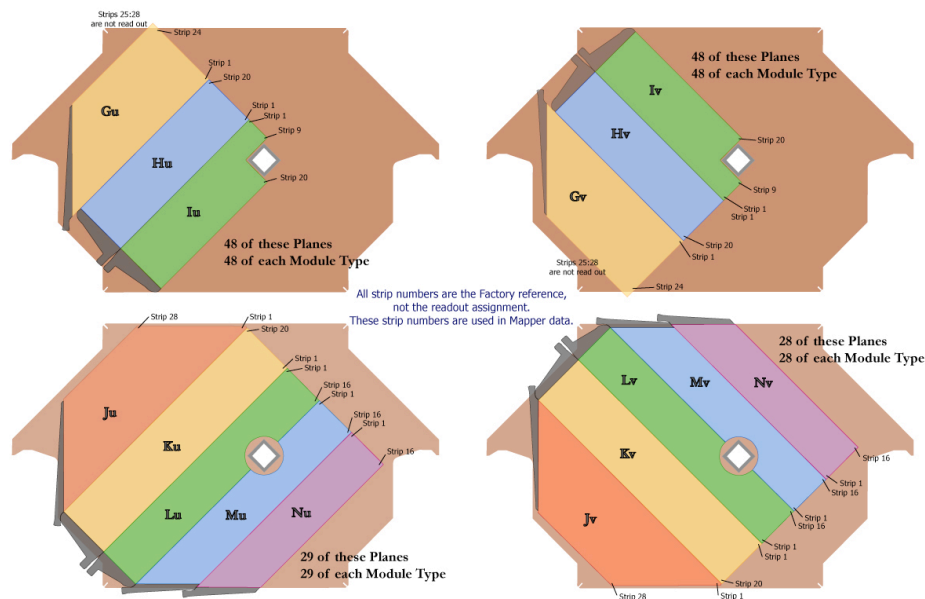


Figure D.2: The types of planes in the Near Detector with their component modules.

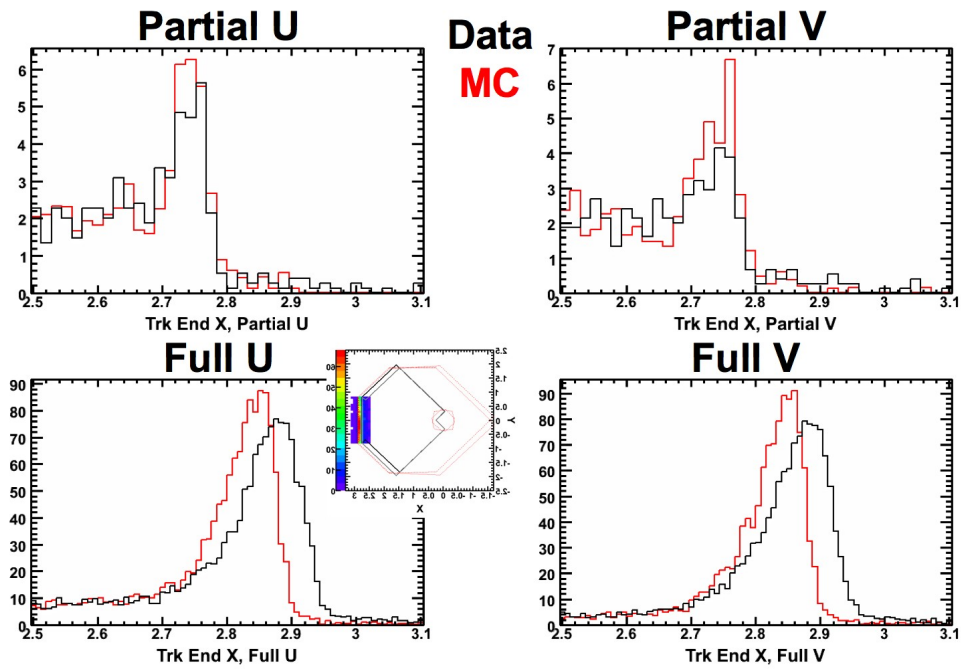


Figure D.3: These plots show Figure D.1 broken up into full and partial U and V planes. As you can see, the discrepancy appears equally in Full U and Full V planes but not in either Partial U or Partial V planes.

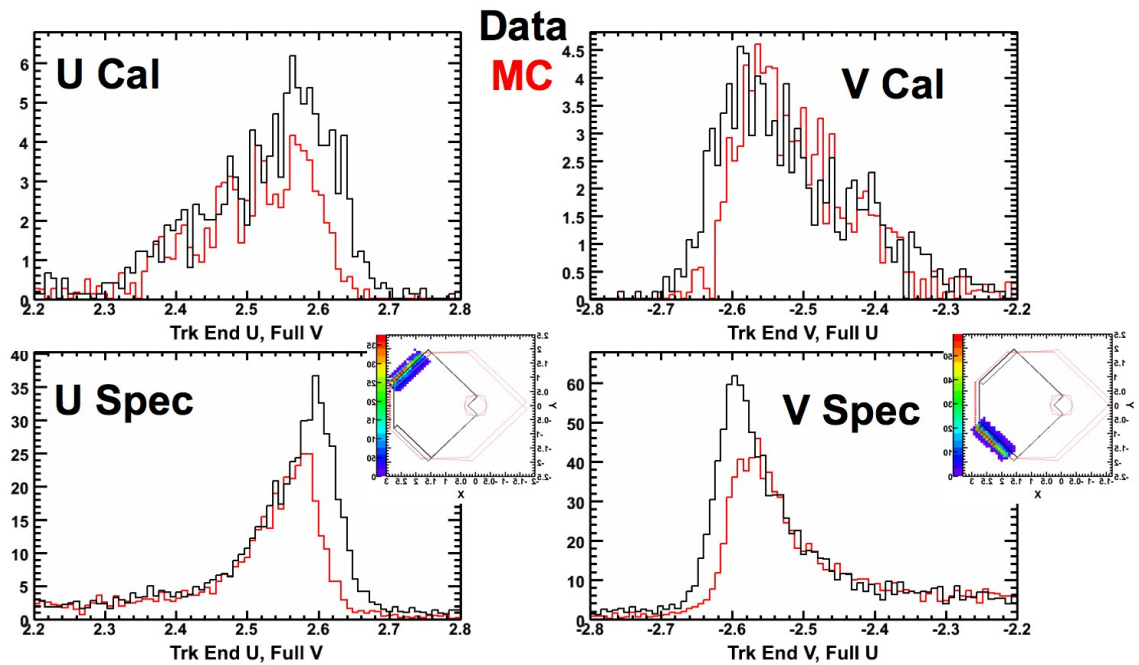


Figure D.4: Here we are looking at V and U planes (from the colored segments of the insets, again looking at the flat edges) broken up into calorimeter and spectrometer. The discrepancy appears in both  $z$  segments, so it is not an effect of the lack of alignment in the spectrometer.

direction not directly measured by the plane. In the dimension the plane actually measures, there is excellent agreement between data and simulation as alignment has been performed on the data in this direction.

There is a second effect apparent in Figure D.5 that strongly suggests a mechanism that might produce the difference in track end position observed. Not only do data hits extend further out in the direction parallel to the strips, there are also *more* data tracks ending on those planes. The effect is striking since the data and simulation follow each other while they rise, the Monte Carlo just hits a cut-off earlier. The combination of these effects suggests that the strips in the simulation do not extend as far out relative the perpendicular plane as they do in the data. Figure D.6 shows an illustration of what the effect might look like.

However, there are several underlying effects that could produce this result. For example, the strips could be too short in the Monte Carlo, or the, strips could be offset along their length, or the plane itself could be offset. Additionally, the extra length could be an effect of the reconstruction (reconstructing points outside the physical detector).




Figure D.5: Viewed in U and V directions (as opposed to the linear combination of the two) features of the discrepancy become apparent. When position is measured along a plane's axis (perpendicular to the direction of the strips) the Data and Monte Carlo agree well on where each strip is located, and thus where the detector edge is. However, when position is measured perpendicular to a plane's axis (along the length of the strip), the discrepancy from Figure D.1 is clear. Note also that there is more Monte Carlo in the top plots and more Data in the bottom plots. Not only is there more data in the bottom plots, but the data and Monte Carlo curves follow each other and the Monte Carlo just falls off sooner than the data.

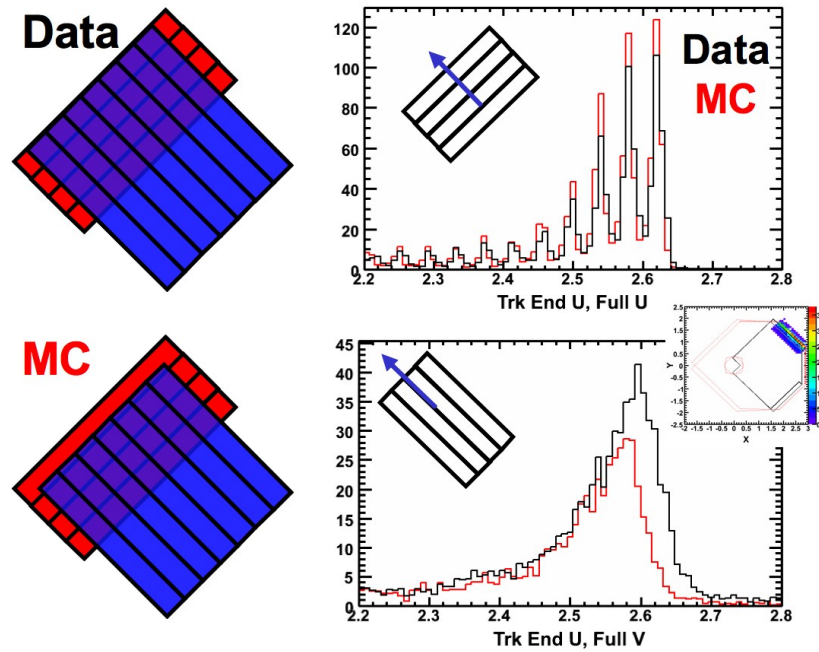


Figure D.6: A proposed explanation of the effects seen in Figure D.5. If the Monte Carlo strips did not quite extend to the edge of plane but the real strips did, it would produce the effects seen. First, the data edge would be further out than the Monte Carlo edge. Also, there would be more data hitting perpendicular strips since there is more strip at that edge than in Monte Carlo. Essentially, the Data particle, after passing through a parallel strip would go on to hit one more perpendicular strip while the Monte Carlo would appear to just stop in the parallel strip since there is no perpendicular Monte Carlo strip to hit that far out.

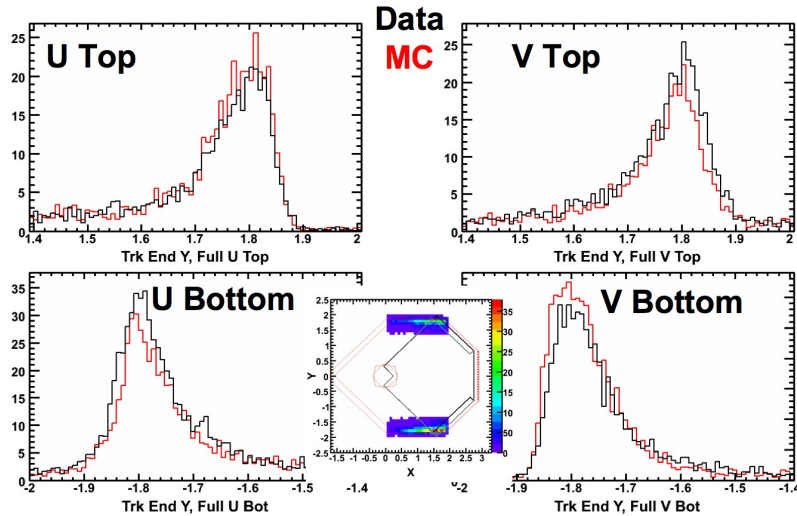


Figure D.7: These are the top and bottom of the detector in  $y$  with a selection on  $x$  as shown in the inset to get the flat segments. Referring to Figure D.2, U Top is the reflective end of the Ju module, V Top is the readout end of Mv and Nv modules, U bottom is the readout end of the Mu and Nu modules, and V bottom is the reflective end of the Jv module. While there is no apparent discrepancy in the  $y$  endpoint, note that there appears to be a small excess of data in the readout ends (V Top and U Bottom) and a small deficit of data in the reflective ends (U Top and V Bottom).

Coordinate	Plane	Modules
U	Full U	Ju
V	Full V	Jv
U	Full V	Kv and Lv
V	Full U	Ku and Lu

Table D.1: The correspondence between plots in Figure D.5 and the modules in Figure D.2.

One way to disambiguate an offset in strip position and an overall change in strip length is by looking at the other end of the strips. Unfortunately there is no data for the far end of the K and L modules, but there is data at the top and bottom of the detector which can tell us about the J modules. The plots can be seen in Figure D.7. There is no clear offset in the top edge of the detector.

## D.4 Physical Measurements

There is an example Near Detector plane in the MINOS surface building at Fermilab. Some measurements were taken of the lengths of modules on that plane to compare with what we are seeing in the data. The length of a full Lu strip was measured to be 376 cm and the distance to the edge from the center of the coil hole (the Near Detector origin) was 262 cm (see Figure D.8). The 376 cm is tolerably close to the value for strip length in the simulated geometry of 375.1 cm. To compare the distance from the origin measurements, look at Figure D.9.




Figure D.8: Measurements taken of the example plane in the MINOS surface building.

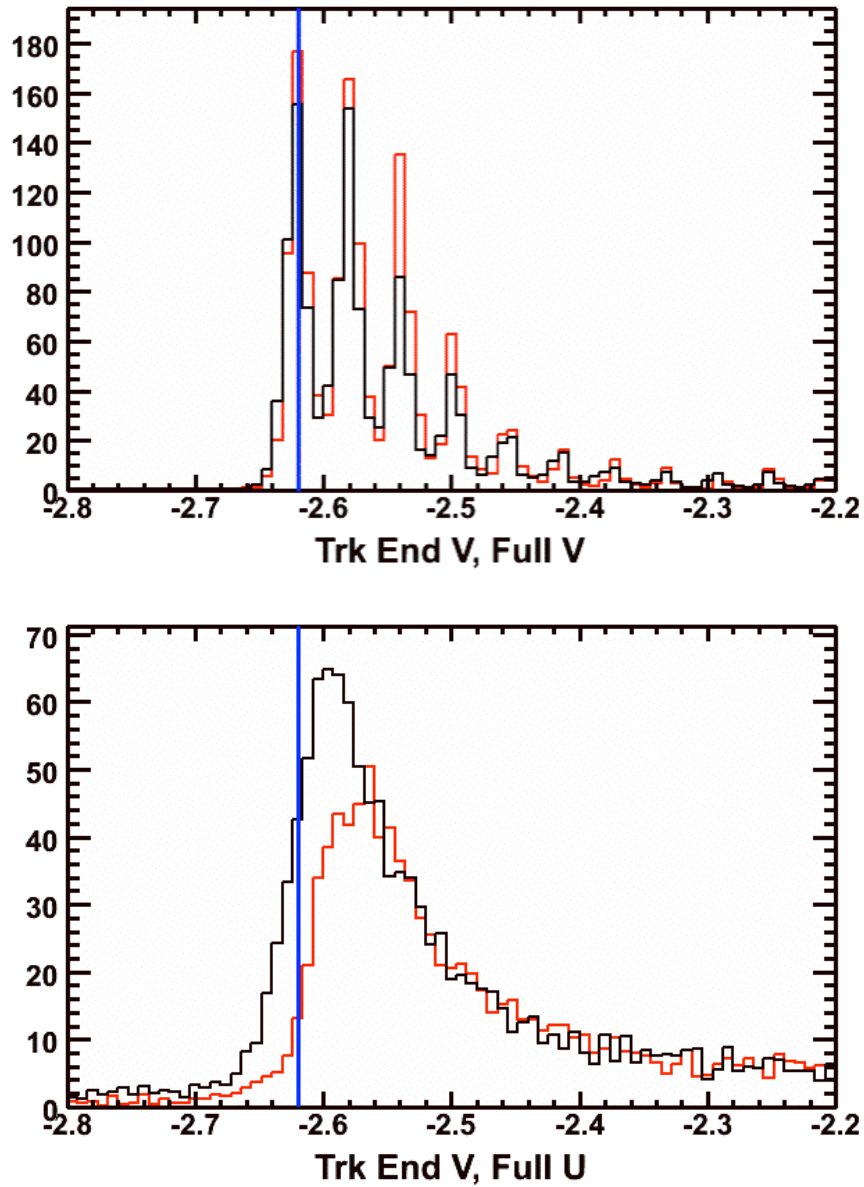


Figure D.9: The plots are identical to those in Figure D.5. The blue lines drawn on the plots mark the measured end of the plane from the surface building. While this measured point lines up with the peak of the parallel strips, it doesn't clearly agree with either the peak or edge of either Data or Monte Carlo for perpendicular strips.

## D.5 Monte Carlo “Measurements”

Some “measurements” were also taken by looking at different variables in the Monte Carlo. The goal of these studies was to determine if the detector edge discrepancy existed at the truth level or if it was only introduced during reconstruction. Two different methods were looked at, both of which show that the Monte Carlo truth agrees with the Monte Carlo reconstruction and they both disagree with the data.

First, true vertices in the scintillator were examined. Since the material of the interaction isn’t stored, only the  $Z$  of the interacting nucleon, the scintillator could not be selected perfectly. Instead, interactions in carbon ( $Z=6$ ) were selected which gives mostly scintillator and a little bit of steel Figure D.10. These results are in Figure D.11.

The second method used was to look at the DigiScintHits arrays from the MCNTuple tree. Since these record the true positions of energy depositions in the detector, they also probe the extent of the detector as seen by the Monte Carlo at the truth level. The results in Figure D.12 clearly show that the Monte Carlo truth is consistent with the Monte Carlo reconstruction on the position of the edge of the detector.




Figure D.10: True vertices whose  $Z=6$ , meaning carbon. This selects primarily interactions in the scintillator (red boxes), but also accepts a few interactions in the steel (blue boxes). Note that the steel extends well out past the end of the scintillator.

## D.6 Conclusions

The large discrepancy in end positions of tracks along the west side of the detector comes from the fact that the simulated strips do not extend as far out relative to perpendicular strips as the data strips do. The best hypothesis for how this occurred is that the simulation is based on the detector






Figure D.11: The true vertex positions in the Monte Carlo. There appears to be a fall off at 2.9 m corresponding to the edge in Monte Carlo consistent with earlier plots (Figure D.1). The few events that extend out in x past 2.9 m are the still interactions (Figure D.10).




Figure D.12: Again the plot from Figure D.1 is shown. Now, overlaid on that plot is the blue shaded region which shows the locations of DigiScintHits. These are the true locations of all energy deposition in the detector. The Monte Carlo appears to be consistent with itself in saying the detector ends at 2.9 m.

as designed, rather than as-built. Speculating a little more, here is a possible mechanism. There was a guide that went through the coil hole cut in each plane for positioning the scintillator on the steel planes they were attached to, but this guide was not the full size of the coil hole. Maybe as each scintillator plane was put in place, it was not centered on the guide but rather pushed all the way to one side, so the guide and coil hole socketed into place. This would introduce a shift in the  $x$ -direction of the planes, but a particular mechanism is difficult to prove.

Whatever the mechanism, the effect makes no difference to the analyses. First, the shift is only apparent along the direction that is *not* used to measure the track position in the planes where it occurs, meaning the effect is not producing shifts in the positions of lots of hits. In the direction where the plane does give position information, everything has been aligned and agrees well between data and simulation. And whether a track ends in strips parallel or perpendicular to the detector edge, it is still going to be classified as an exiting track so the potential extra hit would not effect the measurement of the track's energy.

## Appendix E

# Validation of the Reconstruction Upgrade

Between the neutrino-mode and antineutrino-mode analyses, the reconstruction package was upgraded from version “Cedar” to version “Dogwood.” One of the major upgrades in Dogwood was removing the “range bias” from the track momentum fitter. Previously, curvature measurements were biased by the range measurement because the track fitting algorithm would use some range information when determining track momentum from curvature (for more details see Section E.2). This feature, while it improved the track momentum resolution, made it more difficult to use the comparison between range and curvature to determine the systematic error on the muon momentum. In order to help in validation studies, both the range-biased and the unbiased fitters were run in

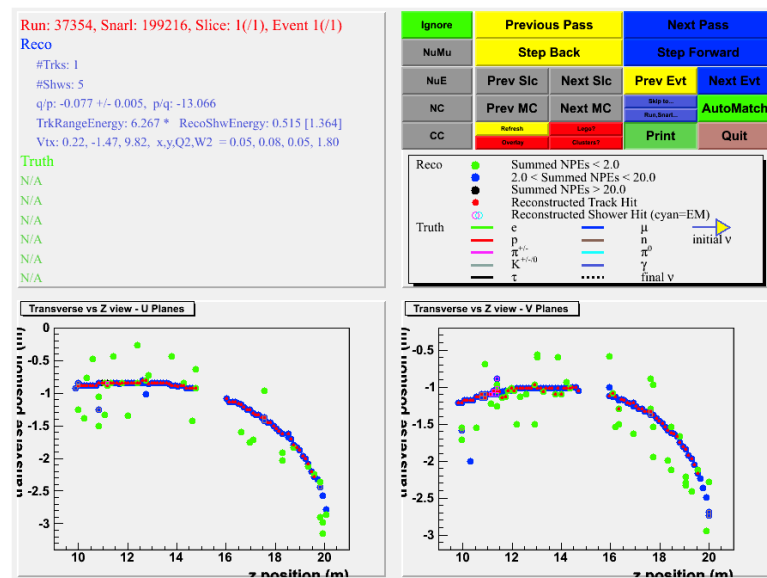


Figure E.1: Event view of a catastrophic charge mis-id in Dogwood from the Far Detector data. While the event is clearly curving away from the coil hole in both  $U$  and  $V$  views, it is reconstructed with  $p/q = -13$  GeV.

Dogwood so their output could be compared.

## E.1 The Charge Sign Problem

The problem was this: in Dogwood1, long, clean tracks were being reconstructed with the wrong charge. The problem was first noticed in the Far Detector data with the event shown in Figure E.1. This event is clearly positive to the eye and to the Cedar reconstruction ( $p/q = 6.41$  GeV), but Dogwood assigned it  $p/q = -13$  GeV. The problem occurred in about 2% of events with tracks in the Far Detector – the ratio of positive to negative events in Dogwood became 82% : 18% instead of 80% : 20% as we had in Cedar. This led to a 2% decrease in the CC signal and a 28% increase in the NuMuBar background. The tracks affected were typically long, with high  $q/p$  significance and were easily identified by alternative charge-sign selectors like relative angle. Note that this problem only affected the new unbiased fitter, not the legacy biased fitter (see Figure E.2).

## E.2 The Charge Sign Fix

After some detailed investigation of the failed events by Jim Musser, the cause was eventually determined. Part of the change in going to an unbiased fitter was removing a 250 MeV cut-off in energy for a hit to be included in the fit. That is, as the fitter moved along a track if it decided that a track's energy had dropped below this cut-off energy, it would stop trying to fit any further hits. However, sometimes the last hit or two on the track (which would previously have been cut) would unduly influence the Kalman filter and flip the charge of the fitted track so strongly that the backward pass along the track could not recover the correct charge. This was at least partially caused by a bug in the  $dE/dx$  formula used by the Kalman fitter which only became significant at low energy.

This fix became known as Dogwood2. However, while it fixed the charge-id problem, it introduced another significant problem with track length. You can see an event-by-event comparison of track lengths in Dogwood2 and Dogwood1 in Figure E.3. In the Dogwood2-Dogwood1 comparison (shown in blue) there are clearly two peaks – a narrower one I would call the “expected disagreement” and the large peak with 10 or more planes missing in Dogwood2.

To understand where this problem comes from some more details on the track fitting algorithm are necessary. The track fitter begins with a *found* track that an early rough reconstruction stage identifies as a candidate. The Kalman fitter then travels back and forth twice along this track, deciding as it goes whether or not the hits in the found track actually belong on the *fitted* track. In principle, and frequently in practice, the fitter does not use all of the hits on the track – this is in fact the essence of removing the range bias. What the range bias did was force the Kalman fitter

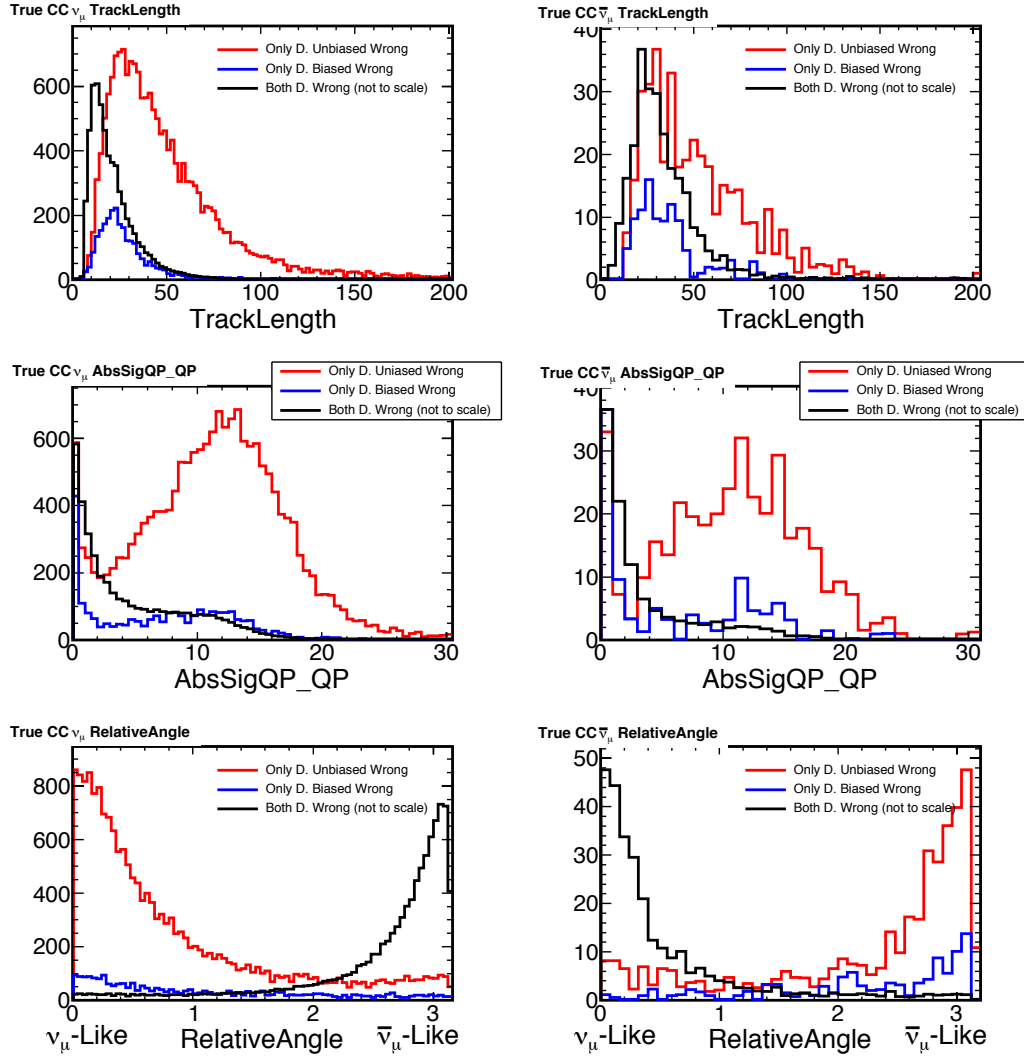


Figure E.2: Shown in color above are the track length (top row),  $|q/p|/\sigma_{q/p}$  (called ‘AbsSigQP\_QP’, middle row), and relative angle (bottom row) distributions for events where the new unbiased and ‘Cedar’-like biased fitters disagree about the charge of the track. Also, not to scale, are the distributions for events where both fitters get the charge sign wrong (this is how these distributions look for ‘typical’ charge mis-ids). They are divided into true (Monte Carlo)  $CC\nu_\mu$  (left column) and  $CC\bar{\nu}_\mu$  (right column). When the two fitters disagree about the charge sign of the track (red and blue histograms), the unbiased fitter (red) is wrong far more often. Beyond that, the tracks the biased fitter is getting wrong tend to have longer track lengths, higher  $|q/p|/\sigma_{q/p}$ , and more signal-like relative angle than typical mis-ids. That is, the biased fitter is getting the charge wrong on good-looking tracks.

to keep adding to its estimated track energy until it made it to the end of the found track. However, none of this *should* affect the track momentum from range since that is based on the found track, not the fitted track.

But clearly the momentum from range was somehow being affected. After some more investigation, Jim Musser was able to identify the problem. Sometimes the Kalman fitter decides that the last point on the found track is not actually part of the track. In this situation, the found track should be unaffected, or at worst just lose one hit. However, due to a quirk in how the found and fitted tracks were stored, when the last hit was dropped it caused the found track to only contain the hits in the fitted track. After a change in how track hits were stored, this problem was also eliminated.

This fix became Dogwood3; you can see its effect in the red curve on Figure E.3. Now in Dogwood3 rather than having a lot of tracks that were way too short, we in fact had tracks coming out longer than they had been in Dogwood1. When we looked specifically at these longer tracks, it turned out the energy resolution was better than it was before, meaning the longer length is in fact the correct one. Apparently this track length bug had existed previously at a much lower level and simply wasn't noticed until Dogwood2 brought it to prominence. Previously the effect would have been much smaller since the range bias caused fits to almost always use all of the found track.

### E.3 Dogwood3

We have seen that Dogwood3 fixes the range problem that Dogwood2 introduced, but have we fixed the original Dogwood1 charge sign problem? The answer can be found in Figure E.4. These figures show a direct comparison between Monte Carlo events matched event-by-event between Cedar and Dogwood (unlike Figure E.2 which compares two different Dogwood fitters). You can see in both the Near and Far Detectors, in both track length and  $(q/p) / \sigma_{q/p}$ , that the red curve (Dogwood1 gets the charge wrong and Cedar gets the charge right) has an extra component at longer lengths and higher confidence than its converse shown in black. This is the charge-sign problem we want to fix. By comparison, the blue curve (the Dogwood3 equivalent of the red curve) is in good agreement both with its converse (shown in green) and the Dogwood1 black curve.

### E.4 Ancillary Problems

There was an additional tracking bug that was discovered as part of this process. It turned out that the legacy biased fitter was ignoring all the hits in the Near Detector spectrometer. This, as might be expected, caused the fitter performance to degrade drastically the more time a track spent in the spectrometer. This bug is also now fixed in Dogwood3.

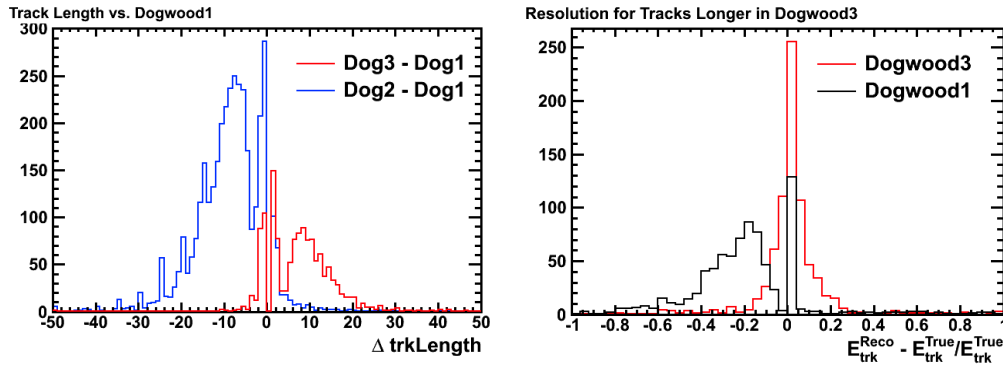


Figure E.3: On the left in blue is the event-by-event difference in track length compared between two identical MC samples reconstructed in Dogwood2 and Dogwood1. In red is the same comparison, but between Dogwood3 and Dogwood1. On the right is the energy resolution in Dogwood1 (black) and Dogwood3 (red) for those tracks that got longer in Dogwood3 (the small red peak on the left). It is clear that Dogwood3 is reconstructing their energy better than Dogwood1 is, so the longer track length must be the right one.

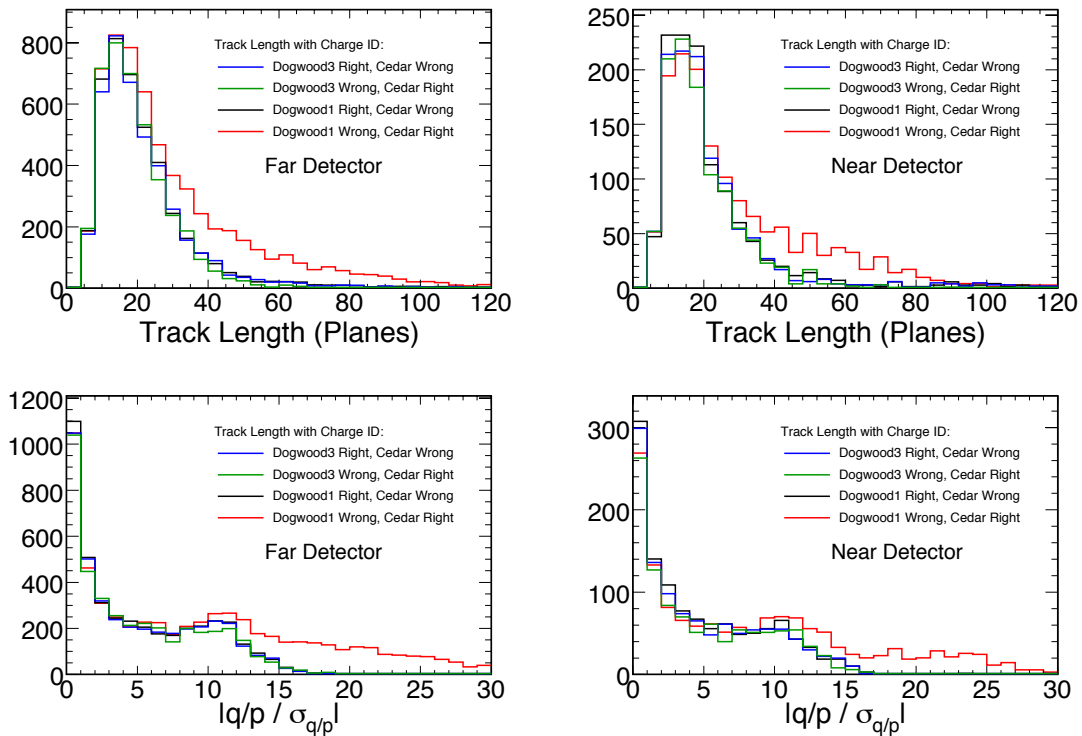


Figure E.4: Above are the track length and  $(q/p) / \sigma_{q/p}$  distributions for events where Cedar and Dogwood disagree about the track charge sign. The asymmetry we see in Dogwood1 (the red vs. black), where there is a subset of events with long tracks and high  $(q/p) / \sigma_{q/p}$  that only Dogwood gets wrong, is now gone in Dogwood3 (blue vs. green).





## Appendix F

# Technical Aspects of Flugg

### F.1 Introduction

This appendix describes the technical details of the newest version of the flux simulation for the NuMI beam line. Section 3.1.4 gives an overview of simulating the NuMI beam in general as well as the motivations for upgrading from the previous simulation. This appendix gives detailed instructions for installing and running the NuMI Flugg simulation, detailed information about the organization and structure of the source code, and tables describing the output data files. It also includes validation studies which compare it in detail to the previous NuMI simulation.

It introduces two major changes. The previous flux simulation versions used a combination of Fluka and GEANT3: Fluka to simulate the target and GEANT3 to swim the particles coming off the target through the rest of the beamline geometry. The new simulation also uses a combination of Fluka and GEANT, but uses the Flugg [125, 126, 127] package to simulate a GEANT4 geometry in Fluka. Instead of having a two step simulation process, now all of the geometry is handled by GEANT and all of the interactions are handled by Fluka. Interactions in the downstream regions are now handled by the same, more trusted, hadronic interaction model as interactions in the target. The largest advantage of this change is in the simulation of helium interactions in the decay pipe. The GEANT4 geometry used is also a significant upgrade from the earlier version. It includes more detailed descriptions of many of the beamline components (horns, decay pipe window, hadron absorber, muon monitors) as well as more accurate dimensions based on as-built information where available.

## F.2 How to Run

### F.2.1 Setup and Installation

First, you will need working copies of GEANT4,<sup>1</sup> Fluka,<sup>2</sup> and Flugg;<sup>3</sup> instructions for installing all three can be found online. You will need to install the first two before installing Flugg. The code required for g4numi in Flugg can be found in two directories in CVS:

```
numisoft/g4numi
numisoft/g4numi_flugg
```

After checking out the two directories, run `g4numi_flugg/scripts/setup.sh`. This sets up the symlinks from the geometry source files in `g4numi` to `g4numi_flugg`. It also replaces a few source files in the main Flugg source directory, fixing some bugs that prevent the simulation from running smoothly.<sup>4</sup> Be sure to set up the Flugg environment first. This is necessary so the script can put the patched Flugg source files in the right location. After this, compile the executable by running `gmake` in `g4numi_flugg/`.

### F.2.2 Running the Simulation

The process of running Flugg can be somewhat onerous so a script was written that automates the process. You can find it here:

```
g4numi_flugg/scripts/g4numi_fluka.sh
```

In order to simplify running on a batch system, it does not take command line arguments but instead uses environmental variables to configure it. Table F.1 lists the required variables and Table F.2 lists the optional ones. The script will automatically create a parent directory based on the beam configuration and then a sub directory for each run like this:

```
$DATA/flugg\-le<z position>z<current>i\_run<run period>/Run<job number>
```

for example,

```
$DATA/flugg\-le010z185i\_run2/Run109
```

Since the random seed for each run is based off the run number, individual jobs can be rerun and will return the same results. Table F.3 shows the files you can expect to find in the job directory of a successful run.

---

<sup>1</sup>GEANT4: <http://geant4.web.cern.ch/>

<sup>2</sup>Fluka: <http://www.fluka.org/>

<sup>3</sup>Flugg: <http://www.fluka.org/content/tools/flugg/>

<sup>4</sup>This is no longer the default behavior. The patched files are still distributed but due to frequent version clashes they are not automatically copied into the Flugg source tree.

Variable	Example	Description
RUN	109	The run number for this job. It defines the output names and the random seed.
EVTS	500000	The number of protons-on-target to simulate.
PERIOD	2	Run Period (1-4). Periods 3 and 4 have Helium automatically.
DATA	<code>\$HOME/data/flugg</code>	Directory where files are output.

Table F.1: Required environmental variables for `g4numi_fluka.sh`.

Variable	Default	Description
CURRENT	185	The nominal horn current in (kA). It is converted to the actual value internally (e.g. 185 $\rightarrow$ 182.1 kA).
TARGETZ	10	The nominal target z position in -cm. 10 becomes -8.9 cm automatically for runs II and III.
STEPL	1.0	The maximum step size in cm. Be sure to include the decimal point, for example <code>export STEPL="1.0"</code> .
SPECIAL	-	The contents of special is added to the run type. For example, if <code>SPECIAL="_sh"</code> , then the files would output to <code>flugg_le010z185i_run1_sh/</code> .
LOWTH	-	Set to "yes" to remove the 1 GeV tracking threshold.
TARGFILE	-	Set to "yes" to produce target hadron ntuples in addition to the neutrino files.

Table F.2: Optional environmental variables for `g4numi_fluka.sh`.

### F.2.3 Inside the Scripts

In addition to basic bookkeeping tasks (creating and moving to directories, renaming files, etc.), the script does three major things. First, it creates the Fluka input file (`g4numi.inp`) that configures many aspects of the simulation. Second, it runs the executable using the `rfluka` script provided in the Fluka distribution. Third, it conserves disk space by converting the output from ascii to the two binary file formats we use (hbook and root).

The command used to run the executable is:

```
$FLUPRO/flutil/rfluka -e $FLUGGINSTALL/bin/Linux-g++/mainG4NuMI -NO -M1 g4numi
```

This tells `rfluka` to use the `mainG4NuMI` executable, run 1 job starting at run number 0 (this is so that the random seed given in the input file is used and Fluka doesn't look for an already existing random seed) and to use the `g4numi.inp` input file. `rfluka` creates a temporary directory with the name `fluka_<pid>` in the job directory (`.../Run<job number>`). This is where you will find the log files and partially completed output files for jobs that crash or are still in progress.<sup>5</sup> Some of the filenames will be different from those in a finished run because `rfluka` renames the files when the job completes and `g4numi_fluka.sh` renames some of them again, but it should still be relatively

<sup>5</sup>You can identify the crashes since they will have core dumps (`core.<pid>`) if your system allows them.

File	Description
<code>flugg_le010z185i_run2_109.ntp</code>	The neutrino ntuple in hbook format.
<code>flugg_le010z185i_run2_109.root</code>	The neutrino ntuple in root format.
<code>target109.ntp</code>	The target hadron ntuple in hbook form. (Optional)
<code>target109.root</code>	The target hadron ntuple in root form. (Optional)
<code>flugg_le010z185i_run2_109.out</code>	The standard out from the Fluka parts of the simulation.
<code>flugg_le010z185i_run2_109.err</code>	The standard error from the Fluka parts of the simulation.
<code>flugg_le010z185i_run2_109.log</code>	The standard out and error from the Flugg/GEANT parts of the simulation.
<code>g4numi.inp</code>	The Fluka input file created by <code>g4numi_fluka.sh</code> .
<code>runConfig.inp</code>	The geometry configuration file used by the (GEANT part) of the executable.
<code>g4numi001_flukaMat.inp</code>	The Fluka material configuration created by Flugg (see Section F.2.3).
<code>g4numi001_Volumes_index.inp</code>	File created by Flugg that gives the association between Fluka region numbers and GEANT names.

Table F.3: Output files found in the directory of a successful job. As an example the files for Run 109 of a le010z185i, Run II job are shown.

simple to identify the files you care about.

The Fluka simulation does not write directly to binary data, it instead writes one or two ascii files (depending on whether the target ntuple is being produced) which are then converted to the two binary formats with the following scripts/programs in `g4numi_flugg`: `root/fill_flux.C`, `root/fill_target.C`, `hbook/ascii2hbook.beam`, `hbook/ascii2hbook.target`. The former two are root scripts and the latter two are actually small compiled Fortran programs (they should get compiled when you compile the main executable so you should not have to worry about them). The ascii files are only removed when both the hbook and root files are produced successfully.

The input file created by this script configures many aspects of the simulation. The input file starts with a series of cryptic commands (called ‘cards’ in Fluka jargon), like `USERDUMP` and `MAT-PROP` which are included in order to activate pieces of user written code. There are also cards to turn off particles we do not care about (electrons and photons). This part of the input file is read in from `g4numi_flugg/g4numi.start` and is unlikely to need to be changed by the user.

The next part of the input file assigns materials to regions of the geometry. This part is read in from a file generated automatically by Flugg. This is why, under normal circumstances, Flugg simulations need to be run twice, first to produce this file and second to do the full simulation.<sup>6</sup> For the geometry as it exists in CVS, these material assignment files are generated ahead of time

<sup>6</sup>This is described in more detail at the Flugg website, <http://www.fluka.org/content/tools/flugg/node8.html>

and can be found in `g4numi_flgug/g4numi_flukaMat.<name>.inp`. If you make a change to the geometry that changes the number of regions (adds/removes) or changes the material of a region, you will need to add your own `flukaMat` file and edit `g4numi_fluka.sh` to read that file in instead. The easiest way to get the file would be to run the script as usual and stop the processing by hand after 30 seconds or so (no reason to run it longer since its results will be nonsense with the wrong materials file). You can then find the correct `flukaMat.inp` file in the `<flugg_run_dir>/Run<process no>/fluka_<pid>` directory.

After the material definitions comes the proton beam configuration. The energy and divergence are always the same, though the size of the beamspot changes for different run periods. The beam starts at  $-6$  m to make sure it starts ahead of the target for all beam configurations. After the beam are a few settings turned on or off by some of the environmental variables described in Table F.2. These cards do things such as set the step size, turn off multiple scattering, etc. The final entries of the input file set the random seed, set the number of protons to simulate, and call the user code that opens and closes the output files.

In addition to the main `g4numi.inp` file, `g4numi_fluka.sh` also produces a very small configuration file called `runConfig.inp` that the executable reads to determine some parameters that affect how the geometry is constructed (run period, target position, horn current).

## F.3 The Code

### F.3.1 Overview

The Flugg simulation code falls into two broad categories: the main simulation, written in Fortran using Fluka, and the geometry description, written in C++ using GEANT4. The basic idea for the interaction between these two parts is this: Fluka performs the actual simulation (performs particle interactions, tracks particle properties, writes output files), but when Fluka goes to make a query about the geometry, the Flugg code uses wrappers to pass questions like “Where am I?” and “How far is it to the region boundary?” off to GEANT4.

Here, in broad strokes, is how the simulation proceeds from primary proton to output neutrino. For more details see the referenced sections. The simulation begins with a primary proton defined in `g4numi.inp` (Section F.2.3). When the primary proton enters the target region, its initial properties are recorded by `FLUSCW` and are used for all the particles which come from this proton. While in the target, the proton interacts, producing potentially many secondaries. As these secondaries leave the target volume, their properties are stored in the (NUMI) common block, again by `FLUSCW`. All the later generations of particles following from a particular target parent keep a reference to the target parent’s entry in the common block in an `ISPUSR` variable. This way, when we get a neutrino we have the information about the particle that produced it as it left the target. This functionality

reproduces the behavior of earlier simulations where the target was simulated separately from the rest of the beamline.

Whenever a secondary is produced, inside the target or out, the `STUPRF` user routine is called. It performs two important functions: it calls the importance weighting function (see Section F.3.2) and it keeps track of the particle's parent information (see Section F.3.3). When the particles leave the target, they and any secondaries they produce are tracked through the rest of the beamline geometry. The new `g4numi` geometry that is used by Flugg contains a much more complete and detailed geometry description than the previous simulations, including the detailed structure of the horns and the decay pipe window. It also includes simulations of the hadron and muon monitors.

Whenever one of these tracked particles decays, the `USDRAW` user routine gets called. This routine loops through all of the created secondaries. For each secondary that is a neutrino, an entry is written to the neutrino `ntuple`. Note, we do not generally use the neutrino direction chosen randomly by the decay within the simulation but instead use the neutrino's parent information to calculate a weight (or probability) for that neutrino to reach the center of the Near or Far Detectors. This weight is calculated by the `NUWEIGHT` function. The entries in these `ntuples` are described in detail in Section F.4.2.

### F.3.2 Importance Weighting and Thresholds

The basic idea behind importance weighting is that we can save disk space and processing time by not simulating every single one of the lower energy particles but instead only keeping a fraction of them, but weight that fraction up so that the weighted flux is unchanged. The idea is to more evenly distribute the statistics and processing time across energies, even though we produce many fewer high energy particles. A parent's weight is passed down to its children, which each may apply their own additional weighting to get their total weight. Every particle has a weight of at least 1, and we set a maximum weight of 100 to avoid having very rare, ultra high weight events.

Importance weighting is performed by the `NUWEIGHT` function which is called by `STUPRF`. The algorithm is as follows. First, the code decides whether or not a particle should be importance weighted.  $\nu$ 's,  $\mu$ 's, and  $K$ 's are not given any additional weight (though they may carry a weight greater than 1 from their parents). Particles with a total momentum above 30 GeV are also not given any additional weight. For the remaining particles, a candidate weight,  $W$ , is calculated as

$$W = 30/|P_{tot}|.$$

We then check to see if this weight would push the particle's total weight ( $W_{tot} = W \times W_{parent}$ ) above 100. If it does,  $W$  is set such that  $W_{tot} = 100$ . We then decide whether to simulate or discard this particle. We do this by choosing a random number  $R$  `Uniform(0, 1)` and if  $R > 1/W$  we discard

the particle.

There is a general tracking threshold throughout the simulation of 1 GeV. That is, if a particle ever has an energy of below 1 GeV, it is discarded. While this is fine for MINOS, for other off-axis experiments lower energy fluxes may be needed. This threshold can be turned off completely from the run scripts with the `LOWTH` variable (see Table F.2 in Section F.2.2). Alternatively, a different threshold can be set by editing the `PART-THR` card which is defined in `g4numi_fluka.sh`.

### F.3.3 Particle History Tracking

`STUPRF` is called whenever a particle is added to the “stack,” the collection of particles that have been produced but have not been tracked yet. When `STUPRF` is called, we have both the parent properties, stored in the `TRACKR` block, and the child properties, stored in `GENSTK` block. It is here that the spare tracking variables described in Table F.4 are filled. The current position, which is the production point of the secondary, is stored in `SPAUSR(1-3)`. Since each secondary is really just a potential neutrino parent, this point is also referred to as the ‘(neutrino) parent production point.’ The current momentum of the secondary, that is the momentum at its production point, is also stored here in `SPAUSR(7-9)`. The parent production medium is recorded based on the properties of the current region and the parent species is also recorded.

Finally, the `SPAUSR(4-6)` variables are set based on the parent of the current secondary (or potential neutrino grandparent), i.e. using the `TRACKR` variables. If the current secondary is a muon, these variables are set to the grandparent’s decay momentum (the current momentum in `TRACKR`). Otherwise, the variables are set to the values in `SPAUSR(7-9)`, that is the production point information of the grandparent. This dichotomy occurs because these same three slots need to be used for two different purposes due to having a limited number of available `SPAUSR` variables. For muons, the grandparent decay information is needed in order to calculate the effects of muon polarization on the neutrino weight at the detectors. For non-muon parents, the grandparent production point information is needed to fill the `tgpr*` variables in the ntuple in a way equivalent to earlier simulations.<sup>7</sup>

### F.3.4 Fluka Source Files

`for/(NUMI)`

A common block to store target parent information.

`for/fluscw.f`

The `FLUSCW` routine is called in a number of circumstances, but the one used is that it is called whenever a boundary is crossed. In particular, it is used to catch when particles enter or leave

<sup>7</sup>For the production run of files, a bug caused only the decay point to be recorded, not the production point. This means that the `mupar*` variables and the `tgpr*` variables all refer to the decay point for all particles.

Variable	Description
SPAUSR(1-3)	x, y, z at neutrino parent production (grandparent decay) point
SPAUSR(4-6)	px, py, pz of the neutrino grandparent at the decay point (muon parents) or production point (hadron parents)
SPAUSR(7-9)	px, py, pz at the neutrino parent production point
SPAUSR(10)	Particle importance weight
ISPUSR(1)	Index in (NUMI) of the target parent
ISPUSR(2)	Neutrino grandparent species
ISPUSR(3)	Neutrino great-grandparent species
ISPUSR(4)	Flag to kill particles. Set to 1 to keep a particle and 2 to kill it.
ISPUSR(5)	1 if SPAUSR(4-6) refers to a muon parent
ISPUSR(6)	Parent production mode flag (1 = inelastic, 2 = decay, 0 = other)
ISPUSR(7)	Parent production medium

Table F.4: Describes how all of the Fluka spare tracking variables `ISPUSR[]` and `SPAUSR[]` are used.

the target region. Whenever a proton enters the target region, its information is stored and used for all target secondaries coming from that proton. Whenever a particle leaves the target region, its properties are stored in an entry in the (NUMI) common block (the block is reset at each new proton). The entry number is stored in an `ISPUSR` variable for that particle which is then passed on to all later generations coming from that particle. That way, when a neutrino is created we can access the target parent information. If the target hadron ntuple is being written out, that happens here as well (the hadron ntuple is described in Table F.8 in Section F.4.2). It contains all the target secondaries that are added to the (NUMI) common block.

#### `for/gcode.f`

A routine that converts Fluka particle codes to GEANT particle codes. It is called while the neutrino ntuple is being written out in `MGDRAW`.

#### `for/impwgt.f`

Implements the importance weighting scheme. See Section F.3.2 for more details.

#### `for/magfld.f`

Passes the Fluka magnetic field call off to the GEANT4 wrapper so that the GEANT4 field is used. This should not require modification by the user.

#### `for/mcode.f`

A routine that converts Fluka material codes to the Gnumi material codes. If a given Fluka material does not have an equivalent in the Gnumi scheme, the function returns the original Fluka material multiplied by  $-1$ . It is called while the neutrino ntuple is being written out in `MGDRAW`.

#### `for/mgdraw.f`

`MGDRAW` has numerous subroutines. The one used in this simulation is `USDRAW`. It is called after



every particle interaction. The process of interest is `ICODE = 102` which refers to decay. The neutrino ntuple is written out here (the neutrino ntuple is described in Table F.6 and Table F.7 in Section F.4.2). Whenever there is a decay, the function loops through the secondaries. For each neutrino found, an entry is written to the ntuple. Note, when the ntuple is being written the `TRACKR` variables still refer to the neutrino *parent* at its decay point, not the neutrino.

#### `for/nuweight.f`

The `NUWEIGHT` routine defined in this file calculates the Near and Far Detector (or any other position) weights. It takes polarization into account for muons.

#### `for/stuprf.f`

The `STUPRF` routine is called whenever a new particle is added to the stack, i.e. when we get new particles from a decay or an inelastic interaction. This routine is responsible for the history tracking or “latching.” The particles own production point information is stored (position and momentum) along with the momentum of the particle’s parent at either production (muons) or decay (hadrons)(at the decay point in production files – see footnote on page 173). The variables are described in Table F.4.

#### `for/usrini.f`

This user routine is called when the executable first begins to run. It opens the needed output ascii files.

#### `for/usrmed.f`

`USRMED` is used normally to apply custom material properties. Here it is used to allow importance weighting to kill particles. If another part of the code (e.g. `IMPWGT`) decides to kill a particle, it sets `ISPUSR(4) = 2` which tells `USRMED` to zero its weight, stopping it from being further simulated.

#### `for/usrout.f`

This user routine is called at the end of the Fluka run. It closes the opened output files.

### F.3.5 GEANT Source Files

#### `include/NumiDataInput.hh, src/NumiDataInput.cc`

The `NumiDataInput` class creates a singleton object that holds all of the configurable parameters of the beamline geometry. It reproduces the functionality of the input files used in the GEANT3-based Gnumi. It is this class that reads in the `runConfig.inp` file described in Section F.2.3. It is used by all the other geometry files.

#### `include/NumiDetectorConstruction.hh, src/NumiDetectorConstruction.cc`

This is the starting point for constructing the NuMI geometry and contains the method actually

called by Flugg (`NumiDetectorConstruction::Construct()`). `NumiDetectorConstruction.cc` is itself mostly just a jumping off point that calls separate methods contained in separate files that produce the rest of the geometry. These files are listed below and their names are self-explanatory. Since they are all technically just a part of the `NumiDetectorConstruction` class they do not have their own associated header files.

```
src/NumiBaffle.cc
src/NumiDecayPipe.cc
src/NumiHadronAbsorber.cc
src/NumiHorn1.cc
src/NumiHorn2.cc
src/NumiMaterials.cc
src/NumiSecMonitors.cc
src/NumiTarget.cc
src/NumiTargetHall.cc
```

`include/NumiHornSpiderSupport.hh`, `src/NumiHornSpiderSupport.cc`

A specialized object used for creating the horn support structures.

`include/NumiMagneticField.hh`, `src/NumiMagneticField.cc`

Actually contains three classes that describe the magnetic field in the inner conductor, in the outer conductor, and in the bulk of the horn volume.

Note: These are not all of the source files in `g4numi` but rather just the subset that are used by Flugg.

## F.4 The Output Ntuples

### F.4.1 Available Files

The Flugg flux files can be found at Fermilab. They are available in `/minos/data/flux/flugg/` with subdirectories for each run configuration. The available run configurations are described in Table F.5.

Beam	No. Files	Target	Target $Z$ (cm)	Current (kA)	Decay Pipe
le010z000i_run1	317	NT01	-10	0	Vacuum
le010z000i_run2	449	NT02	-8.9	0	Vacuum
le010z000i_run3	440	NT02	-8.9	0	Helium
le010z170i_run1	441	NT01	-10	167.3	Vacuum
le010z185i_run1	441	NT01	-10	182.1	Vacuum
le010z185i_run2	431	NT02	-8.9	182.1	Vacuum
le010z185i_run3	442	NT02	-8.9	182.1	Helium
le010z185i_run4	442	NT03	-10	-182.1	Helium
le010z200i_run1	445	NT01	-10	196.9	Vacuum
le100z200i_run1	446	NT01	-100	196.9	Vacuum
le150z200i_run2	445	NT02	-150	196.9	Vacuum
le250z200i_run1	445	NT01	-250	196.9	Vacuum
le250z200i_run2	448	NT02	-250	196.9	Vacuum

Table F.5: Available files at FNAL. You can access them at `/minos/data/flux/flugg/`.

### F.4.2 File Contents

Variable	Description
run	Run number (not used)
evtno	Event number (proton on target)
Ndxdz Ndydz	Neutrino direction slopes for a random decay
Npz	Neutrino momentum (GeV/c) along the $z$ -axis (beam axis)
Nenergy	Neutrino energy (GeV) for a random decay
NdxdzNea NdydzNea	Direction slopes for a neutrino forced towards the center of the Near Detector
NenergyN	Energy for a neutrino forced towards the center of the Near Detector
NWtNear	Weight for a neutrino forced towards the center of the Near Detector
NdxdzFar NdydzFar	Direction slopes for a neutrino forced towards the center of the Far Detector
NenergyF	Neutrino energy (GeV) for a decay forced to the center of the Far Detector
NWtFar	Neutrino weight for a decay forced to the center of the Far Detector
Norig	Not used in flux file
Ndecay	Decay process that produced the neutrino, see Table F.11
Ntype	Neutrino flavor. $\nu_\mu = 56, \bar{\nu}_\mu = 55, \nu_e = 53, \bar{\nu}_e = 52$
Vx Vy Vz	Neutrino production vertex (cm)
pdPx pdPy pdPz	Momentum (GeV/c) of the neutrino parent at the neutrino production vertex (parent decay point)
ppdxdz ppdydz	Direction of the neutrino parent at its production point (which may be in the target)
pppz	$z$ momentum (GeV/c) of the neutrino parent at its production point
ppenergy	Energy (GeV) of the neutrino parent at its production point
ppmedium	Code for the material the neutrino parent was produced in (see Table F.10)
ptype	Neutrino parent species (GEANT codes)
ppvx ppvy ppvz	Production vertex (cm) of the neutrino parent
muparpx muparpy muparpz	Momentum (GeV/c) of the neutrino grandparent at the grandparent decay point (muons) or grandparent production point (hadrons) (at the decay point in production files – see footnote on page 173)
mupare	Energy (GeV) of the neutrino grandparent, as above
Necm	Neutrino energy (GeV) in the center-of-mass frame
Nimpwt	Importance weight of the neutrino
xpoint ypoint zpoint	Debugging hook – unused

Table F.6: The entries stored in the neutrino ntuple files. There is one entry for every neutrino produced.

Variable	Description
tvx	Position (cm) of the neutrino ancestor as it exits target (possibly, but not necessarily, the direct neutrino parent)
tvz	
tpx	
tpy	Momentum (GeV/c) of the ancestor as it exits target
tpz	
tptype	
tgen	Neutrino parent generation in cascade. 1 = primary proton, 2 = particles produced by proton interaction, 3 = particles from
tgptype	Species of the parent of the particle exiting the target (GEANT codes)
tgppx	Momentum (GeV/c) of the parent of the particle exiting the target at the parent production point (at the decay point in production files – see footnote on page 173)
tgppy	
tgppz	
tprivx	Primary particle interaction vertex (not used)
tprivy	
tprivz	
beamx	Primary proton origin (cm)
beamy	
beamz	
beampx	Primary proton momentum (GeV/c)
beampy	
beampz	

Table F.7: The entries stored in the neutrino ntuple files. There is one entry for every neutrino produced.

Variable	Description
x	Position (cm) of the particle as it exits target
y	
z	
px	Momentum (GeV/c) of the parent as it exits target
py	
pz	
type	Species of the particle leaving the target (Fluka codes)
weight	Weight
gener	Generation
momtype	Species of the parent of the particle exiting the target (Fluka codes)
momp	Momentum (GeV/c) of the parent of the particle exiting the target at the parent decay point
px	
py	
protvx	Primary particle interaction vertex (not used)
protvy	
protvz	
protx	Primary proton origin (cm)
proty	
protz	
protpx	Primary proton momentum (GeV/c)
protpy	
protpz	
event	Event number

Table F.8: The entries stored in the hadron ntuple files. There is one entry for each hadron that exits the target volume.

## F.4.3 Useful Codes

Particle	Fluka	GEANT	PDG
	Hadron File	Flux File	SNTF
$\gamma$	7	1	22
$e^+$	4	2	-11
$e^-$	3	3	11
$\mu^+$	10	5	-13
$\mu^-$	11	6	13
$\pi^0$	23	7	111
$\pi^+$	13	8	211
$\pi^-$	14	9	-211
$K_L^0$	12	10	130
$K^0$	24	10/16	311
$\bar{K}^0$	25	10/16	-311
$K^+$	15	11	321
$K^-$	16	12	-321
$n$	8	13	2112
$p$	1	14	2212
$\bar{p}$	2	15	-2212
$K_S^0$	19	16	310
$\Lambda$	17	18	3122
$\Sigma^+$	21	19	3222
$\Sigma^0$	22	20	3212
$\Sigma^-$	20	21	3112
$\Xi^0$	34	22	3322
$\Xi^-$	36	23	3312
$\Omega^-$	38	24	3334
$\bar{n}$	9	25	-2112
$\bar{\Lambda}$	18	26	-3122
$\bar{\Sigma}^-$	31	27	-3222
$\bar{\Sigma}^0$	32	28	-3212
$\bar{\Sigma}^+$	33	29	-3112
$\bar{\Xi}^0$	35	30	-3322
$\bar{\Xi}^+$	37	31	-3312
$\Omega^+$	39	32	-3334
$\tau^+$	41	33	-15
$\tau^-$	42	34	15
$\bar{\nu}_e$	6	52	-12
$\nu_e$	5	53	12
$\bar{\nu}_\mu$	28	55	-14
$\nu_\mu$	27	56	14
$\bar{\nu}_\tau$	44	-	-16
$\nu_\tau$	43	-	16

Table F.9: The particle codes across the three schemes used in MINOS.

Code	Material
5	Beryllium
6	Carbon
9	Aluminum
10	Iron
11	Slab Steel
12	Blu Steel
15	Air
16	Vacuum
17	Concrete
18	Target
19	Rebar Concrete
20	Shotcrete
21	Variable Density Aluminum
22	Variable Density Steel
23	1018 Steel
24	A500 Steel
25	Water
26	M1018 Steel
28	Decay Pipe Vacuum
31	CT852

Table F.10: The material codes as defined by Gnuni and used in the fluxfiles, old and current.

Ndecay	Process
1	$K_L^0 \rightarrow \nu_e + \pi^- + e^+$
2	$K_L^0 \rightarrow \bar{\nu}_e + \pi^+ + e^-$
3	$K_L^0 \rightarrow \nu_\mu + \pi^- + \mu^+$
4	$K_L^0 \rightarrow \bar{\nu}_\mu + \pi^+ + \mu^-$
5	$K^+ \rightarrow \nu_\mu + \mu^+$
6	$K^+ \rightarrow \nu_e + \pi^0 + e^+$
7	$K^+ \rightarrow \nu_\mu + \pi^0 + \mu^+$
8	$K^- \rightarrow \bar{\nu}_\mu + \mu^-$
9	$K^- \rightarrow \bar{\nu}_e + \pi^0 + e^-$
10	$K^- \rightarrow \bar{\nu}_\mu + \pi^0 + \mu^-$
11	$\mu^+ \rightarrow \bar{\nu}_\mu + \nu_e + e^+$
12	$\mu^- \rightarrow \nu_\mu + \bar{\nu}_e + e^-$
13	$\pi^+ \rightarrow \nu_\mu + \mu^+$
14	$\pi^- \rightarrow \bar{\nu}_\mu + \mu^-$
999	Other

Table F.11: The decay codes stored in Ndecay.

## F.5 Validation

The validation section consists of comparisons of the new Flugg flux (v20) to the previous Gnumi flux (v19). The first subsection discusses known changes and their effects on the flux. In this section, if a plot has more than one histogram, the red refers to Flugg and the black refers to Gnumi. If a plot has only one red histogram, then it is the ratio of Flugg to Gnumi. The second subsection includes a more comprehensive set of plots looking at the Near Detector energy spectrum and Far-over-Near ratio for the following:  $\nu_\mu$  in every beam configuration and normal low energy (le010z185i) for every neutrino species ( $\nu_\mu, \bar{\nu}_\mu, \nu_e, \bar{\nu}_e$ ) and every major parent species ( $\pi^\pm, K^\pm, K_L^0, \mu^\pm$ ).

### F.5.1 The Focusing Peak

There are, in fact, two observable changes in the focusing peak: a decrease in the height of the peak and a shift towards higher energy. We can separate out these two effects, and thus get a hint of their cause, by looking at where the neutrino was produced. Figure F.1 shows the focusing peak for neutrinos with vertices in the chase and Figure F.2 shows the peak for neutrinos with vertices in the decay pipe. We see in the former figure that there is just a reduction for a certain energy range without any shift in the spectrum. What is happening is that some parents that would decay to neutrinos are instead being absorbed by the additional horn material that has been added in this simulation. It happens only for a particular range of energy because not all initial parent momenta will have their paths focused to pass through the extra material. This hypothesis has been confirmed with a special simulation run without the extra horn material – the peak decrease is not present in this special run.




Figure F.1: The focusing peak for neutrinos with vertices in the chase ( $V_z < 45$  m). We see an energy-dependent decrease in the peak height without any horizontal shift. The integrated number of these events changes by about 7%. Neutrinos with vertices in the chase makes up about 18% of the focusing peak.






Figure F.2: The focusing peak for neutrinos with vertices in the decay pipe ( $V_z > 45$  m). We see a horizontal shift the peak without a change in its width or height. The mean shifts by 2.3% while the RMS and total integral change by 0.6% and 0.1% respectively.

### F.5.2 Helium and Other Downstream Production

Production due to interactions outside the target was one of the primary motivators for developing the Flugg fluxes, since this is where the previous version of the simulation had difficulty. Specifically, the problem was that there was too much high- $x_f$  production from interactions in the helium – much more than was consistent with the data. This was because downstream production was handled by the obsolete Gfluka package, rather than a modern Fluka version as was used in the target. Now, with a consistent hadron production model, the new Flugg fluxes do, indeed, show a decrease in downstream production, especially at higher energies. It is seen for both for  $\nu_\mu$ 's (Figure F.3) and  $\bar{\nu}_\mu$ 's (Figure F.4). We also see the benefit of this decrease in better modeling of the change in the flux do to helium. This can be seen, both for  $\nu_\mu$ 's and  $\bar{\nu}_\mu$ 's in Figure F.5.




Figure F.3: The Near Detector low energy spectrum for  $\nu_\mu$ 's whose parents were produced outside the target. We see a significant decrease (24%), especially towards higher energies.




Figure F.4: The Near Detector low energy spectrum for  $\bar{\nu}_\mu$ 's whose parents were produced outside the target. As with the  $\nu_\mu$ 's, we see a significant decrease (20%).




Figure F.5: The ratio of the Near Detector spectrum with helium in the decay pipe to the spectrum with an evacuated decay pipe for neutrinos (left) and antineutrinos (right). The black line represents the older GNUMI simulation, the red line represents the newer Flugg simulation, and the blue points represent the data. As shown, the Flugg-based simulation is significantly better at reproducing the effects of helium seen in the data.

### F.5.3 The Near Detector $\nu_e$ Flux

Perhaps the most significant unexpected change in the switch to the new flux was a large (14%) increase the  $\nu_e$  flux at the Near Detector, as seen in Figure F.6. It turns out the change is more generally for all neutrinos from muon decay (16% increase). The other  $\nu_e$  component, coming from kaon decays, shows an increase of about 3%, both Near and Far.<sup>8</sup> The muon increase is primarily in the Near Detector: the Far Detector and the non-detector weighted fluxes show more modest changes (8% increase and 0.1% increase, respectively).

This behavior is suggestive – the fact that the increase is only apparent for detector-weighted spectra, and is larger for the Near Detector, points us towards the immediate cause: a change in the  $V_z$  (neutrino  $z$  vertex) distribution. Figure F.8 shows the Near Detector-weighted  $V_z$  distribution. We can see that the increase in Near Detector flux comes entirely at high- $V_z$ . Figure F.9 shows the same distribution but without weighting to the Near Detector. We see that what appeared to be an

<sup>8</sup>Which is actually quite small compared to the 20% increase in low-energy kaon production due to Flugg's better handling of low energy interactions compared to Gnumi.

increase at high- $V_z$  is actually just a shift from lower  $V_z$  to higher  $V_z$  without a significant change in overall amount. The change in the overall Near Detector flux comes from the fact that neutrinos produced towards the end of the decay pipe have a higher probability of hitting the Near Detector (and thus a higher Near Detector weight) than those produced at the beginning of the decay pipe. This change in the  $V_z$  distribution is not limited to muons, as can be seen in Figures F.10-F.11; however, the affect on  $\nu_\mu$ 's is not as pronounced since the pions tend to decay earlier in the decay pipe, making the shift towards higher  $V_z$  less pronounced.

The question now becomes why is there a change in the  $V_z$  distribution? Unfortunately, this is a difficult question to answer. However, we do have some strong circumstantial evidence that the cause is a change in horn focusing. First, a note about the changes in the Flugg horns. The extra horn material added to Flugg is at the far end in  $z$  and at high  $r$ . This means that there is an additional focused component that is going straighter but with a wider divergence. This additional 'straightness' might be the cause of the change in the  $V_z$  distribution. We also see some evidence of the wider divergence. Figure F.12 shows a variable named 'divx' that is calculated by taking the neutrino parent at its decay point and using its final momentum to project it back to what its  $x$  coordinate would have been at  $z = 45$  m. Both of these distributions do, indeed, show that Flugg has a wider distribution. However, because of the way the variable is constructed, particles that decay closer to the end of the decay pipe will tend to have larger values since they have to project farther back and thus this variable does not necessarily tell us that the raw flux leaving the horns is wider. It gives us the correlation, but it cannot demonstrate causation.

A second piece of circumstantial evidence comes from looking at the relative muon flux, separated by charge, in the different beam configurations. You can see this plotted in Figure F.13 and Figure F.14. For  $\mu^+$ 's, as the focused energy increases the Near Detector excess decreases until there is actually a small deficit in the pseudo high-energy configuration. For the  $\mu^-$ 's, on the other hand, there is no statistically-significant excess in any focused beam configuration.<sup>9</sup> The fact that changing the horn focusing does affect the flux of  $\mu^+$ 's but not the flux of  $\mu^-$ 's strongly suggests (but cannot *prove*) that the horn focusing is the cause of the changes between Flugg and Gnumi. It also cannot tell us what it is about the focusing that makes the difference.

We are limited by not having any information about the immediate muon parent and also by only having information about particles that go on to produce neutrinos. In order to do conclusively prove this hypothesis, information beyond what is stored in the flux files is required, both in Flugg and in Gnumi. This includes more information about the muon parent or a dump of all particles that pass through a plane in the simulation, regardless of whether or not they decay to neutrinos.

---

<sup>9</sup>However, there is a significant excess in the horn-off configurations that is not, at this time, understood.




Figure F.6: At left is the Near Detector  $\nu_e$  and  $\bar{\nu}_e$  energy spectrum. It shows a 14% overall increase compared to Gnumi, concentrated at lower energies. At right is the Near Detector energy spectrum for all neutrino types from  $\mu^\pm$  decay. Here there is a 16% increase which is relatively consistent for all energies.

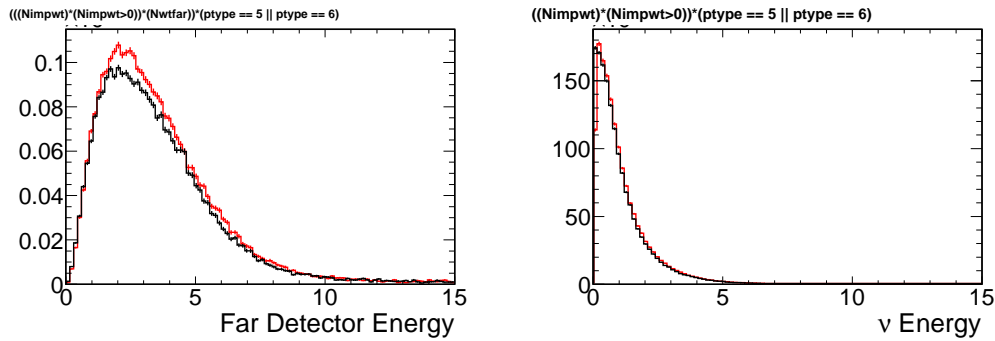


Figure F.7: Both figures show the energy spectrum for neutrinos from  $\mu^\pm$  decay. At left is the Far Detector spectrum, where Flugg shows an 8% increase relative to gnumi. At right is the non-detector (importance-only) weighted muon flux. Here the total number of muons produced is identical to within 0.1%.




Figure F.8: The Near Detector weighted  $V_z$  distribution for neutrinos from muon decay. We can see that the 16% increase in flux is concentrated at higher  $V_z$ .




Figure F.9: The importance (non-detector) weighted  $V_z$  distribution for neutrinos from muon decay. We can see that the increase shown in Figure F.8 is actually a shift from lower to higher  $V_z$  with little change in overall flux when not detector weighted. The increased overall flux comes because neutrinos produced at higher  $V_z$  (closer to the Near Detector) have higher Near Detector weights (or equivalently, are more likely to hit the Near Detector).




Figure F.10: The Near Detector weighted  $V_z$  distribution for  $\nu_\mu$ 's. While the shift from low to high is still present, the change in overall number of events is much smaller because the distribution is already shifted heavily towards the front of the decay pipe.




Figure F.11: The importance (non-detector) weighted  $V_z$  distribution for  $\nu_\mu$ 's.




Figure F.12: Both of these plots show the *divx* variable – this is the projected *x*-position the neutrino parent would have passed through at  $z = 45$  m based on its final position and momentum. At left is the distribution for pion parents and at right is the distribution for muons parents. The pion RMS is about 3% larger and the muon RMS is about 8% larger. Unfortunately, because the flux files only contain parents that produce neutrinos, these effects may be due to the correlation between this variable and the  $V_z$  distribution and cannot demonstrate that Flugg has a wider flux leaving the horns.




Figure F.13: The ratio of integrated  $\mu^+$  Near Detector fluxes from Flugg and Gnumi in a range of beam configurations. The focused energy increases from left to right.




Figure F.14: The ratio of integrated  $\mu^-$  Near Detector fluxes from Flugg and Gnumi in a range of beam configurations. The focused energy increases from left to right.

### F.5.4 Low Angle Scatters

The divergence plots developed to study the  $\nu_e$  flux show another interesting effect. While the effect pertains to too small a number of events to explain significant spectral differences, it is representative of the types of underlying differences between the behavior two simulations (this one just happened to get tracked down). Using the *divr* variable (the extrapolated radial position at  $z = 45$  m, we find that about 5% of the events with parents produced in the target and neutrino vertices in the decay pipe have ‘impossible’ values. That is, they have values of *divr*  $> 100$  cm, which is the radius of the decay pipe walls. If we plot the *divx* and *divy* for parents that decay in the decay pipe, this class of events can be seen as a halo surrounding the decay pipe entrance (see Figure F.15). Since these particles were produced in the target, they must have passed through the central window. Since the halo begins at the radius of the decay pipe, these events seem to be from glancing scatters off the walls. What makes these events interesting is that there are twice as many of them in Flugg and they tend to produce neutrinos farther down in the decay pipe, indicating perhaps smaller scattering angles.




Figure F.15: At left the *divy* vs. *divx* positions of parents produced in the target that decay in the decay pipe is shown. The rectangle in the center is the space in which the flux can pass into the decay pipe, defined by the chase shielding. The halo surrounding this window must come from particles deflecting off the decay pipe walls since the halo begins at  $r = 100$  cm, which is the radius of the decay pipe. At right are the  $V_z$  distributions for events that appear to scatter off the walls of the decay pipe. Flugg has twice as many events and they are shifted to larger  $V_z$  values.



### F.5.5 Multiple Scattering in the Chase

Another Flugg-Gnumi difference that was found was in the amount of multiple-scattering in the air in the Chase. Specifically, there appears to be none in Gnumi and there is an appreciable amount in Flugg. The effect was discovered looking at a special run with monoenergetic pions traveling directly down the center of the horns. By looking at the neutrino production vertices, we can map out the path taken by these pions. So, by looking at  $V_x$  at a slice in  $z$ , we can compare the amount of scattering we see with the amount that is calculated for pions traveling through air. This is shown for  $z = 45$  m in Figure F.16. The calculated  $1\sigma$  width due to multiple is 1.35 cm and is marked with dashed lines on the plot. You can see that this agrees well with the distribution produced by Flugg.

Further investigation revealed that the Gnumi simulation was originally written with vacuum in the chase. When the simulation was updated to have an air filled chase, the step size in the region was never updated accordingly. So, in Gnumi these particles were taking steps so large that they had no chance for scattering to occur. Again, this difference is too small to affect the flux in a significant way; however, its resolution gives confidence in the Flugg simulation.

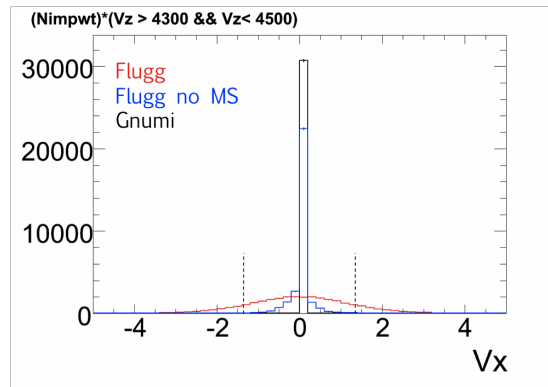


Figure F.16: This figure comes from a special run with monoenergetic, neck-to-neck pions. We use the neutrino vertices to show the path of these pions. The figure shows  $V_x$  distribution at a slice at  $z = 45$  m. You can see that while Gnumi (in black) shows no evidence of scattering, the Flugg distribution matches the calculated  $1\sigma$  scattering width (1.35 cm - marked with dashed lines) quite well.



## Appendix G

# The Water-Scintillator Detector

As neutrino and cosmic ray physics have progressed, ever larger detectors have been required to make new discoveries. Future neutrino experiments being considered, such as LBNE [157], will require near megaton-scale detectors that are sensitive to low-energy neutrino interactions. Cherenkov light detectors, such as the detector used in Super-Kamiokande [52], use water as the target medium. Water is inexpensive by weight, but the amount of Cherenkov light released in neutrino interactions requires large, expensive photodetectors to collect. It is also not sensitive to particularly low energies, with a threshold of approximately 8 MeV. Scintillator produces significantly more light (requiring less expensive photodetectors) and has a lower threshold, but requires media such as plastic (MINOS [131]) or mineral oil (NO $\nu$ A [75]) that are expensive per unit mass. The ideal solution would be a water-soluble scintillator, but unfortunately no practical scintillators of this type are known.

The Caltech neutrino group has been exploring alternatives that could take advantage of the light-producing properties of scintillator while keeping a low cost-per-weight. The basic concept is to use water as the bulk target, but to fill a small fraction of the detector volume with plastic scintillator that will react to the particles produced by interactions with the water. As a prototype,




Figure G.1: The water-scintillator detector prototype at Caltech. The green strands are WLS fiber for readout and the strands extending down from the metal cross-pieces (visible only on the right) are extruded plastic scintillator.




Figure G.2: The water-scintillator detector was read out in eight geometrical regions onto eight phototubes, arranged in light-tight boxes near the detector. The detector was triggered by muon paddles placed above and below the detector.

a  $1 \text{ m}^3$  cube steel tank was fabricated, pictured in Figure G.1. The walls of the tank were coated in highly-reflective titanium-dioxide paint and the volume was strung with 1.5 mm extruded strands of plastic scintillator before being filled with water. The scintillator strands are visible reaching down into the tank in Figure G.1. The detector was read out using wavelength-shifting (WLS) fiber to transport light from different regions of the detector to phototubes. The WLS fibers are the green strands in Figure G.1, and the phototubes are housed in the metal boxes sitting beside the detector in Figure G.2. Each of the eight phototubes receive light from a particular geometrical region of the detector, pictured in Figure G.2. The detector is triggered on cosmic ray muons using plastic scintillator paddles placed above and below the detector in an overlapping pattern, also shown in Figure G.2. The signature trigger of a through-going muon is the coincidence of two hits above and two hits below the same quadrant of the detector from the paddles.

Figure G.3 shows histograms of the responses in each of the geometrical channels when the detector was triggered above the channel 5 region. The regions closest to the triggered paddles show the largest signals, with smaller signals in more distant regions. Thus, some track localization is possible in this detector arrangement – we estimate on the order of tens of centimeters. Figure G.4 gives the summed signal across all channels for through-going muon triggers, showing an average total response of 30 PE. A typical cosmic-ray muon will deposit 150 MeV of energy while traversing 1 m of water [35]. This gives a detector response of 1 PE per 5 MeV of deposited energy, which is similar to the light levels achieved in MINOS.

The light levels and materials costs of this detector technology are promising, but further work is required, particularly on cost-effective construction methods, before it is shown to be a viable choice for a megaton scale detector.

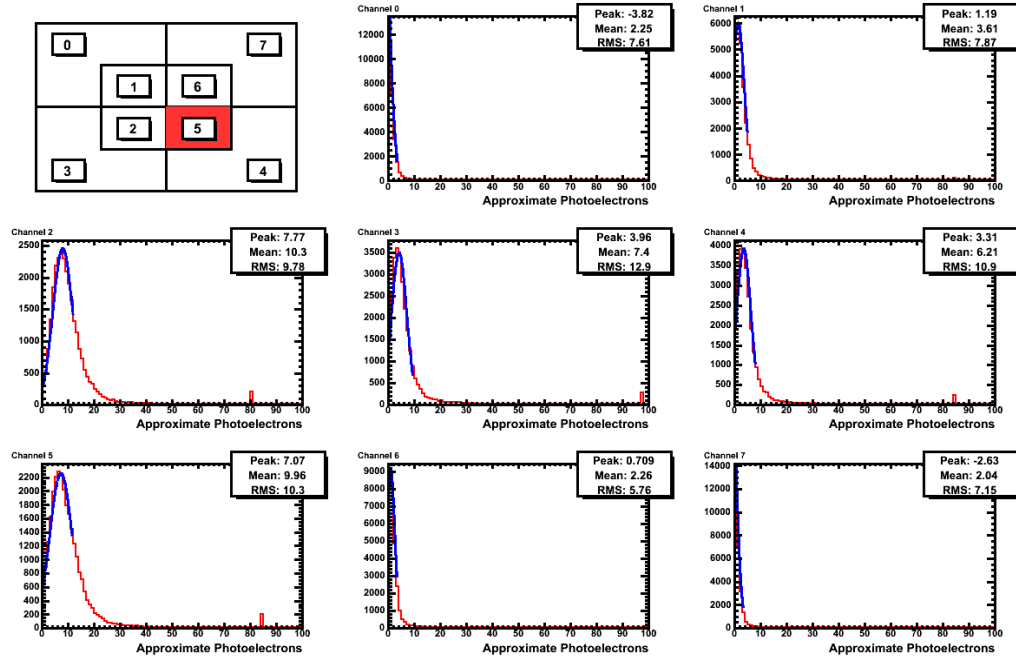


Figure G.3: Histograms of the detector response in each of the eight channels when there was a trigger over the channel 5 region. The closest regions show large signals while more distant ones show only small signals.




Figure G.4: The summed response of all channels to a through-going muon trigger. Approximately 30 PE equivalent is observed in response to the typical energy depositions of 150 MeV.



# Bibliography

- [1] J. Chadwick. Distribution in intensity in the magnetic spectrum of the  $\beta$ -rays of radium. *Ver. Dtsch. Physik. Ges.*, 16:383–391, 1914.
- [2] C. D. Ellis and W. A. Wooster. The average energy of disintegration of radium E. *Proc. of the Royal Society of London: A*, 117(776):109–123, 1927.
- [3] N. Bohr. Faraday lecture, May 8 1930.
- [4] Wolfgang Pauli. Letter to the physical society of Tübingen. 1930.
- [5] J. Chadwick. Possible existence of a neutron. *Nature*, 129:312, 1932.
- [6] E. Fermi. An attempt of a theory of beta radiation. *Z. Phys.*, 88:161–177, 1934.
- [7] Carl D. Anderson. The positive electron. *Phys. Rev.*, 43(6):491–494, 1933.
- [8] P. A. M. Dirac. The quantum theory of the electron. Part II. *Proc. of the Royal Society of London: A*, 118:351–361, 1928.
- [9] G. Gamow and E. Teller. Selection rules for the  $\beta$ -disintegration. *Phys. Rev.*, 49:895, 1936.
- [10] H. Bethe and R. Peierls. The ‘neutrino’. *Nature*, 133:532, 1934.
- [11] C. Powell, P. Fowler, and D. Perkins. The study of elementary particles by the photographic method. *Nature*, 159:694, 1947.
- [12] C. L. Cowan Jr., F. Reines, et al. Detection of the free neutrino: a confirmation. *Science*, 124:103–104, 1956.
- [13] B. Pontecorvo. The neutrino and the recoil of nuclei in beta disintegrations. *Rep. Prog. Phys.*, 11:32, 1947.
- [14] R. Davis and D. Harmer. Attempt to observe the  $^{37}\text{Cl}(\bar{\nu}, e^-)^{37}\text{Ar}$  reaction induced by reactor antineutrinos. *Bull. Am. Phys. Soc.*, 4:217, 1959.
- [15] B. Pontecorvo. Mesonium and antimesonium. *Sov. Phys. JETP*, 6:429, 1957.

- [16] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9(1):36–44, 1962.
- [17] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104(1):254–258, 1956.
- [18] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105(4):1413–1415, 1957.
- [19] Richard L. Garwin, Leon M. Lederman, and Marcel Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Phys. Rev.*, 105(4):1415–1417, 1957.
- [20] T. D. Lee and C. N. Yang. Parity nonconservation and a two-component theory of the neutrino. *Phys. Rev.*, 105(5):1671–1675, 1957.
- [21] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109(3):1015–1017, 1958.
- [22] W. B. Herrmannsfeldt, D. R. Maxson, P. Stähelin, and J. S. Allen. Electron-neutrino angular correlation in the positron decay of argon 35. *Phys. Rev.*, 107(2):641–643, 1957.
- [23] F. Boehm, T. B. Novey, C. A. Barnes, and B. Stech. Positron polarization in a mirror transition. *Phys. Rev.*, 108(6):1497–1499, 1957.
- [24] M. Deutsch, B. Gittelman, R. W. Bauer, L. Grodzins, and A. W. Sunyar. Polarized positrons from  $^{66}\text{Ga}$  and  $^{34}\text{Cl}$ . *Phys. Rev.*, 107(6):1733–1734, 1957.
- [25] S. S. Hanna and R. S. Preston. Positron polarization demonstrated by annihilation in magnetized iron. *Phys. Rev.*, 106(6):1363–1364, 1957.
- [26] L. Landau. On the conservation laws for weak interactions. *Nucl. Phys.*, 3(1):127–131, 1957.
- [27] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the  $2\pi$  decay of the  $K_2^0$  meson. *Phys. Rev. Lett.*, 13(4):138–140, 1964.
- [28] Julian Schwinger. The theory of quantized fields. II. *Phys. Rev.*, 91(3):713–728, 1953.
- [29] Gerhart Lüders and Bruno Zumino. Connection between spin and statistics. *Phys. Rev.*, 110(6):1450–1453, 1958.
- [30] W. Pauli. *Niels Bohr and the Development of Physics, Essays dedicated to Niels Bohr on the occasion of his seventieth birthday.*, chapter Exclusion principle, Lorentz group and reflection of space–time and charge. Pergamon Press, 1955.



- [31] J. S. Bell. Time reversal in field theory. *Proc. of the Royal Society of London: A*, 231(1187):479–495, 1955.
- [32] Sheldon L. Glashow. Partial-symmetries of weak interactions. *Nucl. Phys.*, 22(4):579–588, 1961.
- [33] Steven Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19(21):1264–1266, 1967.
- [34] A. Salam and J. C. Ward. Electromagnetic and weak interactions. *Phys. Lett.*, 13:168–171, 1964.
- [35] K. Nakamura and the Particle Data Group. Review of Particle Physics. *Journal of Physics G*, 37(7A):0705021, 2010.
- [36] Peter W. Higgs. Spontaneous symmetry breakdown without massless bosons. *Phys. Rev.*, 145(4):1156–1163, 1966.
- [37] M.E. Peskin and D.V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.
- [38] Samoil Bilenyk. *Introduction to the Physics of Massive and Mixed Neutrinos*. Springer-Verlag, 2010.
- [39] J. Goldstone. Field theories with “Superconductor” solutions. *Il Nuovo Cimento*, 19:154, 1961.
- [40] Nicola Cabibbo. Unitary symmetry and leptonic decays. *Phys. Rev. Lett.*, 10(12):531–533, 1963.
- [41] Makoto Kobayashi and Toshihide Maskawa.  $\mathcal{CP}$ -violation in the renormalizable theory of weak interaction. *Progress of Theoretical Physics*, 49(2):652–657, 1973.
- [42] Tadao Nakano and Kazuhiko Nishijima. Charge independence for  $v$ -particles. *Progress of Theoretical Physics*, 10(5):581–582, 1953.
- [43] M. Gell-Mann. The interpretation of the new particles as displaced charge multiplets. *Il Nuovo Cimento*, 4:848–866, 1956.
- [44] P. Adamson et al. Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS Near Detector. *Phys. Rev. D*, 81:072002, 2010.
- [45] Richard P. Feynman. Very high-energy collisions of hadrons. *Phys. Rev. Lett.*, 23(24):1415–1417, Dec 1969.
- [46] Martin Mihaylov Tzanov. *Precise Measurement of Neutrino and Anti-neutrino Differential Cross Sections on Iron*. PhD thesis, University of Pittsburgh, 2005.

- [47] John N. Bahcall. Solar Neutrinos. I. Theoretical. *Phys. Rev. Lett.*, 12(11):300–302, 1964.
- [48] Bruce T. Cleveland, Timothy Daily, Jr. Raymond Davis, et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496(1):505, 1998.
- [49] John N. Bahcall and Andrew Ulmer. Temperature dependence of solar neutrino fluxes. *Phys. Rev. D*, 53(8):4202–4210, 1996.
- [50] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81(8):1562–1567, 1998.
- [51] Q. R. Ahmad et al. Measurement of the rate of  $\nu_e + d \rightarrow p + p + e$  interactions produced by  $^8\text{B}$  solar neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87(7):071301, 2001.
- [52] S. Fukuda et al. The Super-Kamiokande detector. *Nucl. Instr. and Meth. A*, 501(2-3):418–462, 2003.
- [53] B. Aharmim et al. Measurement of the  $\nu_e$  and total B-8 solar neutrino fluxes with the Sudbury Neutrino Observatory phase I data set. *Phys. Rev. C*, 75:045502, 2007.
- [54] J. Boger et al. The sudbury neutrino observatory. *Nucl. Instr. and Meth. A*, 449(1-2):172–207, 2000.
- [55] John N. Bahcall.  $^7\text{Be}$  solar neutrino line: A reflection of the central temperature distribution of the Sun. *Phys. Rev. D*, 49(8):3923–3945, 1994.
- [56] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962.
- [57] M. Grunewald et al. Precision electroweak measurements on the  $Z$  resonance. *Phys. Rept.*, 427(5-6):257–454, 2006.
- [58] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17(9):2369–2374, 1978.
- [59] S. Mikheyev and A. Smirnov. Resonant amplification of  $\nu$  oscillations in matter and solar-neutrino spectroscopy. *Il Nuovo Cimento C*, 9:17–26, 1986.
- [60] K. S. Hirata et al. Real time, directional measurement of  $^8\text{B}$  solar neutrinos in the Kamiokande-II detector. *Phys. Rev. D*, 44:2241–2260, 1991.
- [61] J. N. Abdurashitov et al. Solar neutrino flux measurements by the Soviet-American gallium experiment (SAGE) for half the 22-year solar cycle. *Journal of Experimental and Theoretical Physics*, 95(2):181–193, 2002.

- [62] W. Hampel et al. GALLEX solar neutrino observations: Results for GALLEX IV. *Phys. Lett. B*, 447:127–133, 1999.
- [63] B. Aharmim et al. Low energy threshold analysis of the phase I and phase II data sets of the Sudbury Neutrino Observatory. *Phys. Rev. C*, 81:055504, 2010.
- [64] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of  $^8\text{B}$  solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys. Rev. C*, 72(5):055502, 2005.
- [65] S. Abe et al. Precision measurement of neutrino oscillation parameters with KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.
- [66] R. Wendell et al. Atmospheric neutrino oscillation analysis with subleading effects in Super-Kamiokande I, II, and III. *Phys. Rev. D*, 81(9):092004, 2010.
- [67] P. Adamson et al. Measurement of the neutrino mass splitting and flavor mixing by MINOS. *Phys. Rev. Lett.*, 106:181801, 2011.
- [68] K. Abe et al. Search for matter-dependent atmospheric neutrino oscillations in Super-Kamiokande. *Phys. Rev. D*, 77(5):052001, 2008.
- [69] M. H. Ahn et al. Measurement of neutrino oscillation by the K2K experiment. *Phys. Rev. D*, 74(7):072003, 2006.
- [70] Markus A. Luty. Baryogenesis via leptogenesis. *Phys. Rev. D*, 45(2):455–465, Jan 1992.
- [71] M. Goodman and T. Lasserre. Double Chooz, a search for the neutrino mixing angle theta-13. arXiv:hep-ex/0606.025v4, 2006.
- [72] Daya Bay Collaboration. A precision measurement of the neutrino mixing angle theta-13 using reactor antineutrinos at daya bay. arXiv:hep-ex/0701.029v1, BNL-77369-2006-IR, LBNL-62137, TUHEP-EX-06-003, 2006.
- [73] S-B Kim and The RENO Collaboration. Reno: reactor experiment for neutrino oscillation at Yonggwang. *Journal of Physics: Conference Series*, 120(5):052025, 2008.
- [74] Y. Hayato. T2K at J-PARC. *Nucl. Phys. B - Proceedings Supplements*, 143:269–276, 2005. NEUTRINO 2004.
- [75] D. Ayres and The NO $\nu$ A Collaboration. NO $\nu$ A proposal to build a 30 kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline. arXiv:hep-ex/0503.053v1, Fermilab-Proposal-0929, 2005.

- [76] M.C. Gonzalez-Garcia and M. Maltoni. Phenomenology with massive neutrinos. *Phys. Rept.*, 460(1-3):1–129, 2008.
- [77] P. Adamson et al. Measurement of neutrino oscillations with the MINOS detectors in the NuMI beam. *Phys. Rev. Lett.*, 101:131802, 2008.
- [78] Julian Heeck and Werner Rodejohann. Gauged  $L_\mu - L_\tau$  and different muon neutrino and anti-neutrino oscillations: MINOS and beyond. arXiv:1007.2655v2, 2010.
- [79] Joachim Kopp, Pedro A. N. Machado, and Stephen J. Parke. Interpretation of MINOS data in terms of nonstandard neutrino interactions. *Phys. Rev. D*, 82(11):113002, 2010.
- [80] N. Engelhardt, A. E. Nelson, and J. R. Walsh. Apparent  $CPT$  violation in neutrino oscillation experiments. *Phys. Rev. D*, 81(11):113001, 2010.
- [81] Mattias Blennow, Tommy Ohlsson, and Julian Skrotzki. Effects of non-standard interactions in the MINOS experiment. *Phys. Lett. B*, 660(5):522–528, 2008.
- [82] Carla Biggio, Mattias Blennow, and Enrique Fernández-Martínez. General bounds on non-standard neutrino interactions. *Journal of High Energy Physics*, 2009(08):090, 2009.
- [83] W. Anthony Mann, Daniel Cherdack, Wojciech Musial, and Tomas Kafka. Apparent multiple  $\Delta m_{32}^2$  in  $\bar{\nu}_\mu$  and  $\nu_\mu$  survival oscillations from nonstandard interaction matter effect. *Phys. Rev. D*, 82(11):113010, 2010.
- [84] O. W. Greenberg.  $CPT$  violation implies violation of Lorentz invariance. *Phys. Rev. Lett.*, 89(23):231602, 2002.
- [85] Steven Weinberg. *The Quantum Theory of Fields, Vol. I*. Cambridge University Press, 1995.
- [86] C. Athanassopoulos et al. Evidence for  $\nu_\mu \rightarrow \nu_e$  neutrino oscillations from LSND. *Phys. Rev. Lett.*, 81:1774–1784, 1997.
- [87] H. Murayama and T. Yanagida. LSND, SN1987A, and  $CPT$  violation. *Phys. Lett. B*, 520(3-4):263–268, 2001.
- [88] Gabriela Barenboim, Liubo Borisso, Joseph Lykken, and Alexei Smirnov. Neutrinos as the messengers of  $CPT$  violation. *Journal of High Energy Physics*, 2002(10):001, 2002.
- [89] G. Barenboim, L. Borisso, and J. Lykken. Neutrinos that violate  $CPT$ , and the experiments that love them. *Phys. Lett. B*, 534(1-4):106–113, 2002.
- [90] Gabriela Barenboim and Joseph Lykken. A model of  $CPT$  violation for neutrinos. *Phys. Lett. B*, 554(1-2):73–80, 2003.

- [91] Irina Mocioiu and Maxim Pospelov. Indirect limits on the  $CPT$  violating background in the neutrino sector. *Phys. Lett. B*, 534(1-4):114–119, 2002.
- [92] André de Gouvêa. Can a  $CPT$  violating ether solve all electron (anti)neutrino puzzles? *Phys. Rev. D*, 66(7):076005, 2002.
- [93] Gabriela Barenboim and Joseph D. Lykken. MINOS and  $CPT$ -violating neutrinos. *Phys. Rev. D*, 80(11):113008, 2009.
- [94] J. S. Díaz, V. A. Kostelecký, and M. Mewes. Perturbative Lorentz and  $CPT$  violation for neutrino and antineutrino oscillations. *Phys. Rev. D*, 80(7):076007, 2009.
- [95] J. S. Díaz and V. A. Kostelecký. A simple Lorentz-violating texture for neutrino mixing. arXiv:hep-ph/1012.5985v1, 2010.
- [96] Antonio Di Domenico et al.  $CPT$  symmetry and quantum mechanics tests in the neutral kaon system at KLOE. *Foundations of Physics*, 40:852–866, 2010. 10.1007/s10701-009-9366-x.
- [97] J. M. Link et al. Charm system tests of  $CPT$  and Lorentz invariance with FOCUS. *Phys. Lett. B*, 556:7–13, 2003.
- [98] B. Aubert et al. Search for  $CPT$  and Lorentz violation in  $B_0$ - $B_0$ bar oscillations with dilepton events. *Phys. Rev. Lett.*, 100:131802, 2008.
- [99] C. Gemmel et al. Limit on Lorentz and  $CPT$  violation of the bound neutron using a free precession  $^3\text{He}/^{129}\text{Xe}$  comagnetometer. *Phys. Rev. D*, 82(11):111901, 2010.
- [100] G. W. Bennett et al. Search for Lorentz and  $CPT$  violation effects in muon spin precession. *Phys. Rev. Lett.*, 100(9):091602, 2008.
- [101] H. Dehmelt, R. Mittleman, R. S. Van Dyck, and P. Schwinberg. Past electron-positron  $g - 2$  experiments yielded sharpest bound on  $CPT$  violation for point particles. *Phys. Rev. Lett.*, 83(23):4694–4696, 1999.
- [102] P. Adamson et al. A search for Lorentz invariance and  $CPT$  violation with the MINOS Far Detector. *Phys. Rev. Lett.*, 105:151601, 2010.
- [103] P. Adamson et al. Testing Lorentz invariance and  $CPT$  conservation with NuMI neutrinos in the MINOS Near Detector. *Phys. Rev. Lett.*, 101:151601, 2008.
- [104] John N. Bahcall and Henry Primakoff. Neutrino-antineutrino oscillations. *Phys. Rev. D*, 18(9):3463–3466, 1978.
- [105] J. Schechter and J. W. F. Valle. Neutrino-oscillation thought experiment. *Phys. Rev. D*, 23(7):1666–1668, 1981.

- [106] Paul Langacker and Jing Wang. Neutrino-antineutrino transitions. *Phys. Rev. D*, 58(9):093004, 1998.
- [107] André de Gouvêa, Boris Kayser, and Rabindra N. Mohapatra. Manifest  $\mathcal{CP}$  violation from Majorana phases. *Phys. Rev. D*, 67(5):053004, 2003.
- [108] M. A. B. Bégin, R. V. Budny, R. Mohapatra, and A. Sirlin. Manifest left-right symmetry and its experimental consequences. *Phys. Rev. Lett.*, 38(22):1252–1255, 1977.
- [109] A. Jodidio et al. Search for right-handed currents in muon decay. *Phys. Rev. D*, 34(7):1967–1990, 1986.
- [110] F. Bergsma et al. Bounds on supersymmetric particles from a proton beam-dump experiment. *Phys. Lett. B*, 121(6):429–432, 1983.
- [111] S. R. Mishra et al. Search for right handed coupling in  $\nu$ -N scattering. *Phys. Rev. Lett.*, 68(24):3499–3502, 1992.
- [112] H. Abramowicz et al. Experimental study of opposite-sign dimuons produced in neutrino and antineutrino interactions. *Z. Phys. C*, 15:19–31, 1982.
- [113] T. P. Cheng and Wu-Ki Tung. General local interactions and tests of  $V - A$  theory in neutrino scattering processes. *Phys. Rev. D*, 3(3):733–744, 1971.
- [114] M. Jonker et al. Measurement of the polarization of positive muons produced in high-energy antineutrino interactions. *Z. Phys. C*, 17:211–217, 1983.
- [115] A. M. Cooper, J. G. Guy, A. G. Michette, M. Tyndel, and W. Venus. Limits on neutrino-antineutrino transitions from a study of high-energy neutrino interactions. *Phys. Lett. B*, 112(1):97–99, 1982.
- [116] S. Mishra et al. A study of wrong-sign single muon production in  $\nu_\mu$ -nucleon interaction. *Z. Phys. C*, 44:187–198, 1989.
- [117] David Delepine, Vannia González Macías, Shaaban Khalil, and Gabriel López Castro. Probing majorana neutrino  $\mathcal{CP}$  phases and masses in neutrino-antineutrino conversion. *Phys. Lett. B*, 693(4):438–442, 2010.
- [118] J. Hylen et al. Conceptual design for the technical components of the neutrino beam for the Main Injector (NuMI). Fermilab-TM-2018, 1997.
- [119] S. van der Meer. A directive device for charged particles and its use in an enhanced neutrino beam. Cern Report CERN-61-07, 1961.

- [120] S. Shukla, J. Marriner, and J. Griffin. Slip stacking in the Fermilab Main Injector. In *PF / DPB Summer Study on New Directions for High-Energy Physics (Snowmass 96)*, Snowmass, Colorado, 25 Jun - 12 Jul 1996.
- [121] R. R. Wilson. The Tevatron. *Physics Today*, 30N10:23–30, 1977.
- [122] C. C. Jensen and G. A. Krafczyk. Numi proton kicker extraction system. In *Proceedings of the 2005 IEEE Particle Accelerator Conference (PAC 05)*, 21st IEEE Particle Accelerator Conference, page 692, Knoxville, Tennessee, 16-20 May 2005.
- [123] A. G. Abramov et al. Beam optics and target conceptual designs for the NuMI project. *Nucl. Instr. and Meth. A*, 485:209–227, 2002.
- [124] C. Alt et al. Inclusive production of charged pions in p+C collisions at 158 GeV/c beam momentum. *Eur. Phys. J. C*, 49:897–917, 2006.
- [125] M. Campanella, A. Ferrari, P.R. Sala, and S. Vanini. First calorimeter simulation with the FLUGG prototype. ATL-SOFT-99-004, 10 December 1999.
- [126] M. Campanella, A. Ferrari, P.R. Sala, and S. Vanini. Reusing code from FLUKA and GEANT4 geometry. ATL-SOFT-98-039, 13 October 1998.
- [127] G. Battistoni, F. Cerutti, A. Fasso, A. Ferrari, S. Muraro, J. Ranft, S. Roesler, and P. R. Sala. The FLUKA code: description and benchmarking. *AIP Conf. Proc.*, 896:31–49, 2007.
- [128] S. Agostinelli et al. G4—a simulation toolkit. *Nucl. Instr. and Meth. A*, 506(3):250 – 303, 2003.
- [129] A. Fassò, A. Ferrari, J. Ranft, and P.R. Sala. FLUKA: a multi-particle transport code. CERN-2005-10, , INFN/TC 05/11, SLAC-R-773, 2005.
- [130] A. Fassò et al. The physics models of FLUKA: status and recent developments. In *Computing in High Energy and Nuclear Physics (CHEP2003)*, La Jolla, CA, USA, March 24-28 2003. (paper MOMT005), eConf C0303241 (2003), arXiv:hep-ph/0306267.
- [131] D.G. Michael et al. The magnetized steel and scintillator calorimeters of the MINOS experiment. *Nucl. Instr. and Meth. A*, 596(2):190 – 228, 2008.
- [132] J. Oliver et al. Design and performance of the readout system of the MINOS Far Detector. *IEEE Trans. Nucl. Sci.*, 51:2193–2195, 2004.
- [133] T. Cundiff et al. The MINOS Near Detector front end electronics. *IEEE Trans. Nucl. Sci.*, 53:1347–1355, 2006.
- [134] C. Zeitnitz and T.A. Gabriel. The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters. *Nucl. Instr. and Meth. A*, 349(1):106–111, 1994.

- [135] P. Adamson et al. A study of muon neutrino disappearance using the Fermilab Main Injector neutrino beam. *Phys. Rev. D*, 77:072002, 2008.
- [136] P. Adamson et al. New constraints on muon-neutrino to electron-neutrino transitions in MINOS. *Phys. Rev. D*, 82:051102, 2010.
- [137] P. Adamson et al. Search for sterile neutrino mixing in the MINOS long-baseline experiment. *Phys. Rev. D*, 81:052004, 2010.
- [138] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering, Series D*, 82:35–45, 1960.
- [139] D. E. Groom, N. V. Mokhov, and S. I. Striganov. Muon stopping power and range tables 10-MeV to 100-TeV. *Atomic Data and Nuclear Data Tables*, 78:183–356, 2001.
- [140] R. Brun et al. Geant detector description and simulation tool. CERN Program Library Long Writeup W5013, 1994.
- [141] R. M. Sternheimer and R. F. Peierls. General expression for the density effect for the ionization loss of charged particles. *Phys. Rev. B*, 3(11):3681–3692, 1971.
- [142] R.M. Sternheimer, M.J. Berger, and S.M. Seltzer. Density effect for the ionization loss of charged particles in various substances. *Atomic Data and Nuclear Data Tables*, 30(2):261–271, 1984.
- [143] P. Adamson et al. The MINOS calibration detector. *Nucl. Instr. and Meth. A*, 556(1):119–133, 2006.
- [144] H. Gallagher. The NEUGEN neutrino event generator. *Nucl. Phys. B*, 112(1-3):188–194, 2002.
- [145] T. Yang et al. A hadronization model for few-GeV neutrino interactions. arXiv:hep-ph/0904.4043, 2009.
- [146] T. Sjostrand, S. Mrenna, and P. Skands. PYTHIA 6.4 physics and manual. *Journal of High Energy Physics*, 05:026, 2006.
- [147] Z. Koba, H. B. Nielsen, and P. Olesen. Scaling of multiplicity distributions in high-energy hadron collisions. *Nucl. Phys. B*, 40:317–334, 1972.
- [148] R. Merenyi et al. Determination of pion intranuclear rescattering rates in  $\nu_\mu$ - $Ne$  versus  $\nu_\mu$ - $D$  interactions for the atmospheric  $\nu$  flux. *Phys. Rev. D*, 45(3):743–751, 1992.
- [149] The MINOS Collaboration. Calibration position paper. minos-doc-3137, July 2007.



- 
- [150] S. Dytman, H. Gallagher, and M. Kordosky. Shower energy scale uncertainty for the run I+II CC analysis. *minos-doc-4287*, 2008.
- [151] Gary J. Feldman and Robert D. Cousins. A unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.
- [152] T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Trans. Information Theory*, 13(1):21 – 27, 1967.
- [153] The MINOS Collaboration. Run IV calibration position paper. *minos-doc-3137*, May 2010.
- [154] J. J. Evans. Absolute (calorimetric) shower energy uncertainty. *minos-doc-7173*, 2010.
- [155] David Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, 2004.
- [156] A. Fassò, A. Ferrari, J. Ranft, and P.R. Sala. FLUKA: present status and future developments. In A. Menzione and A. Scribano, editors, *IV Int. Conf. on Calorimetry in High Energy Physics*, page 493, La Biodola, Italy, 21-26 Sep 1993. World Scientific.
- [157] V. Barger et al. Report of the US long baseline neutrino experiment study. *arXiv:hep-ex/0705.4396v1*, Fermilab-0801-AD-E, BNL-77973-2007-IR, 2007.