

OBSERVATION OF ELECTRON NEUTRINO
APPEARANCE IN THE NUMI BEAM WITH
THE NOVA EXPERIMENT

Evan Niner

Submitted to the faculty of the University Graduate School

in partial fulfillment of the requirements

for the degree

Doctor of Philosophy

in the Department of Physics,

Indiana University

August, 2015

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the
requirements for the degree of Doctor of Philosophy.

Doctoral Committee

Mark Messier, Ph.D. (Chair)

Jon Urheim, Ph.D.

Rex Tayloe, Ph.D.

Herb Fertig, Ph.D.

August 10, 2015

Copyright © 2015

Evan Niner

To my parents, whose unconditional love, support, and encouragement to dream for the past 28 years has made this work possible.

To Jackie, who has stood by my side through this thesis and always, onward to the next adventure.

ACKNOWLEDGMENTS

More people than I can possibly name have played a role in my life that has culminated with the completion of this thesis. Thanks to all who have been a part of this journey, I am forever grateful.

I would like to start by thanking the two people most responsible for shaping my career as a physicist. Andreas Piepke introduced me to neutrino physics 9 years ago as a freshman at the University of Alabama through work on the EXO experiment. I am thankful for your mentorship and guidance and providing me with opportunities and responsibilities I would not have thought myself capable. Second, I cannot begin to express my thanks to Mark Messier, my thesis adviser these past five years. Beyond teaching me many of the finer points of experimentation and statistics you have been a friend. This thesis would not have been possible without the countless hours of discussion and support.

Thanks to the rest of my thesis committee: Jon Urheim, Rex Tayloe, and Herb Fertig for reading this document. Thanks also to the other faculty of the I.U. NOvA group Jim Musser and Stuart Mufson for your support.

I would like to thank the entire NOvA collaboration, without whose efforts the experimental data would not have existed to write this thesis. The ν_e appearance analysis was an ensemble effort and I would like to thank the group conveners Patricia Vahle and Mayly Sanchez for their guidance. Other members of the group whose discussions, advice, and contributions led to this work including Jianming Bian, Chris Backhouse, Ji Liu, Tian Xin, Joe Lozier, Dan Pershey, Erika Catano-Mur, Alex Radovic, Nitin Yadav, and Hongyue

Duyang. Thanks to Kanika Sachdev for the innumerable conversations/rants during the analysis. Thanks to Xuebing Bu for a heroic effort in leading the day-to-day construction of the Near Detector and allowing me to contribute. A big thanks to Peter Shanahan and Andrew Norman for providing me with opportunities in the DAQ group. Thanks also go to Brian Rebel, Tom Coen, Jon Paley, Mat Muether, Nate Mayer and Gary Feldman. I would like to thank my fellow graduate students Jan Zirnstien, Nick Raddatz, Dominick Rocco, Gareth Kafka, Fernanda Pshihas, Justin Vassel, Ryan Murphy, Jose Sepulveda-Quiroz, and Enrique Arrieta-Diaz for fellowship and beers both at I.U. and the collaboration meetings.

Thanks to many graduate students at I.U. for the friendship, late-night problem solving sessions, late-night venting sessions, and late-night trips for food (R.I.P Happy Pig) and drink: Manny, Jorge, Tessa, Tracy, Aparajita, Hung Chun, Ben, Bo, Sean, Jeff, and $\widehat{\text{Evan}}$. A particular shout-out to my friend and office mate Michael Baird, a cat who won't cop out when there's danger all about (right on). For five years we pushed each other to the thesis finish line, I couldn't imagine it any other way. Thanks also to Gavin Davies, a post doc, for whose tireless efforts to produce the data files, physics advice, and friendship I am grateful. A big thank you to another pot doc in the group, Denver Whittington.

Above all, I would like to thank my family. My parents, Holly and Keith, and sister Beth have been a never ending source of support and encouragement throughout my life. Thanks for always attempting to answer my unending questions of "why" growing up and supporting my passion for science. I will spend the rest of my life thanking my fiance, Jackie, who has been steadfast during this journey. This work would not be possible without you.

Evan Niner

OBSERVATION OF ELECTRON NEUTRINO APPEARANCE IN THE NUMI BEAM
WITH THE NOVA EXPERIMENT

NOvA is a long-baseline neutrino oscillation experiment that uses two functionally identical detectors separated by 810 kilometers at locations 14 milliradians off-axis from the NuMI muon neutrino beam at Fermilab. At these locations the beam energy peaks at 2 GeV. This baseline is the longest in the world for an accelerator-based neutrino oscillation experiment, which enhances the sensitivity to the neutrino mass ordering. The experiment studies oscillations of the muon neutrino and anti-neutrino beam that is produced. Both detectors completed commissioning in the summer of 2014 and continue to collect data. One of the primary physics goals of the experiment is the measurement of electron neutrino appearance in the muon neutrino beam which yields measurements of the oscillation parameters $\sin^2 2\theta_{13}$, δ , and the neutrino mass ordering within the standard model of neutrino oscillations. This thesis presents the analysis of data collected between February 2014 and May 2015, corresponding to 3.52×10^{20} protons-on-target. In this first analysis NOvA recorded 6 electron neutrino candidates which is a 3.3σ observation of electron neutrino appearance. The T2K experiment performs the same measurement on a baseline of 295 kilometers and has a 1σ preference for the normal mass ordering over the inverted ordering over the phase space of the CP violating parameter δ , which is also weakly seen in the NOvA result. By the summer of 2016 NOvA will triple its statistics due to increased beam power and a completed detector. If electron neutrinos continue to be observed at the current rate NOvA will be able to establish a mass ordering preference at a similar confidence level to T2K.

Mark Messier, Ph.D. (Chair)

Jon Urheim, Ph.D.

Rex Tayloe, Ph.D.

Herb Fertig, Ph.D.

CONTENTS

1	Introduction	1
2	Neutrino Oscillation Theory and Experimental Landscape	6
2.1	Neutrino Mass	7
2.2	Oscillation Formalism	10
2.3	PMNS Matrix	13
2.4	ν_e Appearance Oscillation Channel	14
2.5	Matter Effect	16
2.6	Atmospheric Neutrino Oscillations	19
2.6.1	Super-Kamiokande	19
2.6.2	Accelerator Experiments	21
2.7	Solar Neutrino Oscillations	24
2.7.1	SNO	24
2.7.2	KamLAND	26
2.8	Measuring θ_{13}	27
2.8.1	Reactor Experiments	27
2.8.2	θ_{13} at Accelerator Experiments	29
2.9	Summary of Results	31
2.9.1	Three Flavor Analysis of MINOS and T2K	31
2.9.2	Experimental State of the PMNS Matrix	35

2.9.3	Unanswered Questions in Neutrino Physics	36
2.10	Measuring ν_e Appearance with NOvA	37
3	The NOvA Experiment	39
3.1	NuMI Beam	39
3.1.1	Beam Flux and Kinematics	41
3.1.2	NuMI Upgrade	45
3.2	The NOvA Detector Design	46
3.3	The Far Detector	50
3.4	The Near Detector	50
3.5	Energy Calibration	54
3.6	Monte Carlo Simulation	55
4	Timing System and Calibration	58
4.1	Timing System Requirements	59
4.2	Timing System Design	59
4.3	NOvA Time Specification	61
4.4	Timing Synchronization	62
4.5	NuMI Beam Triggering System	65
4.6	Timing Resolution	67
4.7	Offline Timing Calibration	69
4.7.1	Track Selection	71
4.7.2	Calibration Procedure	72
4.7.3	Timing Calibration Results	77
4.7.4	Additional Correction Factors	80

5	Event Reconstruction	85
5.1	Event Topologies	85
5.2	Reconstruction Philosophy	87
5.3	Interaction Separation with DBSCAN	88
5.4	Guidelines with Hough Transform	90
5.5	Vertex Identification with ElasticArms	94
5.5.1	Energy Function	96
5.5.2	Seeding and Minimization	97
5.5.3	Performance of Vertexing	99
5.5.4	Improvements to Vertexing	103
5.6	Prong Formation with Fuzzy K-Means	106
5.6.1	Classic Fuzzy K-Means	108
5.6.2	Possibilistic Fuzzy K-Means	109
5.6.3	Seeding	111
5.6.4	Optimizing Prongs and Cleanup	112
5.6.5	Forming 3D Prongs	114
5.6.6	Performance of Fuzzy K-Means Algorithm	115
5.6.7	Improvements to Prong Formation	118
5.7	Event Classification	125
5.7.1	Hit Redistribution	126
5.7.2	dE/dx and Log-Likelihood Calculations	127
5.7.3	Neural Net Input Variables	130
5.7.4	Artificial Neural Network and Training	133
5.7.5	Performance of Event Classification	136
5.7.6	Improvements to Event Selection	137

6	Event Selection	138
6.1	Analysis Period	139
6.2	Quality Run Selection	140
6.2.1	Far Detector	141
6.2.2	Near Detector	143
6.3	Far Detector Event Selection	144
6.4	Near Detector Event Selection	146
6.5	Far Detector Effective Fiducial Mass	154
6.6	Analysis Exposure	157
7	Electron Neutrino Appearance Analysis	159
7.1	Extrapolation	159
7.2	Sideband Studies	163
7.2.1	Cosmic Bremsstrahlung Efficiency	163
7.2.2	High Energy Neutrino Selection	166
7.2.3	Low LID Sideband	170
7.3	Systematic Uncertainties	170
7.3.1	Beam	171
7.3.2	Birk's Law	172
7.3.3	Calibration	173
7.3.4	Light Levels	174
7.3.5	Neutrino Interaction	174
7.3.6	Near Detector Containment	175
7.3.7	Rock Contamination	176
7.3.8	Decomposition	176
7.3.9	Data and Monte Carlo Statistics	177

7.3.10	Normalization	177
7.3.11	Alignment	178
7.3.12	Bad Channels	178
7.3.13	Systematic Error Table	178
7.4	Prediction	178
8	Results	181
8.1	Statistical treatment of PMNS Parameter Measurements	181
8.2	Far Detector Data Observation	183
8.3	Confidence Intervals	184
8.3.1	Discussion	187
9	Conclusions	192
	Bibliography	194
	CV	

LIST OF FIGURES

2.1	Feynman diagram for $\pi^+ \rightarrow \mu^+ + \nu_\mu$ illustrating the neutrino produced in a definite electro-weak state matching the lepton.	7
2.2	Diagram showing the two possible mass orderings with each state showing the approximate measured flavor combination [29].	14
2.3	Coherent forward scattering for neutrinos on electrons in matter. On left, an electron neutrino scattering via a charged current interaction. On right, any neutrino flavor scattering via a neutral current interaction.	16
2.4	Zenith angle distributions for e-like and μ -like events in Super-Kamiokande with visible energy < 1.33 GeV (sub-GeV) and > 1.33 GeV (multi-GeV). The dotted line shows the un-oscillated Monte Carlo prediction and the solid line is the best-fit under the two flavor oscillation hypothesis. Figure from [28].	20
2.5	The ratio of μ -like events observed in Super-Kamiokande data to the un-oscillated Monte Carlo prediction in bins of L/E . The solid line indicates the best fit under the two flavor oscillation hypothesis. Figure from [34]. . .	20
2.6	In the $\sin^2 2\theta_{23}$, Δm_{32}^2 parameter space the 90% C.L. allowed regions are shown the T2K 2011 [36], T2K 2013 [37], Super-Kamiokande [38], and MINOS [39]. Results shown are from both two and three flavor oscillation fits. The MINOS result assumed the same neutrino and anti-neutrino oscillation parameters. Figure was taken from [36].	21

2.7	The 90% C.L. of the allowed region from a two flavor oscillation fit to the atmospheric mixing angle and mass splitting in MINOS data. The fit was performed both assuming identical neutrino and anti-neutrino oscillations and fitting each sample separately. Figure was taken from [39].	23
2.8	Measurement of the solar neutrino flux from SNO. The x-axis is the flux from electron neutrinos and the y-axis is the flux from muon and tau flavors. The bands indicate the constraints from the different interactions. The grey band is the elastic scattering constraint from Super-Kamiokande. The best fit is shown with 68%, 95%, and 99% C.L. and is in agreement with the Standard model prediction shown with dotted lines. Figure was taken from [15]. . .	25
2.9	a) The ratio of observed $\bar{\nu}_e$ events in KamLAND to the unoscillated prediction (survival probability) as a function of L/E with $L_0 = 180$ km clearly showing the oscillation dip. The blue-line is the best fit to the two-flavor oscillation hypothesis. b) The measured solar neutrino oscillation parameters from KamLAND (colored contours) and SNO (black lines). The two results show agreement, providing complementary information. Both figures from [45].	26
2.10	Ratio of the measured to predicted event rate assuming no oscillations for each detector in Daya Bay. The red curve shows the best fit oscillation curve, inset is the χ^2 vs $\sin^2 2\theta_{13}$. Figure was taken from the first published results in 2012 [47].	28
2.11	Summary of measurements of $\sin^2 2\theta_{13}$ between 2011 and 2013 for both reactor and accelerator experiments (assuming $\delta = 0^\circ$ and $\theta_{23} = 45^\circ$). Daya Bay provides the world-leading constraint on the measurement. Figure was taken from [56].	30

- 2.12 a) T2K 68% and 90% C.L. allowed regions for $\sin^2 2\theta_{13}$ assuming normal (top) or inverted mass ordering as a function of δ . Constraints on $\sin^2 \theta_{23}$ and Δm_{32}^2 come from [37] and the shaded region shows the PDG2012 best-fit $\sin^2 2\theta_{13}$ from the reactor experiments [62]. Image from [59]. b) MINOS 68% and 90% C.L. regions for $2 \sin^2 2\theta_{13} \sin^2 \theta_{23}$ depending on mass ordering choice, δ , and θ_{23} octant. Previous measurements are used to place constraints in the fit, $\sin^2 2\theta_{23} = 0.957_{-0.046}^{+0.035}$ and $|\Delta m_{32}^2| = (2.39_{-0.10}^{+0.09}) \times 10^{-3} \text{eV}^2$. Image from [61]. 32
- 2.13 In the $(\sin^2 \theta_{23}, \Delta m_{32}^2)$ parameter space the 68% (dotted) and 90% (solid) C.L. regions for the normal (top) and inverted (bottom) mass ordering when $\sin^2 2\theta_{13}$ is constrained by the reactor data. The T2K fit is compared to Super Kamiokande [65] and MINOS [63], figure taken from [64]. 33
- 2.14 The $-2\Delta\ln(L)$ as a function of δ after marginalizing over $\sin^2 2\theta_{13}$, $\sin^2 \theta_{23}$, and Δm_{32}^2 for different mass ordering and θ_{23} octant choices. The reactor data provided a constraint in each fit of $\sin^2 2\theta_{13}$. a) T2K fit where the solid (dotted) line with markers is the Feldman-Cousins 90% C.L. limits for the normal (inverted) mass ordering. The constrain on θ_{23} comes from the T2K muon neutrino disappearance measurement [64]. Figure from [59]. b) The MINOS fit where the constraint on θ_{23} comes from the MINOS muon neutrino disappearance measurement [39]. Figure from [63]. 34

2.15	Measured oscillation probabilities in the NOvA experiment for $P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_e}$ vs $P_{\nu_\mu \rightarrow \nu_e}$ assuming 18×10^{20} POT in each neutrino and anti-neutrino running. The central starred point indicates a benchmark value assuming no matter effect, maximal θ_{23} , no CP violation, and $\sin^2 2\theta_{13} = 0.095$. The matter effect splits the oscillation in red (inverted mass ordering) and blue (normal mass ordering) ellipses dependent on the δ phase. The probabilities are further split depending on the θ_{23} octant, with values drawn assuming $\sin^2 2\theta_{23} = 0.97$. The starred point with 1 and 2 σ contours illustrates that mass ordering, δ and octant can be measured simultaneously through these two oscillation channels.	38
3.1	Schematic of the Fermilab accelerator complex [70].	40
3.2	Schematic of the NuMI beam [70].	41
3.3	a)The neutrino flux as function of pion energy viewed from an angle θ with respect to the beam axis. The flux was normalized to 810 km. b)The energy of neutrinos produced at an angle θ relative to the pion direction as a function of pion energy. To select the neutrino energy that maximizes the oscillation of ν_μ 's, NOvA chose a 14 milliradian angle drawn as a horizontal dashed line. Figure from [70].	42
3.4	The top panel depicts the charged-current ν_μ event rates prior to oscillation 810 km from Ferimlab for a detector located at various off-axis angles. The bottom panel depicts the oscillation probability $\nu_\mu \rightarrow \nu_e$ as a function of neutrino energy for that distance [70].	42

3.5	Simulated energy distributions for the ν_e oscillation signal, intrinsic beam ν_e events, neutral-current events and ν_μ charged-current events with and without oscillations. No selection efficiencies or mis-identification rates are taken into account. The simulation used $\delta m_{32}^2 = 2.5 \times 10^{-3} \text{eV}^2$, $\sin^2(2\theta_{13}) = 0.10$. An off-axis angle of 14 milliradians and oscillation distance of 810 km was assumed [70].	43
3.6	Flux times cross-section projections for the NOvA Far Detector (a) and Near Detector (b) of the ν_μ , $\bar{\nu}_\mu$, and $(\nu_e + \bar{\nu}_e)$ components of the beam in forward horn current (neutrino dominant) mode. Oscillation weights have not been applied. Figure from [72].	44
3.7	Flux times cross-section projections for the NOvA Far Detector (a) and Near Detector (b) of the ν_μ , $\bar{\nu}_\mu$, and $(\nu_e + \bar{\nu}_e)$ components of the beam in reverse horn current (anti-neutrino dominant) mode. Oscillation weights have not been applied. Figure from [72].	44
3.8	Timing schematic for creating a NuMI neutrino spill after the beamline upgrade is complete. Figure from [70].	46
3.9	Cutaway diagram of the orthogonal planes of cells in the detector. Not pictured is the 0.7 mm wavelength shifting fiber inside each cell. In reality the cells have rounded edges. Figure from [73].	47
3.10	a) The ends of 32 wavelength-shifting fibers collected at the end of scintillation cells to mount to an APD. b) Front face of an APD that will be pressed against the fiber ends.	48
3.11	The quantum efficiency of APDs (magenta) and PMTs (blue) are shown in comparison to the spectrum of light emitted from the wavelength shifting fiber at distances between 0.5 and 17 meters. Figure from [70].	49

3.12	Pictured at the Far Detector are the tops of vertical 32 cell modules (A), with the end of the horizontal planes visible (B). FEBs (C) are mounted on each module with an APD housed inside. The FEBs transmit data to DCMs (D) and are powered by power supplies (E).	51
3.13	a) Average detector response to muons at the far end of cells in the Far Detector, satisfying the design requirement of 4.7 PE/cm. Figure from [77]. b) Efficiency for registering a hit on a muon track as a function of distance from the APD. Efficiency is greater than 90% for the length of cells, satisfying the design requirement. Figure from [78].	51
3.14	View of the upstream, 3×3 module region of the Near Detector during construction.	52
3.15	A) 3×3 module active region of the Near Detector. B) Muon catcher region with steel plates alternating with scintillation planes, height is two-thirds that of the active region. C) Electronics rack alcove. D) Catwalks. E) Movable access platform [80].	53
3.16	A typical attenuation curve for a cell in the Far Detector where the readout end is at +750 cm and the far end of the cell is at -750 cm. Due to the looped fiber in each cell, the data is fit with a double exponential shown in red for the short and long light paths to the APD [81].	55
3.17	Energy loss with respect to distance from track end of stopping muons identified using their Michel electrons. This plot was made using cosmic data from the Far Detector. Super-imposed is the mean energy loss per cm as a function of distance from the track end. The flat region between 100 and 200 centimeters from the end of the track is used to fix the absolute energy scale [82].	56

4.1	Schematic of the distributive timing system deployed at the Far Detector. Figure from [88].	60
4.2	Schematic of the distributive timing system deployed at the Near Detector.	61
4.3	Sample calculation of the MTDU, STDU and DCM delay values that would be loaded into the firmware registers of a toy system. Figure adapted from [90].	63
4.4	a) Delay measured from a DCM to the end of the timing branch in the Far Detector. DCMs 1-6 are in a branch on top of the detector with 6 being closest to the STDU backbone. DCMs 7-12 are located in a branch on the side, with 7 being closest to the backbone. b) Delay calibrated from a timing unit to the end of the backbone in the Far Detector. The MTDU is located off of the detector hall and has a longer delay, each STDU is uniformly spaced 93.75 ns apart.	64
4.5	Schematic of the spill server system designed to relay accelerator event time stamps from the MTDU's at Fermilab to applications running on the Near and Far detectors triggering data readout. Figure from [89].	66
4.6	An illustration of the DCS algorithm used to trigger readout on the FEB. The black and blue curves represent two different pulses and the markers s_{i-3} , s_{i-2} , s_{i-1} , and s_i are the digitization samples. In single-point mode these two pulses would have the same change and time recorded. In multi- point mode saving all four samples allows a more accurate fit for the pulse start time and peak ADC [93].	68

4.7	The time difference between pairs of cell hits in the same DCM on reconstructed muon tracks as a function of the number of photo-electrons in each hit shown for Far Detector multi-point data. The RMS of each column is the timing resolution.	69
4.8	Timing resolution determined from Far Detector data.	70
4.9	Timing resolution determined from Near Detector data.	70
4.10	A representative example of the measured relative time offset between a pair of DCMs and the expected value based on the detector operating conditions at the time of writing this document.	74
4.11	The matrix of relative timing offsets between DCMs at the Far Detector. Dotted lines indicate the diblock boundaries and blank spaces indicate DCM pairs with insufficient statistics. The structure seen in the plot is due to the repeated pattern of cable delays between DCMs in a diblock, discussed in Figure 4.1.	75
4.12	The matrix of relative timing offsets for the Near Detector. Dotted lines indicate the diblock boundaries. The structure in the plot is different from that of the Far Detector shown in Figure 4.11 because the DCMs are connected in a different pattern, most notably the detector views are not on separate timing branches (Figure 4.2).	75
4.13	The absolute timing calibration for diblocks 2 to 14 in the Far Detector, showing a systematic drift beginning in diblock 7. The red line indicates the measured delays from the TDU and the black points are the calibration result using muons. Dotted lines indicate the diblock boundaries.	78

4.14	The absolute timing calibration for diblocks 2 to 14 in the Far Detector after a drift correction was applied to diblocks 7 to 14. The red line indicates the measured delays from the TDU and the black points are the calibration result using muons. Dotted lines indicate the diblock boundaries.	78
4.15	The errors in the absolute timing calibration are shown before (a) and after (b) a drift correction was applied to diblocks 7-14.	79
4.16	The timing calibration was performed on a set of runs two weeks apart to show stability in the result on a run-by-run basis. Results from two representative DCMs are shown.	81
4.17	a) The absolute calibration of the Near Detector after correcting for diblock drift is shown by the black points with the expected result in red. b) The errors on the calibration result. The blue dotted lines indicate diblock boundaries.	82
4.18	The correction to the reconstructed times as a function of the distance of the energy deposition to the APD for the far (a) and near (b) detectors. The three curves illustrate the effect variations in fiber speed have on the correction. 83	
5.1	Simulated 2.15 GeV neutrino interactions showing a toy ν_μ CC interaction (top panel), ν_e CC interaction (middle panel), and NC interaction (bottom panel). Figure from [95].	86
5.2	Flowchart of the reconstruction algorithms used to identify ν_e CC interactions.	88

5.3	550 μ s Far Detector readout window before (a) and after (b) slicing reconstruction. Hits are colored by time. The top panel shows the vertically oriented planes and the bottom panel the horizontal planes. Hits are colored in bold with the reconstructed slice they are a member of, un-bolded hits are not associated with a physics slice. At ~ 3600 cm in z a contained neutrino interaction is clearly seen.	91
5.4	The completeness and purity of reconstructed slices with at least 10 hits per view for Far Detector cosmic ray simulations (a,b) and Near Detector neutrino simulations (c,d). Figure from [97].	92
5.5	The first and second iteration of the multi-Hough transform on one view of a Far Detector neutrino simulation event. The first selected line is shown in (a) and the corresponding Hough map in (b). After hits associated with this line are removed the second line (c) and corresponding map (d) are calculated. Figure from [99].	93
5.6	The perpendicular distance between the first (a) and second (b) most dominant Hough lines and the true neutrino interaction vertex for simulated ν_e CC, ν_μ CC, and NC interactions in the X-sampling readout view area normalized to one. Figure from [99].	94
5.7	Vertex resolution for simulated ν_e CC events in the Far Detector.	100
5.8	Vertex resolution for simulated ν_μ CC events in the Far Detector.	101
5.9	Vertex resolution for simulated NC events in the Far Detector.	102
5.10	Vertex resolution for simulated ν_μ CC events in the Near Detector.	104
5.11	Comparison of the reconstructed vertex distribution in data and Monte Carlo for the Near Detector. Red bands show the beam flux uncertainty. Figure from [106].	105

5.12	In the left panel are simulated 1 GeV muons in the Far Detector and in the right 2 GeV. The x-axis is the distance in centimeters of the hit to the vertex and the y-axis is the difference between the angle of the hit with respect to the vertex and the initial true momentum of the particle. The function drawn on the plot in red is equation 5.10.	107
5.13	Points 1 and 2 are both equidistant from clusters <i>A</i> and <i>B</i> . In a fuzzy <i>k</i> -means model this means that both points would have a membership of 50% to each cluster even though point 2 is clearly more of an outlier and is better suited to be flagged as noise or form a new cluster.	109
5.14	A simulated ν_e CC quasielastic interaction in the Far Detector with completed 3D prong reconstruction from the fuzzy-k means algorithm. The prong outlined in red in each view is the reconstructed electron and outlined in green is the reconstructed proton. The corresponding energy profile histograms used to compute the suitable 3D prong matches is shown in Figure 5.15.	116
5.15	Cumulative energy profile as a function of path length along a prong for perspective 3D match candidates shown in Figure 5.14. The red curves are for prongs in the XZ (vertical planes) view and the blue is for the YZ (horizontal planes) view. The upper-left and lower-right panels show the preferred matches with similar energy profiles that result in the green and red track respectively. The off-diagonal elements illustrate the difference in energy profile shape for the wrong combinations.	116

5.16	Completeness (a), purity (b), and angular resolution (c) of the reconstructed electron in simulated Far Detector ν_e CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the lepton's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 625$ cm and $100 < z < 5500$ cm.	118
5.17	Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the lepton's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 625$ cm and $100 < z < 5500$ cm.	119
5.18	Completeness (a), purity (b), and angular resolution (c) of the reconstructed electron in simulated Far Detector ν_e CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 625$ cm and $100 < z < 5500$ cm.	120

5.19	Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 625$ cm and $100 < z < 5500$ cm.	121
5.20	Completeness (a), purity (b), and angular resolution (c) of the most energetic reconstructed prong in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the prong's reconstructed energy. A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 625$ cm and $100 < z < 5500$ cm.	122
5.21	Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Near Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $ x, y < 140$ cm and $100 < z < 700$ cm.	123
5.22	Comparison of the number of 3D prongs (a), number of hits in the most energetic prong (b) and length of the longest prong in a slice (c) in data and Monte Carlo for the Near Detector.	124
5.23	The longitudinal dE/dx profile comparing simulated electrons (red) to other particles (blue), specifically photons (a) and muons (b). The left panels show the profile for plane 2 of the shower and the right panels show plane ten. Figure from [113].	129

5.24	The transverse dE/dx profile for simulated electrons (red) and π^0 's (blue). The left panel shows the profile for transverse cell index 0 and the right shows transverse cell index 3. Figure from [113].	129
5.25	The likelihood differences between electrons and other particles: photons (a), muons (b), and protons (c). True electrons are plotted in red, the other particle in blue. The left panels show the longitudinal likelihood difference and the right panels show the transverse. Figure from [113].	131
5.26	The likelihood differences between electrons and other particles: neutrons (a), pions (b), and π^0 's (c). True electrons are plotted in red, the other particle in blue. The left panels show the longitudinal likelihood difference and the right panels show the transverse. Figure from [113].	132
5.27	The structure of the artificial neural network passing from 16 input variables through three hidden layers of neurons to the output. Variable names and importance shown in Table 5.2. Figure from [113].	135
5.28	Event classification output from the LID algorithm where values closer to one represent ν_e CC signal. The plot shows simulated Far Detector events scaled to the analysis exposure of 2.8×10^{20} full detector equivalent POT with standard preselection cuts (a) and a version zoomed into the signal region (b).	136
6.1	Number of consecutive analyzable diblocks in the Far Detector during the analysis period from February 6 th , 2014 until May 15 th , 2015.	139
6.2	Timing diagram of the beam window and out-of-time cosmic sideband. . . .	140
6.3	Distributions of MIP rate for APD pixels in Far Detector data as a function of DCM location. Figure from [117].	142

6.4	Mean slice times for Near Detector Data fit for the rise and fall of the neutrino beam peak. The tail on falling edge is due to delayed neutron shine in the detector from neutrino interactions in the surrounding rock. Figure from [119].	144
6.5	LID spectrum of un-extrapolated Far Detector Monte Carlo scaled to 3.52×10^{20} POT. Also shown is cosmic background data from minimum bias trigger data, scaled to 165 live seconds. The selection at an LID value greater than 0.95 is drawn, which maximizes the FOM. Figure from [123].	147
6.6	Drawing of the fiducial (red) and shower containment boundaries (green) inside the Near Detector (black). Figure from [125].	149
6.7	Slice calorimetric energy versus the number of hits in the slice for both detectors.	150
6.8	Kinematic quantities in lepton-nucleon scattering where k and k' are the four-momenta of the incoming and outgoing leptons, P is the four-momentum of the nucleon with mass M , W is the mass of the hadronic final state, and $q = k - k'$ is the four-momentum transfer to the nucleon. Figure from [28].	151
6.9	In color are the kinematic variables W^2 (a) and Q^2 (b) plotted against the true neutrino energy for events passing selection in Near Detector Monte Carlo. Overlaid with boxes are Far Detector Monte Carlo events with standard oscillation weights passing selection. Both selections are area normalized to one.	151
6.10	LID distribution of Near Detector data and Monte Carlo passing the selection criteria on a log scale (a) and zoomed in near the signal region (b). Plots are normalized to 1.66×10^{20} POT in the data.	152
6.11	Calorimetric energy of slices in Near Detector data and Monte Carlo passing selection and $LID > 0.95$.	154

6.12	The efficiency of the Far Detector ν_e CC selection as a function of detector size with binomial errors shown. The fit is of the form $y = a + \frac{b}{x^{0.993}}$ with $a = 0.509 \pm 0.005$ and $b = -0.557 \pm 0.048$	155
6.13	The cumulative POT exposure and fiducial mass for the analysis period. The exposure is $27.9 \text{ POT} \times kT \times 10^{20}$ which is the equivalent of 2.8×10^{20} POT for a full 10.0 kT fiducial volume detector.	158
6.14	POT accumulated in each detector configuration for pre and post shutdown data (black) and real-conditions Monte Carlo (red). The Monte Carlo is area normalized to the data. Beam and data quality cuts are applied at the spill level to both data and Monte Carlo.	158
7.1	Steps to extrapolate prediction for the $\nu_\mu \rightarrow \nu_e$ signal channel. Starting with the data/MC ratio in the Near Detector (a), using a matrix (b) to convert to a true energy spectrum in the Near Detector (c), re-weighting the Far Detector Monte Carlo in true energy (d), and then applying a matrix (e) to convert the prediction back to units of reconstructed energy (f).	164
7.2	Extrapolated prediction for the major background channels: $\nu_e \rightarrow \nu_e$ (a,b), $\nu_\mu \rightarrow \nu_\mu$ (c,d), and neutral-current (e,f). Left panels show the data and Monte Carlo spectra in the Near Detector with the ν_e selection, right panels show the Far Detector Monte Carlo and extrapolated prediction.	165
7.3	Data and Monte Carlo comparisons of the LID distribution for cosmic muon induced electromagnetic showers in the Far Detector (a) and Near Detector (b). In the Far Detector the distribution is also compared to nu_e CC signal events. Figure from [133].	166

7.4	The cumulative Far Detector exposure overlaid with the accumulation of 27 candidate events in the 3 to 10 GeV sideband passing selection cuts and visually identified as consistent with a neutrino. Applying a Kolmogorov-Smirnov test to the distribution gives a p value of 0.78	168
7.5	Far Detector LID spectrum (a) and calorimetric energy distribution (b) of high energy sideband events selected in data overlaid on the Monte Carlo neutrino prediction and out-of-time cosmic background. Applying a Kolmogorov-Smirnov test to the LID distribution gives a p value of 0.84.	168
7.6	Event displays for the two high energy neutrino events selected with an LID value greater than 0.7. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time.	169
7.7	LID spectrum (a) and calorimetric energy distribution with LID > 0.95 (b) of Near Detector data and Monte Carlo with the beam flux uncertainty band drawn.	172
7.8	Systematic uncertainties in percentages on background (a) and signal (b) in the Far Detector.	179
8.1	a) Accumulation of ν_e candidate events as a function of exposure. b) Calorimetric energy distribution of selected events compared to the ν_e CC signal and beam neutrino background projection with the standard set of oscillation weights assuming no matter effect or CP violation.	184
8.2	Event displays for the LID selected events in the NUMI beam window for 3.52×10^{20} POT of Far Detector data. The reconstructed vertex and 3D reconstructed prongs are drawn. Displays are zoomed on the interaction of interest but preserve the aspect ratio of the full detector. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time with the slice.	185

8.3	More event displays for the LID selected events in the NUMI beam window for 3.52×10^{20} POT of Far Detector data. The reconstructed vertex and 3D reconstructed prongs are drawn. Displays are zoomed on the interaction of interest but preserve the aspect ratio of the full detector. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time with the slice.	186
8.4	a) δ vs $\sin^2 \theta_{13}$ contour of the NOvA result for the normal (red) and inverted (blue) mass ordering. The fit was performed without reactor experiment constraints. b) The 90% percent contours of NOvA are compared to the T2K contours from an analysis of 28 events in 6.6×10^{20} POT.	188
8.5	Best fit of $\sin^2 \theta_{13}$ for the normal (red) and inverted (blue) mass orderings after marginalizing over δ from Figure 8.4. The fit does not include a constraint on θ_{13} from the reactor experiments.	189
8.6	a) Best fit of δ to the NOvA data, constraining θ_{13} from the Daya Bay measurements and the atmospheric parameters from T2K. The normal mass ordering (red) is slightly preferred over the inverted ordering (blue). b) The T2K contours under the same parameter constraints with 28 candidate ν_e CC events from an analysis of 6.6×10^{20} POT [64] is shown on the same plot for comparison.	190
8.7	Best fit of δ to the NOvA data assuming 10.56×10^{20} POT and 18 events, which is three times the dataset used in this analysis, expected by the summer of 2016.	191

LIST OF TABLES

2.1	Best-fit values and 3σ allowed ranges for three-flavor oscillation parameters derived from a global fit [66]. The values (values in parentheses) correspond to normal (inverted) mass ordering choice and $\Delta m^2 = m_3^2 - (m_2^2 - m_1^2) / 2$.	35
4.1	List of acronyms frequently used in this chapter.	58
4.2	Far Detector run periods used for the timing calibration.	77
5.1	Vertex resolution for simulated Far Detector neutrino interactions, all numbers in centimeters. The mean and full-width half-maximum (resolution) of the x , y , and z vertex coordinate and the combined 3D difference between the true and reconstructed vertex. Numbers are derived from Figures 5.7, 5.8, and 5.9.	99
5.2	Near Detector data (1.6×10^{20} POT) and Monte Carlo (normalized to data) event counts passing the levels of preselection cuts. Table from [113].	135
6.1	Far Detector event counts and Figure of Merit ($\text{signal}/\sqrt{\text{background}}$) for out-of-time NuMI cosmic background scaled to 165 liveseconds and NuMI Monte Carlo scaled to 2.8×10^{20} POT. Table from [123].	147
6.2	Near Detector data (1.65×10^{20} POT) and Monte Carlo (normalized to data) event counts passing the levels of preselection.	153
6.3	Effective fiducial mass in kilotons as a function of detector size derived from selection efficiencies in the Monte Carlo.	156

6.4	POT, fiducial mass and livetime accounting for First Analysis period. POT exposure and livetime are also expressed in terms of a full 10.0 fiducial kiloton equivalent.	157
7.1	Extrapolated prediction of Far Detector event counts normalized to 3.52×10^{20} POT and the standard Far Detector Monte Carlo for comparison. . .	163
7.2	Prediction of Far Detector neutrino background and out-of time cosmic data for the 3 to 10 GeV sideband for different LID cuts. Panel A shows the prediction based on Far Detector Monte Carlo, while panel B shows the prediction based on the extrapolation of Near Detector data. In both cases the prediction is normalized to 3.52×10^{20} POT. Panel C shows the measured result in data, where the 57 candidate events were visually hand-scanned and classified as appearing like a neutrino or cosmic ray.	167
7.3	Prediction of Far Detector neutrino background and out-of time cosmic data in the low-LID sideband. Panel A shows the prediction based on the extrapolation of Near Detector data normalized to 3.52×10^{20} POT. Panel B shows the measured result in data, where the 5 candidate events were visually hand-scanned and classified as appearing like a neutrino or cosmic ray. . .	170
7.4	Calibration systematic error relative change from nominal for the signal and background components of the LID selection. Table from [134].	174
7.5	Using the standard Near Detector decomposition technique, the extrapolation is performed with eight different geographic regions of the Near Detector data and Monte Carlo to measure containment effects. Values recorded are the relative percent difference with respect to the nominal full-detector extrapolation. The final result is an absolute average of the uncertainties. . .	175

7.6	Extrapolated prediction of Far Detector events, recorded as the relative percentage change from nominal. Different decomposition methods for the Near Detector data are shown. Any excess/deficit between data and Monte Carlo in an energy bin is either distributed proportionally among the components, or assigned entirely to one component.	176
7.7	Statistical uncertainty on extrapolation of background components in percents. Table from [134].	177
7.8	Final systematic uncertainty (in percentages) on the combined background and signal in the Far Detector for LID. The last row corresponds to the sum in quadrature of all errors in this table. Table from [134].	179
8.1	Summary of oscillation parameters and uncertainties used in the analysis. Table adapted from [152].	183

CHAPTER 1

Introduction

Neutrinos have proved to be an elusive and puzzling particle throughout their history. The existence of neutrinos was first postulated by Wolfgang Pauli in 1930 in a letter to a conference he was unable to attend [1] as a “desperate remedy” to the puzzle of nuclear β decay. At the time β decay was believed to be a two-body process and so the emitted electron in the radioactive decay was expected to have a discrete energy. However, experimental observations, first made in 1914 [2], had shown the spectrum to be continuous. Pauli’s solution was to suggest that a third neutral particle, a “neutron”, was involved in the process and which carried away undetected energy. A formal theory of β decay was proposed by Enrico Fermi in 1934 [3]. In the theory Pauli’s “neutron” was re-coined “the neutrino”, meaning little neutral one, in light of the discovery of the more massive “neutron” by James Chadwick in 1932 [4].

In 1946 Bruno Pontecorvo proposed a method to directly detect the neutrino through the inverse beta decay process, $p + \bar{\nu} \rightarrow n + e^+$ [5]. Although initially considering a nuclear bomb as a source of neutrinos for the experiment, Fred Reines and Clyde Cowan eventually settled on using a nuclear reactor for the neutrino flux. The experiment, named Project Poltergeist, searched for the coincident scintillation light of the positron with the delayed gamma ray from neutron capture first at the Hanford reactor in 1953 [6] and definitively in 1956 at Savannah River [7]. Fred Reines was awarded the Nobel Prize for the discovery in

1995.

A second flavor of neutrino, the muon neutrino, was discovered in 1962 at Brookhaven [8], which won the Nobel Prize in 1998 for Leon Lederman, Melvin Schwartz, and Jack Steinberger. The third generation of neutrino, the tau, was finally observed in 2000 by the DONUT collaboration at Fermilab [9]. Measurements of the width of the decay of the Z boson have shown that there are no more than three generations of light active neutrino flavors [10]. These discoveries complete the Standard Model where there is a neutrino associated with each lepton generation: electron, muon, and tau.

The first experimental hints of the phenomenon that would become known as neutrino oscillations came in 1968. Ray Davis set up an experiment in the Homestake mine in South Dakota to measure the electron neutrino flux from the Sun as direct confirmation of the theoretical models of nuclear fusion and fission in the Sun. The experiment was located deep underground to reduce cosmic-ray backgrounds and counted the number of ^{37}Ar atoms produced in a 390,000 litre container of tetrachloroethylene through the process $\nu_e + ^{37}\text{Cl} \rightarrow e^- + ^{37}\text{Ar}$. The result was a rate of ν_e interactions of about one third of the theoretical predictions. While the neutrino rate discrepancy was initially attributed to errors in either the measurement or the theoretical prediction, this “solar neutrino problem” persisted when a similar measurement was made in 1989 by the Kamiokande-II experiment in a water-Cherenkov detector [11]. The SAGE [12] and GALLEX [13] experiments in 1991 and 1992 saw the same deficit by measuring $\nu_e + ^{71}\text{Ga} \rightarrow e^- + ^{71}\text{Ge}$ in a Gallium detector. The working hypothesis that developed was that neutrinos created in one flavor state oscillated to different flavor states depleting the counts of the original flavor. In the case of the solar neutrino problem the electron neutrinos would be turning into muon and tau neutrinos.

The first definitive measurement of neutrino oscillations came from atmospheric neu-

trinos in the Super-Kamiokande experiment in 1998 [14]. The SNO experiment provided conclusive evidence in 2001 that the solar neutrino problem was also a consequence of neutrino oscillations [15]. Ray Davis and Masatoshi Koshihba were awarded the 2002 Nobel Prize for their pioneering contributions to the detection of cosmic neutrinos.

One consequence of neutrino oscillations is the implication that neutrinos have non-zero mass, which is not predicted in the standard model and requires new physics. One of the theoretical mechanisms for generating neutrino mass leads to an explanation of the matter-antimatter asymmetry of the universe (Section 2.1). Dozens of experiments have developed in the last 20 years to map out the parameters involved in neutrino oscillations (see Chapter 2). Two of the remaining unknowns include the ordering of the three neutrino mass states and the existence of CP violation in the neutrino sector. The NOvA experiment is designed to measure both of these parameters.

NOvA generates a beam peaked at 2 GeV of muon neutrinos with an accelerator at Fermilab. The beam passes through two functionally identical liquid scintillator tracking detectors, one located at Fermilab, the other 810 kilometers away in Ash River, MN. The detectors are relatively fine-grained, consisting of alternating horizontal and vertical planes of plastic cells with a 4×6 cm cross section filled with scintillator. The baseline is the longest used for such an experiment in the world and increases the sensitivity to the ordering of the neutrino masses. The measurement described in this thesis, and one of the primary goals of the experiment, is the observation of the appearance of electron neutrinos in the Far Detector. The design of the experiment, optimized for this oscillation channel, is described in Chapter 3.

The neutrino beam is pulsed, delivering a $9.6 \mu\text{s}$ batch of neutrinos every 1.33 seconds. In order to capture the neutrino beam without being overwhelmed by cosmic-ray backgrounds a fine timing system is required to trigger the readout of data in both detectors precisely

when the beam is present. The large size of the Far Detector ($15.6 \times 15.6 \times 60$ meters) requires a precisely synchronized internal timing system in order to reconstruct particles traversing the detector. A significant personal contribution was made to the experiment developing monitoring and calibration tools to synchronize the detectors in time both internally and with respect to the neutrino beam, elaborated in Chapter 4.

Performing an analysis on the data requires reconstruction and classification of interactions in the the detector as signal (electron neutrinos) or background (cosmic rays, muon neutrinos, and neutral current interactions). A suite of software is used to spatially and temporally separate interactions into reconstructable objects. A major contribution of this thesis to the analysis was the development of a tracking algorithm robust enough to handle electron showers, muon tracks, displaced photons, and small energy deposits from neutrons (Chapter 5). The 3D reconstructed objects from this algorithm are the primary input to an interaction classification algorithm that determines the degree to which the interaction topology is consistent with an electron neutrino.

The Near Detector, which sees a high neutrino interaction rate due to proximity to the beam, is used to measure background event rates in the analysis and extrapolate a background projection to the Far Detector which reduces systematic uncertainties on the measurement. Selection criteria are developed in the Far Detector to identify electron neutrinos in the energy window where oscillations occur while removing backgrounds. In the Near Detector selection is designed to measure backgrounds with topologies present in the Far Detector (Chapter 6). Systematic uncertainties from a variety of effects are measured though a combination of data-driven techniques and simulation studies (Chapter 7). This analysis was performed in a blind fashion, with all event selection criteria and systematic uncertainties set before the Far Detector neutrino beam data was processed. To check the quality of the analysis before un-blinding the result several sideband regions are

examined. The final results, presented in Chapter 8, show a measurement of six events which is a 3.3σ observation electron neutrino appearance at the longest baseline in the world and discuss preliminary measurements of the neutrino mass ordering. The impact of this result on the global neutrino community and the prospects for the next six years of data taking are discussed in Chapter 9.

CHAPTER 2

Neutrino Oscillation Theory and Experimental Landscape

The theoretical picture developed over the past 50 years to explain the solar neutrino problem and experimental observations of atmospheric, reactor and accelerator neutrinos is that, contrary to the standard model, neutrinos have non-zero mass and can therefore change electro-weak flavor state or “oscillate” during flight. The idea that non-zero mass could lead to oscillations was first suggested by Pontecorvo in 1958, although the focus was on neutrino to anti-neutrino oscillations occurring analogous to the kaon system [16]. In 1962 the modern picture of oscillations was put forward by Maki, Nakagawa, and Sakata, after the discovery of the second neutrino flavor, that oscillations occur between flavor states as a result of mixing of independent mass states [17].

In charged-current weak interaction production, a neutrino is created in a definite electro-weak eigenstate with the associated lepton (e, μ, τ) , shown in Figure 2.1, and is a superposition of the three mass states. The relative mixing strengths of each mass state to the flavor states is determined by the PMNS matrix named for Pontecorvo, Maki, Nakagawa, and Sakata. This matrix, U , is unitary such that $UU^\dagger = I$. The electro-weak neutrino flavor α can thus be written as

$$|\nu_\alpha\rangle = \sum_{k=1}^3 U_{\alpha k}^* |\nu_k\rangle, \quad (2.1)$$

where $U_{\alpha k}^*$ is the matrix element with the relative amplitude coupling flavor state α to mass state k . After this superposition of neutrino mass states is created via an electro-weak




Figure 2.1: Feynman diagram for $\pi^+ \rightarrow \mu^+ + \nu_\mu$ illustrating the neutrino produced in a definite electro-weak state matching the lepton.

interaction a phase difference accumulates as the states propagate in energy eigenstates. This phase difference means that when the neutrino interacts again the mass states are in a different superposition and may associate to a different flavor when detected.

2.1 Neutrino Mass

One of the most important consequences of neutrino oscillations is the implication that neutrinos have non-zero mass, which is not predicted by the Standard Model. While the relative mass splittings have been measured through oscillation experiments, the absolute scale is inaccessible. The current best upper limit on the sum of the neutrino masses comes from measurements of the Cosmic Microwave Background by the Plank collaboration combined with other astrophysical data [18]: $\sum m_\nu < 0.23$ eV, 95% C.L. There are ongoing efforts to measure the neutrino mass both through observations of neutrinoless double-beta decay and precision measurements of the endpoint of the beta decay spectrum ($m_\beta = \sqrt{\sum_i |U_{ei}|^2 m_i^2} < 2.0$ eV), summarized in [19].

In the Standard Model neutrinos do not receive mass because there is not a right-handed neutrino field. If a neutrino is treated as a Dirac particle, meaning the neutrino and anti-neutrino are distinct states as with quarks and leptons, then a simple extension to the SM

can be made to add right handed neutrino field allowing a mass term to be written of the form

$$\mathcal{L}_D = -m_D (\bar{\nu}_R \nu_L + \bar{\nu}_L \nu_R). \quad (2.2)$$

In this equation ν_L and ν_R are the three-component column vectors the left and right handed neutrino fields and $m_D = \frac{yv}{\sqrt{2}}$ is the Dirac mass term with y the Yukawa coupling and v the vacuum expectation value of the Higgs field. ν_R is a weak isospin singlet while ν_L is a doublet with the corresponding charged lepton, thus this mass term does not respect the symmetry of the Standard Model. While an equation of this form is possible, it offers no explanation for why the Higgs-neutrino Yukawa coupling would have to be six orders of magnitude smaller than the Higgs-quark coupling to generate the small mass values. In search of more natural explanations for the neutrino mass scale one popular idea is the seesaw mechanism [20, 21].

In the most basic form of the seesaw mechanism a Majorana mass term is added to the Lagrangian. A Majorana particle is defined as a particle that is identical to its anti-particle. This condition is only possible for neutral particles such as neutrinos. Mass terms can be constructed for both right and left-handed fields of the form

$$\mathcal{L}_{M_L} = -\frac{1}{2} m_M^L (\bar{\nu}_L \nu_L^c + \bar{\nu}_L^c \nu_L) \quad (2.3)$$

$$\mathcal{L}_{M_R} = -\frac{1}{2} m_M^R (\bar{\nu}_R \nu_R^c + \bar{\nu}_R^c \nu_R) \quad (2.4)$$

where ν is the right or left handed neutrino field and ν^c is the charge-conjugated field. The left-handed term in Equation 2.3 violates gauge symmetries and is not allowed in the Standard Model, while there are no restrictions on the right handed term. The right-handed Majorana mass term and the Dirac mass term can be combined into one Lagrangian of the

form

$$\mathcal{L} = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\nu}_R^c \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & m_M^R \end{pmatrix} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + h.c. \quad (2.5)$$

The mass matrix can be diagonalized as

$$m_{\pm} = \frac{1}{2} m_M^R \pm \frac{1}{2} \sqrt{(m_M^R)^2 + 4m_D^2}. \quad (2.6)$$

Note that a phase transformation is applied such that for $m_- < 0$ the sign is flipped to positive. The Dirac mass m_D can be generated through the Standard Model Higgs mechanism and is expected to be on the order of the charged lepton mass. The Majorana mass term m_R is not confined to the same scale and plausibly exists at the GUT scale ($10^{14} - 10^{16}$ GeV). In the limit that $m_M^R \gg m_D$ the eigenvectors approximate as

$$\nu_+ \approx (\nu_R + \nu_R^c) + \frac{m_D}{m_M^R} (\nu_L + \nu_L^c) \quad (2.7)$$

$$\nu_- \approx (\nu_L - \nu_L^c) + \frac{m_D}{m_M^R} (\nu_R + \nu_R^c). \quad (2.8)$$

In this solution there is one heavy state $N = \nu_+$ dominated by the sterile right handed term with a mass $m_N \approx m_M^R$ and a light state with $\nu = \nu_-$ and a mass $m_\nu \approx \frac{m_D^2}{m_M^R}$. Now we see that the seesaw mechanism naturally produces the light active neutrino flavors observed today with $m_\nu \approx 0.1$ eV, when $m_D \approx 100$ GeV in line with the quark scale and the massive right handed Majorana neutrinos on order $m_M^R \approx 10^{14}$ GeV. These neutrinos would have masses on order with the GUT scale and could have been found in the early universe due to the high temperatures. The heavy neutrinos would have decayed via the Yukawa coupling, and if there was a CP violating phase in the coupling, would have produced a lepton/anti-lepton asymmetry. Then symmetry breaking decay processes in the hot early universe that violate net lepton and baryon number could have propagated this asymmetry into a

baryon/anti-baryon asymmetry. This scenario for the generation of the currently observed baryon asymmetry from CP violation in the neutrino sector is known as leptogenesis. Recent studies have shown that the CP violating phase δ in the neutrino mixing matrix can provide the necessary CP violation in leptogenesis to explain the total baryon asymmetry in the Universe [22–24].

The seesaw mechanism not only provides a natural explanation for the light neutrino masses, it also provides a theory to explain the matter-antimatter asymmetry. In order for the seesaw mechanism to be plausible neutrinos must be found to be Majorana particles and CP violation in the lepton sector is required [25]. Specifically, CP violation is required in the decay of the heavy Majorana neutrinos, but generally this implies CP violation in light neutrinos observed through neutrino oscillations which can be probed by long-baseline accelerator experiments such as NOvA.

2.2 Oscillation Formalism

In this section the neutrino oscillation formalism is derived following [26] and [27] using a plane-wave approximation. This approximation assumes that all massive neutrino components have the same momentum and that the propagation time t is equal to the distance traveled L . Plane waves occupy all space and time so the second assumption is put in by hand from the more formal wave-packet QFT derivation also found in the reference. Natural units of $\hbar = c = 1$ are used throughout.

The neutrino mass states evolve in time according to the Schrödinger equation

$$i \frac{d}{dt} |\nu_k(t)\rangle = \mathcal{H} |\nu_k(t)\rangle \quad (2.9)$$

with energy eigenvalues

$$E_k = \sqrt{p^2 + m_k^2}, \quad (2.10)$$

which allows for the mass state to evolve as a plane wave

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle. \quad (2.11)$$

Applying Equation 2.11 to 2.1 gives

$$|\nu_\alpha\rangle = \sum_{k=1}^3 U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle. \quad (2.12)$$

Using the unitarity of the PMNS matrix and the inverse of Equation 2.1, we can rewrite the result as

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_{k=1}^3 U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle \quad (2.13)$$

so that the pure flavor state ν_α at $t = 0$ becomes a superposition of flavor states at a later time. The amplitude of the transition $\nu_\alpha \rightarrow \nu_\beta$ as a function of time can be expressed as

$$A_{\nu_\alpha \rightarrow \nu_\beta}(t) \equiv \langle \nu_\beta | \nu_\alpha(t) \rangle = \sum_k U_{\alpha k}^* U_{\beta k} e^{-iE_k t}, \quad (2.14)$$

resulting in a transition probability of

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |A_{\nu_\alpha \rightarrow \nu_\beta}(t)|^2 = \sum_{kj} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (2.15)$$

For ultra-relativistic neutrinos the energy eigenvalues in Equation 2.10 can be approximated as

$$E_k \approx E + \frac{m_k^2}{2E} \quad (2.16)$$

where $E = p$ is the neutrino energy neglecting the mass. Additionally, time is approximated

as the distance traveled L so that the final oscillation probability is

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = \sum_{kj} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{kj}^2 L}{2E}} \quad (2.17)$$

where the neutrino oscillation phase is determined by the ratio of the distance traveled L to the neutrino energy E and $\Delta m_{kj}^2 = m_k^2 - m_j^2$. The exponential in Equation 2.17 can be replaced using Euler's formula and the property $\cos 2\theta = 1 - 2\sin^2 \theta$. Additionally, from the properties of unitary matrices $\sum_k U_{\alpha k} U_{\beta k}^* = \delta_{\alpha\beta}$, so the probability of a neutrino created in electroweak state α being observed in state β can be rewritten as:

$$\begin{aligned} P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) &= \delta_{\alpha\beta} - 4 \sum_{k>j} \Re [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{kj}^2 L}{4E} \right) \\ &\quad + 2 \sum_{k>j} \Im [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin \left(\frac{\Delta m_{kj}^2 L}{2E} \right). \end{aligned} \quad (2.18)$$

If instead of calculating the transition probability we are interested in the survival probability, meaning how likely is a neutrino created in state α also detected in state α , then Equation 2.18 can be simplified since the quartic products are all real:

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - 4 \sum_{k>j} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{\Delta m_{kj}^2 L}{4E} \right). \quad (2.19)$$

Neutrinos and anti-neutrinos can be related through a CP transformation $\nu_\alpha \xleftrightarrow{CP} \bar{\nu}_\alpha$, interchanging the neutrino and anti-neutrino and reversing the helicity. A T transformation exchanges the initial and final states. Thus, survival probabilities are the same for neutrinos and anti-neutrinos due to CPT invariance, $\nu_\alpha \rightarrow \nu_\alpha \xleftrightarrow{CP} \bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha$, and cannot be used to observe δ (Section 2.3). The implication is that if there is CP violation in the neutrino sector it can only be measured through a transitional probability where an experiment would look for $\nu_\alpha \rightarrow \nu_\beta \neq \bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta$.

2.3 PMNS Matrix

In the standard picture of neutrino oscillations with three active neutrino flavors and no sterile states the 3×3 PMNS matrix is written in the form

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (2.20)$$

where each element controls the mixing of an electro-weak flavor state to a mass state. The matrix can be parameterized with three rotation angles θ_{12} , θ_{23} , and θ_{13} and a CP phase δ . The expanded form of the matrix, following the Particle Data Group convention [28] is

$$\begin{aligned} U &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \end{aligned} \quad (2.21)$$

where $s_{ij} = \sin \theta_{ij}$ and $c_{ij} = \cos \theta_{ij}$. The angle $\theta_\odot = \theta_{12}$ is known as the solar mixing angle for historic reasons. This angle is predominantly measured with solar ($L \sim 10^8$ km) and reactor neutrinos ($E \sim 1$ MeV) where the L/E ratio is large. The associated mass splitting is small with $\Delta m_{\odot}^2 = \Delta m_{21}^2 \approx 8 \times 10^{-5}$ eV². Similarly, θ_{23} is known as the atmospheric mixing angle and is probed with atmospheric and accelerator neutrinos where the baseline is relatively short ($L \leq 10^4$ km) and the energies are high ($E \sim 1$ GeV) such that L/E is comparatively small. The atmospheric sector is associated with a large mass splitting $\Delta m_{32}^2 \approx \Delta m_{31}^2 \approx 3 \times 10^{-3}$ eV². For δ to possibly be non-zero all three mixing angles must be non-zero. It should be noted that the PMNS matrix also has two Majorana phases α_1

and α_2 that are factored into a fourth sub-matrix, but these phases have no impact on the oscillation results and have been thus suppressed in Equation 2.21. While it is known (see Section 2.7.1) that $m_2 > m_1$, the sign of the large atmospheric mass splitting has not been measured. It remains to be seen if the mass ordering is normal ($m_3 > m_2 > m_1$) or inverted ($m_2 > m_1 > m_3$) as shown in Figure 2.2.

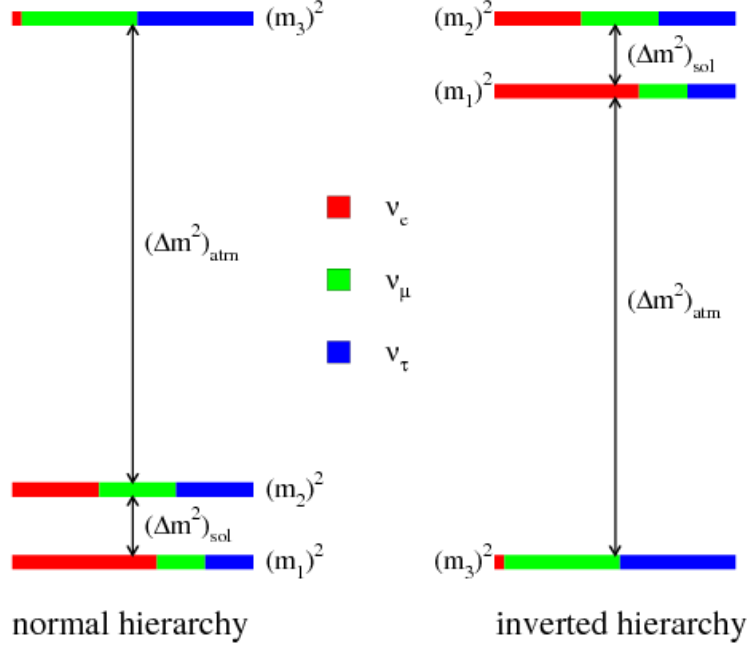


Figure 2.2: Diagram showing the two possible mass orderings with each state showing the approximate measured flavor combination [29].

2.4 ν_e Appearance Oscillation Channel

The oscillation channel that is the subject of this thesis is muon neutrinos oscillating into electron neutrinos ($\nu_\mu \rightarrow \nu_e$), commonly referred to as the ν_e appearance channel. From Equation 2.15 this probability can be written as

$$P_{\nu_\mu \rightarrow \nu_e} = |U_{\mu 1}^* e^{-i \frac{m_1^2 L}{2E}} U_{e 1} + U_{\mu 2}^* e^{-i \frac{m_2^2 L}{2E}} U_{e 2} + U_{\mu 3}^* e^{-i \frac{m_3^2 L}{2E}} U_{e 3}|^2. \quad (2.22)$$

The orthogonality of the rows of the PMNS matrix can be used to eliminate that first term and after simplification the probability becomes

$$P_{\nu_\mu \rightarrow \nu_e} = \left| 2U_{\mu 3}^* U_{e 3} \sin \frac{\Delta m_{31}^2 L}{4E} e^{-i \frac{\Delta m_{32}^2 L}{4E}} + 2U_{\mu 2}^* U_{e 2} \sin \frac{\Delta m_{21}^2 L}{4E} \right|^2. \quad (2.23)$$

After substituting in Equation 2.21, and approximating $\theta_{13} \ll 1$ we arrive at

$$P_{\nu_\mu \rightarrow \nu_e} = \left| \sqrt{P_{atm}} e^{-i \left(\frac{\Delta m_{32}^2 L}{4E} \pm \delta \right)} + \sqrt{P_{sol}} \right|^2 \quad (2.24)$$

$$= P_{atm} + 2\sqrt{P_{atm}}\sqrt{P_{sol}} (\cos \Delta m_{32}^2 \cos \delta \mp \sin \Delta m_{32}^2 \sin \delta) + P_{sol} \quad (2.25)$$

where

$$\sqrt{P_{atm}} = \sin \theta_{23} \sin 2\theta_{13} \sin \left(\frac{\Delta m_{31}^2 L}{4E} \right) \quad (2.26)$$

$$\sqrt{P_{sol}} = \cos \theta_{23} \sin 2\theta_{12} \left(\frac{\Delta m_{31}^2 L}{4E} \right). \quad (2.27)$$

The derived oscillation probability depends on all three mixing angles, both the solar and atmospheric mass splittings and δ . The middle term in Equation 2.25 has both a CP conserving and CP violating component with the - sign used for neutrinos and + for anti-neutrinos. The oscillation probability effectively has two scales depending on the product of $\Delta m^2 \frac{L}{E}$ at the solar and atmospheric mass-splitting. For the NOvA baseline (810 km) and energy (2 GeV) P_{sol} is a 1% effect. In general, due to the differences in the mass splittings and the small value of θ_{13} most experiments are designed in an L/E region so that to first order the result can be evaluated in a simplified two-flavor approximation. It is only in recent years, since θ_{13} was measured, that tests of the full three-flavor parameterization are being performed as discussed in Section 2.9.1. The relations of each term the oscillation formula and the effects on the NOvA measurements will be discussed further in Section




Figure 2.3: Coherent forward scattering for neutrinos on electrons in matter. On left, an electron neutrino scattering via a charged current interaction. On right, any neutrino flavor scattering via a neutral current interaction.

2.10.

2.5 Matter Effect

The oscillation probability derived in the previous section holds true only for a vacuum state. For neutrinos traveling through matter, relevant both for solar neutrinos and accelerator experiments such as NOvA, the oscillation probabilities must be modified. This modification comes from the electron content in matter which provides an additional charged-current channel for coherent forward scattering for electron neutrinos as illustrated in Figure 2.3. This phenomenon was first proposed by Wolfenstein [30] and is now known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect.

To derive the MSW effect for the simplified case of two neutrinos (ν_μ, ν_e) we start by expressing the Schrödinger equation as:

$$i \frac{d}{dL} \begin{bmatrix} \nu_e \\ \nu_\mu \end{bmatrix} = \left[U \begin{pmatrix} \frac{m_1^2}{2E} & 0 \\ 0 & \frac{m_2^2}{2E} \end{pmatrix} U^* + \begin{pmatrix} V_{CC} & 0 \\ 0 & 0 \end{pmatrix} \right] \begin{bmatrix} \nu_e \\ \nu_\mu \end{bmatrix} \quad (2.28)$$

$$= \frac{1}{4E} \begin{bmatrix} -\Delta m^2 \cos 2\theta + 4EV_{CC} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta \end{bmatrix} \begin{bmatrix} \nu_e \\ \nu_\mu \end{bmatrix} \quad (2.29)$$

where V_{CC} is the the additional component to the Hamiltonian from the charged-current

coherent forward scattering of electron neutrons off electrons. This factor takes the form

$$V_{CC} = \pm\sqrt{2}G_F N_e \quad (2.30)$$

with N_e being the number density of electrons in the medium, G_F is Fermi's constant, and the positive sign for neutrinos, the negative for anti-neutrinos. The matrix in Equation 2.29 is the flavor basis Hamiltonian and can be diagonalized according to

$$U_M^T \mathcal{H}_F U_M = \mathcal{H}_M, \quad (2.31)$$

where \mathcal{H}_M is the effective Hamiltonian in the mass basis in matter,

$$\mathcal{H}_M = \frac{1}{4E} \begin{bmatrix} -\Delta m_M^2 & 0 \\ 0 & \Delta m_M^2 \end{bmatrix}, \quad (2.32)$$

and U_M is the effective mixing matrix in matter,

$$U_M = \begin{bmatrix} \cos \theta_M & \sin \theta_M \\ -\sin \theta_M & \cos \theta_M \end{bmatrix}. \quad (2.33)$$

From Equations 2.32 and 2.33

$$\Delta m_M^2 = \sqrt{(\Delta m_0^2 \cos 2\theta_0 \mp 2EV_{CC})^2 + (\Delta m_0^2 \sin 2\theta_0)^2} \quad (2.34)$$

$$\tan 2\theta_M = \frac{\tan 2\theta_0}{1 \mp \frac{2EV_{CC}}{\Delta m_0^2 \cos 2\theta_0}}, \quad (2.35)$$

where Δm_0^2 and θ_0 denote the mass splitting and mixing angle in a vacuum and the minus sign for neutrinos, plus for anti-neutrinos. So the matter effect modifies the oscillation parameters based on the matter density and neutrino energy.

Mikheyev and Smirnov noted [31] that, depending on density, a resonance exists where

the mixing becomes maximal ($\theta_M = 45^\circ$) no matter how small θ_0 is. The resonance condition is satisfied when the matter density is

$$N_e^{res} = \frac{\Delta m_0^2 \cos 2\theta_0}{2\sqrt{2}EG_F}. \quad (2.36)$$

The resonance effect is of particular relevance for neutrinos generated in the Sun or other high density regions. For neutrinos passing through the Earth's crust, as in the NOvA experiment, the density is not near the resonance. The mixing angles in the Earth can be approximated [32], using the fact that both $\frac{\Delta m_{21}^2}{\Delta m_{31}^2}$ and $\sin^2 \theta_{13}$ are small, as

$$\begin{aligned} \Delta m_{31}^2|_M &\approx \Delta m_{31}^2 \mp 2\sqrt{2}G_F N_e E \\ \Delta m_{21}^2|_M &\approx \mp 2\sqrt{2}G_F N_e E \\ \Delta m_{32}^2|_M &\approx \Delta m_{32}^2 \end{aligned} \quad (2.37)$$

with the minus sign for neutrinos, plus for anti-neutrinos. After substituting into Equation 2.25, the full oscillation probability for muon neutrinos oscillating into electron neutrinos is

$$\begin{aligned} P_{\nu_\mu \rightarrow \nu_e} &= \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2 (\Delta m_{31}^2 \mp aL)}{(\Delta m_{31}^2 \mp aL)^2} \Delta m_{31}^4 \\ &+ \sin 2\theta_{23} \sin 2\theta_{13} \sin 2\theta_{12} \frac{\sin (\Delta m_{31}^2 \mp aL)}{(\Delta m_{31}^2 \mp aL)} \Delta m_{31}^2 \frac{\sin (\mp aL)}{\mp aL} \Delta m_{21}^2 \\ &\quad (\cos \Delta m_{32}^2 \cos \delta \mp \sin \Delta m_{32}^2 \sin \delta) \\ &+ \cos^2 \theta_{23} \sin^2 2\theta_{12} \frac{\sin^2 (\mp aL)}{(\mp aL)^2} \Delta m_{21}^4, \end{aligned} \quad (2.38)$$

where the minus sign is for neutrinos, plus for anti-neutrinos and

$$a = \frac{G_F N_e}{\sqrt{2}} \approx \frac{1}{3500 \text{ km}} \quad (2.39)$$

assuming a uniform density in the crust of $3 \text{ g} \times \text{cm}^{-3}$. The result is that the electron neutrino appearance probability is sensitive to the relative sign difference between Δm_{31}^2 and aL and can measure the neutrino mass ordering, discussed in Section 2.10. For the NOvA baseline of 810 kilometers the result is a $\pm 19\%$ effect on the $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation probabilities.

2.6 Atmospheric Neutrino Oscillations

2.6.1 Super-Kamiokande

The first definitive evidence of neutrino oscillations came from the Super-Kamiokande experiment in 1998 [14]. This experiment instrumented a 50 kiloton tank of pure water located one kilometer underground with $\sim 11,000$ photo-multiplier tubes. The experiment measured the flux of atmospheric electron and muon neutrinos both from above (~ 10 km baseline) and below (baseline $\sim 13,000$ km) by detecting the Cherenkov radiation light cones produced by the charged lepton product in a neutrino interaction. The result was an observation of a deficit of muon neutrinos traveling upward through the Earth, while no deficit was seen compared to the prediction for the downward-going muon neutrinos or in either electron neutrino population. This result was interpreted as evidence of the oscillation of muon neutrinos into tau neutrinos since no change in the electron neutrino flux was observed. The zenith angle distributions of e-like and μ -like events for the Super-Kamiokande measurement [33] are in Figure 2.4. In an alternative analysis the data can be binned in units of L/E [34] which highlights the oscillation dip, shown in Figure 2.5. Both approaches produce complementary measurements of the oscillation parameters θ_{23} and Δm_{32}^2 , shown in combination with the accelerator experiments in Figure 2.6.

The oscillation hypothesis is further strengthened by evidence of the appearance of tau neutrinos. This measurement is complicated by the fact that the energy threshold

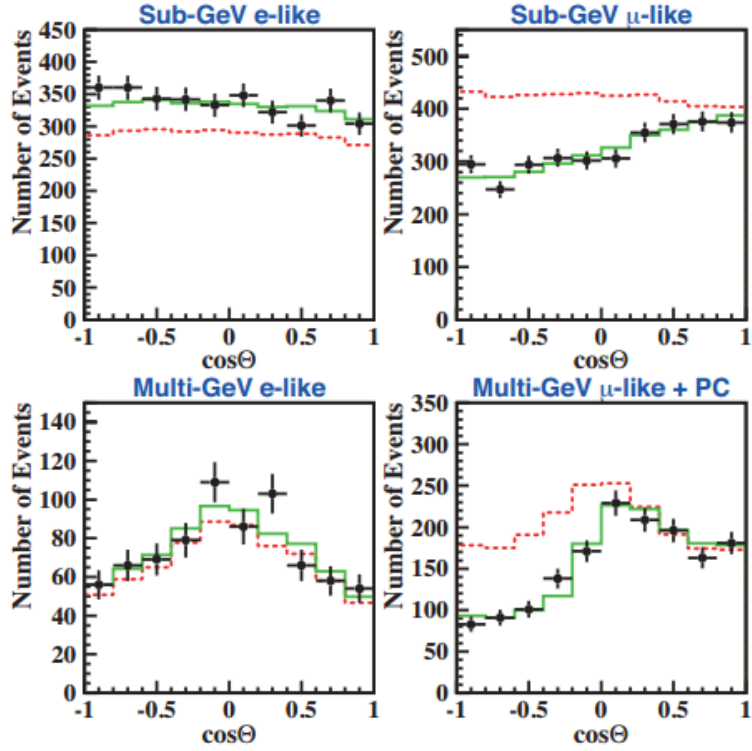


Figure 2.4: Zenith angle distributions for e-like and μ -like events in Super-Kamiokande with visible energy < 1.33 GeV (sub-GeV) and > 1.33 GeV (multi-GeV). The dotted line shows the un-oscillated Monte Carlo prediction and the solid line is the best-fit under the two flavor oscillation hypothesis. Figure from [28].




Figure 2.5: The ratio of μ -like events observed in Super-Kamiokande data to the un-oscillated Monte Carlo prediction in bins of L/E . The solid line indicates the best fit under the two flavor oscillation hypothesis. Figure from [34].

for a charged current ν_τ interaction is 3.4 GeV and the resultant tau lepton immediately decays into a final state with other leptons and mesons. Super-Kamiokande is not able to distinguish a ν_τ interaction on an event-by-event basis, however a multi-variate analysis on the zenith angle distributions of high energy events has shown 3.8σ evidence of ν_τ appearance [35].




Figure 2.6: In the $\sin^2 2\theta_{23}$, Δm_{32}^2 parameter space the 90% C.L. allowed regions are shown the T2K 2011 [36], T2K 2013 [37], Super-Kamiokande [38], and MINOS [39]. Results shown are from both two and three flavor oscillation fits. The MINOS result assumed the same neutrino and anti-neutrino oscillation parameters. Figure was taken from [36].

2.6.2 Accelerator Experiments

Another method to study the atmospheric oscillation scale is to artificially create a beam of neutrinos (or anti-neutrinos) in an accelerator aimed at a detector. This method has the advantage that the baseline and initial neutrino flavor and energy is precisely known. The first experiment to make use of a neutrino beam over a several hundred kilometer baseline was K2K (KEK-to-Kamioka), which produced a 1.3 GeV muon neutrino beam at the KEK accelerator and passed it through both a near detector to measure the flux and Super-Kamiokande 250 kilometers away. The experiment ran from 1999 to 2005 and confirmed, although with lower precision, the oscillation parameters measured by Super-

Kamiokande [40].

The next long-baseline neutrino beam experiment was MINOS (Main Injector Neutrino Oscillation Search). This experiment, first operational in 2005, sent a ~ 3 GeV neutrino beam produced in the NuMI (Neutrinos at the Main Injector) beamline at Fermilab through a near detector and far detector located 735 km away in a mine in Soudan, MN. The beamline used magnetic focusing horns to be able to run in either a muon neutrino or anti-neutrino dominated mode. The near and far detectors are functionally identical tracking-calorimeters consisting of alternating planes of plastic scintillator and steel orthogonal to the beam direction. The detectors have a toroidal magnetic field used to determine the charge and momentum of muons in CC interactions. Over the course of the experiment data was collected in both the ν_μ and $\bar{\nu}_\mu$ disappearance channels with the NuMI beam and atmospheric events in order to precisely measure the atmospheric oscillation parameters and test the assumption that $P_{\nu_\mu \rightarrow \nu_\mu} = P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu}$, shown in Figure 2.7. The result, while still consistent with other experiments, slightly disfavors a maximal value of θ_{23} and shows agreement between neutrinos and anti-neutrinos [39]. The MINOS experiment is discussed further in Section 2.8.2 for the ν_e appearance channel and Section 2.9.1 for the full three flavor oscillation fits.

The first off-axis long-baseline experiment was T2K (Tokai to Kamioka) which began operation in 2010. This experiment operates with a 295 km baseline and again makes use of Super-Kamiokande for the far detector. The muon neutrino beam is produced at the J-PARC main ring, with the detectors located 2.5° off-axis to produce a narrow-band energy beam at the oscillation maximum. This off-axis technique (see Section 3.1.1) is also employed by the NOvA experiment. T2K was designed to measure θ_{13} through the ν_e appearance channel, but also measures the atmospheric oscillation parameters with recent results favoring a maximal value θ_{23} [37]. Measurements of θ_{13} and three flavor oscillation




Figure 2.7: The 90% C.L. of the allowed region from a two flavor oscillation fit to the atmospheric mixing angle and mass splitting in MINOS data. The fit was performed both assuming identical neutrino and anti-neutrino oscillations and fitting each sample separately. Figure was taken from [39].

fits will be discussed in later sections.

One final long-baseline accelerator experiment is the OPERA (Oscillation Project with Emulsion- tRacking Apparatus) experiment. Unlike the other experiments, OPERA was designed to search for the appearance of ν_τ from a ν_μ beam. The experiment uses the 17 GeV CNGS neutrino beam at CERN and located a detector 730 km away in Gran Sasso, Italy. To observe the very short lived tau particle the detector core uses layers of emulsion film and lead to detect the kink where the tau decays into a muon. In the most recent analysis, examining data from 2008 to 2012, OPERA found five tau neutrino events which constitutes 5.1σ evidence of $\nu_\mu \rightarrow \nu_\tau$ appearance [41].

2.7 Solar Neutrino Oscillations

2.7.1 SNO

The definitive experiment in confirming the Standard Model solar neutrino flux prediction and attributing the solar neutrino problem (first described in Chapter 1) to oscillations was SNO (Sudbury Neutrino Observatory), which turned on in 1999. The experiment uses a water Cherenkov detector consisting of 1000 tons of pure heavy water surrounded by photomultiplier tubes 6010 meters water equivalent underground. The experiment measured the solar neutrino flux through three different interactions

$$\nu_e + d \rightarrow p + p + e^- \quad (CC) \quad (2.40)$$

$$\nu_x + d \rightarrow p + n + \nu_x \quad (NC) \quad (2.41)$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (ES) \quad (2.42)$$

where the neutral current interaction is sensitive to all neutrino flavors as is elastic scattering (although ν_e 's dominate the cross-section due to the additional s-channel available), while the charged-current interaction is only relevant for ν_e 's. In order to improve the measurement of NC events the detector ran two years with a “salt phase” by adding two tons of NaCl. The advantages were the higher neutron-capture cross section of chlorine, higher energy photons produced in neutron capture, and isotropic production of photons that contrasts with the directional electrons produced in the other interaction types. The result was that SNO saw the expected ν_e deficit in the CC channel, but in the NC channel the neutrino flux agreed with the Standard Model prediction [15], as shown in Figure 2.8.

Interpreting the SNO result is complicated by the matter effects in the Sun, but results favor the large mixing angle (LMA) solution. A combined fit of global solar neutrino data now shows that the LMA solution uniquely solves the solar neutrino problem at more




Figure 2.8: Measurement of the solar neutrino flux from SNO. The x-axis is the flux from electron neutrinos and the y-axis is the flux from muon and tau flavors. The bands indicate the constraints from the different interactions. The grey band is the elastic scattering constraint from Super-Kamiokande. The best fit is shown with 68%, 95%, and 99% C.L. and is in agreement with the Standard model prediction shown with dotted lines. Figure was taken from [15].

than 5σ [28]. Measurements of the day/night asymmetry in solar ^8B neutrinos in Super Kamiokande which provide evidence of the terrestrial matter effect [42] further constrain the LMA solution. The Borexino experiment measured the lack of a day/night effect in the lower energy ^7Be neutrinos which is also in line with the LMA solution [43]. As the density at the core of the sun is much greater than N_e^{res} , so from Equation 2.35 θ_M is virtually 90° in this solution and thus ν_e 's produced inside the core are nearly entirely in the mass state ν_2 . As neutrinos travel outward through the Sun the density changes smoothly and so the resonance is passed adiabatically, keeping neutrinos in the ν_2 state when they exit. At this point the composition of the solar neutrino flux is $|\nu_2\rangle = \sin\theta_\odot |\nu_e\rangle + \cos\theta_\odot |\nu_x\rangle$, a result that holds upon detection at Earth where oscillation effects have averaged out over the very long baseline. One consequence of the measurement is the solution only works for a specific sign of the mass ordering, by convention ν_2 is chosen as the heavier state and ν_1 the lighter of the pair.

2.7.2 KamLAND

The solar mixing angle and mass splitting were measured in a terrestrial setting with the KamLAND experiment. This experiment used a one kiloton liquid scintillator detector surrounded by photo-multipliers located at the Kamiokande site to measure the flux of $\bar{\nu}_e$'s from nearby nuclear reactors at an average baseline of 180 km. In the solar L/E region KamLAND observed the disappearance of anti-electron neutrinos which combined with the global solar neutrino data to confirm LMA as a unique solution at the 5σ level [44]. KamLAND provided complimentary information to SNO and the solar experiments, precisely measuring the mass splitting while having an ambiguity in the mixing angle octant. Ultimately, KamLAND was able to observe a large enough portion of L/E space to trace out the rise and fall of the oscillation probability, shown in Figure 2.9 [45].

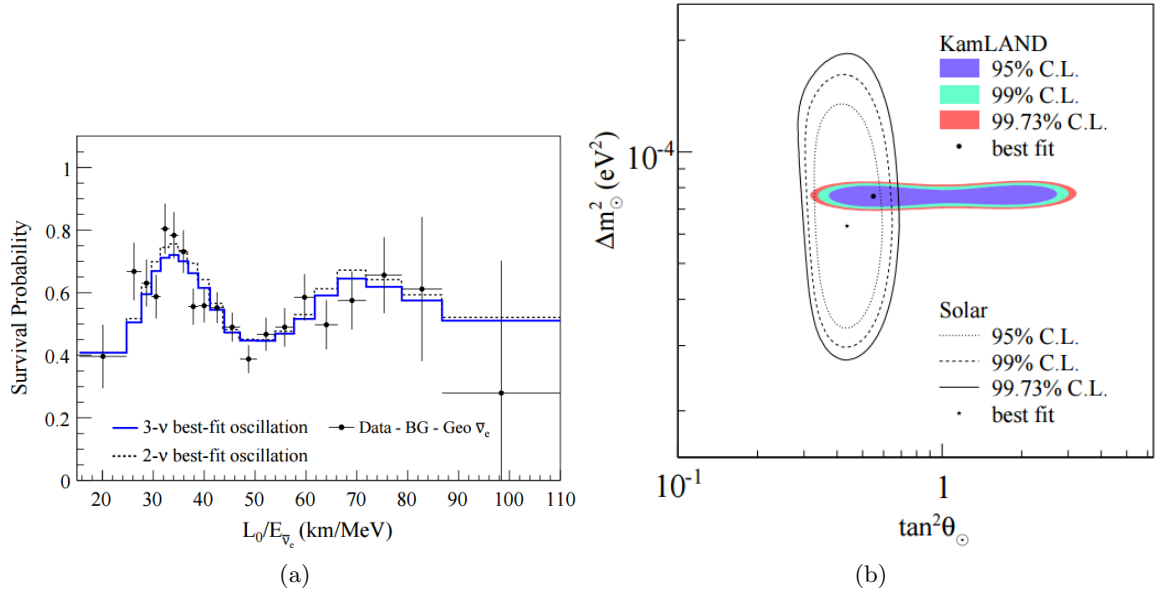


Figure 2.9: a) The ratio of observed $\bar{\nu}_e$ events in KamLAND to the unoscillated prediction (survival probability) as a function of L/E with $L_0 = 180$ km clearly showing the oscillation dip. The blue-line is the best fit to the two-flavor oscillation hypothesis. b) The measured solar neutrino oscillation parameters from KamLAND (colored contours) and SNO (black lines). The two results show agreement, providing complementary information. Both figures from [45].

2.8 Measuring θ_{13}

The last mixing angle to be measured, although it is now the most precisely known of the angles, is θ_{13} . Prior to this measurement it was known that this angle was small compared to θ_{23} and θ_{12} and possibly 0. It has been probed at the atmospheric L/E scale both through reactor experiments measuring $\bar{\nu}_e \rightarrow \bar{\nu}_e$ disappearance and accelerator experiments measuring $\nu_\mu \rightarrow \nu_e$ appearance. Measuring this angle was important as all three mixing angles need be non-zero in order for CP violation in the neutrino sector to be possible.

2.8.1 Reactor Experiments

In KamLAND the $\bar{\nu}_e \rightarrow \bar{\nu}_e$ disappearance channel was studied with nuclear reactors at a baseline where the solar L/E dominated and the atmospheric terms were negligible. Over a shorter baseline (~ 1 km) in the energy range of 3 MeV the opposite is true, allowing the disappearance probability to simplify to

$$P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} \approx 1 - \sin^2 2\theta_{13} \sin^2 \left(\frac{1.267 \Delta m_{31}^2 L}{E} \right) \quad (2.43)$$

in order to probe θ_{13} . The first such reactor experiment was Chooz, named for the nuclear reactor it is situated near in the north of France. The experiment measured $\bar{\nu}_e$'s through the inverse beta decay process with a detector that consisted of a gadolinium doped liquid scintillator inner region to maximize neutron capture and an outer un-doped liquid scintillator region surrounded by photo-multiplier tubes at a baseline of one kilometer. The signal is prompt photons from the positron annihilation followed by delayed gamma-rays with the neutron capture. After two years of running the experiment set a limit in 1999 of $\sin^2 2\theta_{13} \leq 0.1$ at 90% C.L. [46].

A trio of second-generation reactor experiments (Daya Bay, RENO, and Double Chooz)

began publishing results in 2012. Although the detectors vary in size and number, all experiments utilize a very similar detector design consisting of an inner gadolinium-doped liquid scintillator volume surrounded by a gadolinium-free liquid scintillator “ γ -catcher” volume, and then an outer veto layer of water (Daya Bay and RENO) or liquid scintillator (Double Chooz). The first and most precise result came from Daya Bay. The Daya Bay experiment measures the neutrino flux from six reactor cores with six functionally identical detectors spread between two near detector halls (at a flux weighted average of 470 and 576 meters) and a far detector location at a flux weighted average of 1648 meters. The first result in March 2012 with 55 days of data was a measurement of $\sin^2 2\theta_{13} = 0.092 \pm 0.016$ (stat) ± 0.005 (syst) which was 5.2σ evidence that θ_{13} was non-zero [47], shown in Figure 2.10. In 2013 Daya Bay further constrained the measurement with more data to $\sin^2 2\theta_{13} =$




Figure 2.10: Ratio of the measured to predicted event rate assuming no oscillations for each detector in Daya Bay. The red curve shows the best fit oscillation curve, inset is the χ^2 vs $\sin^2 2\theta_{13}$. Figure was taken from the first published results in 2012 [47].

$0.090^{+0.008}_{-0.009}$ ($\theta_{13} = 8.7^\circ \pm 0.4^\circ$) [48] and in 2014 with 621 days of data showed a result at the Neutrino 2014 conference of $\sin^2 2\theta_{13} = 0.084 \pm 0.005$ [49]. The experiment expects a final precision after two to three additional years of running of 2-3% (0.2°).

The RENO (Reactor Experiment for Neutrino Oscillation) experiment uses six reactor cores in Yonggwang, South Korea spaced equidistant along 1280 meters with a near detector located 290 meters perpendicular to the center of the array and a far detector at 1380 meters. In April 2012 with 229 days of data the experiment measured $\sin^2 2\theta_{13} = 0.113 \pm 0.013 \pm 0.019$ which rejected the null hypothesis at 4.9σ [50]. The experiment presented updated results in 2013 [51] and 2014 [52] for a current result with ~ 795 livedays of data of $\sin^2 2\theta_{13} = 0.101 \pm 0.008 \pm 0.010$.

Double Chooz features an upgraded detector design from the Chooz experiment with a far detector located 1050 meters from the two reactor cores and a near detector at 415 meters that came online in late 2014. The first results in November 2011 with 101 days of far detector data ruled out the null oscillation hypothesis at 1.9σ with a value of $\sin^2 2\theta_{13} = 0.086 \pm 0.041 \pm 0.030$ [53] and improved the measurement in June 2012 with 229 days of data to $\sin^2 2\theta_{13} = 0.109 \pm 0.030 \pm 0.025$ [54]. One unique ability of Double Chooz is to simultaneously turn off both reactor cores, not likely to be done at the larger power plants used by RENO and Daya Bay, and take a background only measurement to improve the result. In 2014 with 467 days of data the experiment improved the measurement to $\sin^2 2\theta_{13} = 0.090^{+0.032}_{-0.029}$ [55] which is consistent with the other reactor measurements and expected to improve with the addition of data from the new near detector to reduce uncertainties. A summary of the progression of measurements of $\sin^2 2\theta_{13}$ in recent years can be seen in Figure 2.11.

2.8.2 θ_{13} at Accelerator Experiments

For long-baseline accelerator experiments (currently T2K, MINOS+, and NOvA) θ_{13} is measured through the oscillation channel $\nu_\mu \rightarrow \nu_e$ shown in Equation 2.38. This probability depends on the mass ordering, θ_{23} and δ which can significantly alter the probability. While




Figure 2.11: Summary of measurements of $\sin^2 2\theta_{13}$ between 2011 and 2013 for both reactor and accelerator experiments (assuming $\delta = 0^\circ$ and $\theta_{23} = 45^\circ$). Daya Bay provides the world-leading constraint on the measurement. Figure was taken from [56].

θ_{13} cannot be precisely measured as with reactor experiments, it can be constrained to be non-zero. The measured value of θ_{13} from the reactor experiments can be used as a constraint in fitting the data from accelerator experiments to probe the mass ordering and δ phase, discussed further in Section 2.9.1.

T2K produced the first evidence of ν_e appearance in an accelerator beam in June 2011 when six events were observed in 1.43×10^{20} protons on target (POT) with a background expectation of 1.5 ± 0.3 events, giving 2.5σ evidence of a non-zero θ_{13} [57]. Updated results were produced in 2013 [58,59] with a current measurement of 28 events in 6.57×10^{20} POT with a background prediction of 4.92 ± 0.55 events. Assuming $\sin^2 \theta_{23} = 0.5$ and $\delta = 0$ the best-fit for the normal (inverted) mass ordering is $\sin^2 2\theta_{13} = 0.140_{-0.032}^{+0.038}$ ($\sin^2 2\theta_{13} = 0.170_{-0.037}^{+0.045}$) which is 7σ evidence of non-zero θ_{13} , shown in Figure 2.12. With the reactor measured value of $\sin^2 2\theta_{13} = 0.09$ and assuming no matter effect or CP violation the expectation would have been an observation of 15 signal events. This tension is resolved in Section 2.9.1 by fitting for mass ordering and CP violation using reactor constraints on θ_{13} and atmospheric

constraints on θ_{23} and $|\Delta m_{32}^2|$.

MINOS is less suited to measure ν_e appearance as it sees a broad-band beam with a large neutral-current background and the dense detector design with steel plates compresses the electron showers making them harder to distinguish from that background. Nonetheless in 2011 MINOS observed 62 events on a background of 50 to disfavor $\theta_{13} = 0$ at the 89% C.L. [60]. In 2013 the dataset was expanded to include 10.6×10^{20} POT in neutrino running and 3.3×10^{20} POT in anti-neutrino running. Combining neutrino and anti-neutrino data 172 events were observed with a background expectation of 145 [61], shown in Figure 2.12. If a normal mass ordering, maximal θ_{23} , no CP violation, and $\sin^2 2\theta_{13} = 0.1$ were assumed the expectation would have been 183 events, causing slight tension with the T2K result which has a discrepancy in the opposite direction although the disagreement is below 1σ . The MINOS detectors are still operating but the experiment is now known as MINOS+ since the NuMI beamline was upgraded to produce a higher energy spectrum optimized for the NOvA experiment. NOvA is collecting data now and the first results on $\nu_\mu \rightarrow \nu_e$ appearance are the subject of this thesis.

2.9 Summary of Results

2.9.1 Three Flavor Analysis of MINOS and T2K

With both the MINOS and T2K experiments the ν_e appearance and ν_μ disappearance data can be fit jointly with θ_{12} and Δm_{21}^2 input from the solar experiments and the precise reactor experiments used to constrain θ_{13} . The resulting fits can explore the three-parameter interplay of δ , θ_{23} octant, and the mass ordering. The current best-fit results from MINOS [63] and T2K [64] exhibit a small amount of tension in both the atmospheric parameter space (Δm_{32}^2 , $\sin^2 \theta_{23}$) shown in Figure 2.13, and the choice of mass ordering and δ shown in Figure 2.14.




Figure 2.12: a) T2K 68% and 90% C.L. allowed regions for $\sin^2 2\theta_{13}$ assuming normal (top) or inverted mass ordering as a function of δ . Constraints on $\sin^2 \theta_{23}$ and Δm_{32}^2 come from [37] and the shaded region shows the PDG2012 best-fit $\sin^2 2\theta_{13}$ from the reactor experiments [62]. Image from [59]. b) MINOS 68% and 90% C.L. regions for $2\sin^2 2\theta_{13}\sin^2 \theta_{23}$ depending on mass ordering choice, δ , and θ_{23} octant. Previous measurements are used to place constraints in the fit, $\sin^2 2\theta_{23} = 0.957^{+0.035}_{-0.046}$ and $|\Delta m_{32}^2| = (2.39^{+0.09}_{-0.10}) \times 10^{-3} \text{eV}^2$. Image from [61].




Figure 2.13: In the $(\sin^2 \theta_{23}, \Delta m^2_{32})$ parameter space the 68% (dotted) and 90% (solid) C.L. regions for the normal (top) and inverted (bottom) mass ordering when $\sin^2 2\theta_{13}$ is constrained by the reactor data. The T2K fit is compared to Super Kamiokande [65] and MINOS [63], figure taken from [64].




Figure 2.14: The $-2\Delta\ln(L)$ as a function of δ after marginalizing over $\sin^2 2\theta_{13}$, $\sin^2 \theta_{23}$, and Δm_{32}^2 for different mass ordering and θ_{23} octant choices. The reactor data provided a constraint in each fit of $\sin^2 2\theta_{13}$. a) T2K fit where the solid (dotted) line with markers is the Feldman-Cousins 90% C.L. limits for the normal (inverted) mass ordering. The constrain on θ_{23} comes from the T2K muon neutrino disappearance measurement [64]. Figure from [59]. b) The MINOS fit where the constraint on θ_{23} comes from the MINOS muon neutrino disappearance measurement [39]. Figure from [63].

The MINOS experiment has a slight preference for an inverted mass ordering (0.23 units of $-2\Delta\log(L)$) and the lower octant with $\sin^2 \theta_{23} = 0.41$, however the 90% C.L. is still broad ($\sin^2 \theta_{23} = 0.34 - 0.67$) and maximal mixing ($\sin^2 \theta_{23} = 0.5$) is only disfavored at $-2\Delta\log(L) = 1.54$. For the best-fit combination of inverted mass ordering and the lower octant $\delta/\pi = 0.62$ although no value of δ is ruled out at 90%. T2K has the most precise measurement on $\sin^2 \theta_{23}$ and favors near maximal mixing, $\sin^2 \theta_{23} = 0.514^{+0.055}_{-0.056}$. T2K favors the normal mass ordering near the 1σ level and $\delta = 3\pi/2$ in order to account for the excess ν_e appearance signal compared to the reactor experiment measurement. At the 90% C.L. the δ range from 0.35π to 0.63π (0.09π to 0.90π) is ruled out for the normal (inverted) mass ordering. So the T2K fit disfavors the MINOS mass ordering and CP phase choice at the 90% level. Additional data from both experiments and results from NOvA will be important to clarify these discrepancies.

2.9.2 Experimental State of the PMNS Matrix

A global fit of existing solar, atmospheric, accelerator, and reactor experimental data in late 2013 was performed [66] to jointly fit all oscillation parameters in light of all three mixing angles being non-zero with results shown in Table 2.1. While θ_{23} was the first measured mixing angle, it is now the least precisely known and at the same time θ_{13} has become the most precisely constrained parameter in only two years. MINOS and atmospheric neutrino experiments have a preference for θ_{23} in the lower octant, but T2K data pushes the fit such that maximal mixing cannot be eliminated. The joint analysis prefers the T2K solution of $\delta \approx 3\pi/2$ which would indicate maximum CP violation, but much of the parameter space is still viable.

Parameter	best-fit ($\pm 1\sigma$)	3σ
Δm_{21}^2 [10^{-5}eV^2]	$7.54_{-0.22}^{+0.26}$	6.99 - 8.18
$ \Delta m^2 $ [10^{-3}eV^2]	2.43 ± 0.06 (2.38 ± 0.06)	2.23 - 2.61 (2.19 - 2.56)
$\sin^2 \theta_{12}$	0.308 ± 0.017	0.259 - 0.359
$\sin^2 \theta_{23}, \Delta m^2 > 0$	$0.437_{-0.023}^{+0.033}$	0.374 - 0.628
$\sin^2 \theta_{23}, \Delta m^2 < 0$	$0.455_{-0.031}^{+0.039}$	0.380 - 0.641
$\sin^2 \theta_{13}, \Delta m^2 > 0$	$0.0234_{-0.0019}^{+0.0020}$	0.0176 - 0.0295
$\sin^2 \theta_{13}, \Delta m^2 < 0$	$0.0240_{-0.0022}^{+0.0019}$	0.0178 - 0.0298
δ/π (2σ range quoted)	$1.39_{-0.27}^{+0.38}$ ($1.31_{-0.33}^{+0.29}$)	(0.00 - 0.16) \oplus (0.86 - 2.00) ((0.00 - 0.02) \oplus (0.70-2.00))

Table 2.1: Best-fit values and 3σ allowed ranges for three-flavor oscillation parameters derived from a global fit [66]. The values (values in parentheses) correspond to normal (inverted) mass ordering choice and $\Delta m^2 = m_3^2 - (m_2^2 - m_1^2)/2$.

The current large-angle mixing picture of the PMNS matrix with θ_{23} approximately

maximal ($\pi/4$), $\theta_{12} \cong \pi/5.4$ and $\theta_{13} \cong \pi/20$ is in contrast to the CKM matrix for quark mixing which features small angle mixing. For comparison, the magnitude of CP violation can be described in a standard way for both matrices with the Jarlskog invariant J_{CP} [67]. In the general oscillation formula (Equation 2.18) CP violation is generated through the imaginary term. For three-flavor oscillations all forms of that term where $\alpha \neq \beta$ and $k \neq j$ coincide up to a sign due to the matrix unitarity. A rephasing invariant can be factored away from the baseline dependence on L/E :

$$J_{CP} = \Im (U_{\mu 3} U_{e 3}^* U_{e 2} U_{\mu 2}^*) \quad (2.44)$$

$$= \frac{1}{8} \cos \theta_{13} \sin 2\theta_{12} \sin 2\theta_{23} \sin 2\theta_{13} \sin \delta. \quad (2.45)$$

For the current best-fit values in Table 2.1 J_{CP} for the PMNS matrix is -0.032 (-0.029) for the normal (inverted) mass ordering. If it remains true that δ is a nearly maximal effect then CP violation in the neutrino sector would be relatively large, approximately 1000 times the magnitude of CP violation in the quark sector ($J_{CP}^{CKM} = (3.06_{-0.20}^{+0.21}) \times 10^{-5}$ [28]).

2.9.3 Unanswered Questions in Neutrino Physics

Neutrino physics has been a rapidly progressing field with definitive proof of oscillations coming less than 20 years ago and the final mixing angle first measured in the past three years. As neutrino oscillation experiments enter the precision era there are several important questions that remain to be answered,

- Is the mass ordering normal or inverted? This question is important for narrowing the parameter space being searched by neutrinoless double beta decay experiments to probe both the Majorana nature of the particle and the absolute mass scale.
- Is δ non-zero? As discussed previously, this condition is required for the seesaw mechanism to be a viable theory for neutrino mass generation and matter-antimatter

asymmetry.

- Is θ_{23} maximal (45°) and if not does ν_3 mix more with ν_μ ($\theta_{23} > 45^\circ$) or ν_τ ($\theta_{23} < 45^\circ$)?

This is important for understanding the texture of the PMNS matrix.

- Are neutrinos Dirac or Majorana particles? This question is being studied through neutrinoless double beta decay and the other requirement for the seesaw theory is that neutrinos be Majorana particles.
- What is the absolute value of the neutrino mass scale? Currently only the relative mass differences are known and upper limits have been set.
- Are there sterile neutrino flavors in addition to the three active neutrinos?

By measuring $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$ (as well as the anti-neutrino channels) NOvA is directly sensitive to the first three questions.

2.10 Measuring ν_e Appearance with NOvA

As discussed previously, the oscillation probability for $\nu_\mu \rightarrow \nu_e$ (Equation 2.38) depends on all three mixing angles, both mass splittings, δ , and the mass ordering choice. The event rates observed in the NOvA experiment could vary by as much as $\pm 60\%$ depending on the interplay of these parameters as shown in Figure 2.15. Some parameters alter the oscillation probabilities measured for neutrinos and anti-neutrinos. At the NOvA baseline of 810 kilometers the matter effect provides a 20% enhancement (suppression) of the electron neutrino (anti-neutrino) appearance probability when compared to the vacuum benchmark. The effect is the opposite for the inverted mass ordering. In a similar fashion δ varies the probabilities for neutrinos and anti-neutrinos. As illustrated in the figure there are some combinations, such as normal mass ordering and $\delta = 3\pi/2$, where both effects go in the same direction making it possible to measure both parameters (illustrated with the star point). Other degenerate regions exist where the oscillation parameters cannot be decoupled by

NOvA. Additionally, the θ_{23} value also slides the oscillation probabilities. To have a chance at uniquely determining all three parameters an experiment must operate in both neutrino and anti-neutrino mode. The figure assumes a standard NOvA running plan of three years (18×10^{20} POT) in each configuration.




Figure 2.15: Measured oscillation probabilities in the NOvA experiment for $P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_e}$ vs $P_{\nu_\mu \rightarrow \nu_e}$ assuming 18×10^{20} POT in each neutrino and anti-neutrino running. The central starred point indicates a benchmark value assuming no matter effect, maximal θ_{23} , no CP violation, and $\sin^2 2\theta_{13} = 0.095$. The matter effect splits the oscillation in red (inverted mass ordering) and blue (normal mass ordering) ellipses dependent on the δ phase. The probabilities are further split depending on the θ_{23} octant, with values drawn assuming $\sin^2 2\theta_{23} = 0.97$. The starred point with 1 and 2 σ contours illustrates that mass ordering, δ and octant can be measured simultaneously through these two oscillation channels.

CHAPTER 3

The NOvA Experiment

3.1 NuMI Beam

The experiment makes use of the NuMI (Neutrinos at the Main Injector) beamline at Fermilab to provide the neutrino beam. The NuMI beam was originally constructed for the MINOS experiment which began collecting data in 2005 [68]. At the Fermilab accelerator complex, drawn in Figure 3.1, the origin of the beam is a gaseous hydrogen source where H^- ions are produced and accelerated to 400 MeV in the linac before being fed into the Booster ring. Here the electrons are stripped and the protons are accelerated to 8 GeV in the 75.47 meter radius synchrotron. The beam in the Booster is bunched at 53 MHz to produce batches of $\approx 4 \times 10^{12}$ protons. These batches are then injected into the Main Injector (MI) synchrotron using a slip-stacking method where six successive batches are injected in a train followed by six more in a different orbit [69]. Once this twelve batch group is loaded into the MI it is accelerated to 120 GeV in the 528.30 meter radius synchrotron and then extracted to the NuMI beamline. The MI injection and acceleration cycle is 1.3 seconds following upgrades discussed in Section 3.1.2. The numbers in this section are the final performance specifications after upgrades are finished in 2016.

Once the beam is extracted from the MI it is bent downward at an angle of 58 milliradians in the direction of the on-axis MINOS detector located in the Soudan mine in Minnesota. The protons collide with a graphite target producing a secondary meson beam




Figure 3.2: Schematic of the NuMI beam [70].

3.1.1 Beam Flux and Kinematics

The NOvA detectors are located 14 milliradians off-axis from the NuMI beamline as opposed to the on-axis MINOS experiment. This choice was made due to the decay kinematics of pions and kaons in order to optimize sensitivity to the $\nu_\mu \rightarrow \nu_e$ oscillation channel.

In the rest frame the pions and kaons decay isotropically, $(\pi, K) \rightarrow \mu + \nu$, producing mono-energetic neutrinos. When boosted into the lab frame the flux, F , and energy, E_ν , of neutrinos produced from the meson decay in flight and measured at a detector of area A located at a distance z is:

$$F = \left(\frac{2\gamma}{1 + \gamma^2\theta^2} \right)^2 \frac{A}{4\pi z^2} \quad (3.1)$$

$$E_\nu = \frac{1 - \frac{m_\mu^2}{m_{\pi,K}^2}}{1 + \gamma^2\theta^2} \quad (3.2)$$

where θ is the angle between the incoming meson and outgoing neutrino and $\gamma = \frac{1}{\sqrt{1-\beta^2}}$. The result of these kinematics is that the outgoing neutrino energy distribution is relatively flat across a broad range of meson energies for small angles, as shown in Figure 3.3. For the selected detector angle of 14 milliradians the NuMI beam produces a relatively narrow flux peaked at 2 GeV with roughly five times more events than would be seen at that energy on-axis as shown in Figure 3.4. This is important because the neutrino energy that maximizes the oscillation probability of electron neutrino appearance in a muon neutrino beam is approximately 1.6 GeV for the NOvA baseline of 810 km.

Beyond increasing the rate of signal events, two backgrounds are reduced by using




Figure 3.3: a) The neutrino flux as function of pion energy viewed from an angle θ with respect to the beam axis. The flux was normalized to 810 km. b) The energy of neutrinos produced at an angle θ relative to the pion direction as a function of pion energy. To select the neutrino energy that maximizes the oscillation of ν_μ 's, NOvA chose a 14 milliradian angle drawn as a horizontal dashed line. Figure from [70].




Figure 3.4: The top panel depicts the charged-current ν_μ event rates prior to oscillation 810 km from Fermilab for a detector located at various off-axis angles. The bottom panel depicts the oscillation probability $\nu_\mu \rightarrow \nu_e$ as a function of neutrino energy for that distance [70].

the narrow off-axis beam. The first is neutral current events where the outgoing lepton (the neutrino) is not observed, carrying away an unknown amount of energy. These event topologies, which frequently contain electromagnetic showers from $\pi^0 \rightarrow \gamma\gamma$ decay, can give the appearance of ν_e signal events, as discussed further in Chapter 5. Due to the missing neutrino energy, higher energy neutral current events can shift down into the expected signal region. The off-axis detector location sees a narrow beam in the signal region with a reduced high energy tail, which results in much of the neutral current background shifting below the signal region, as shown in Figure 3.5. The second background of concern is the intrinsic ν_e component of the neutrino beam before oscillations that comes from muon and kaon decay. These are three-body decay processes which produce a broader spectrum than the relatively narrow signal of ν_e appearance when located off-axis.




Figure 3.5: Simulated energy distributions for the ν_e oscillation signal, intrinsic beam ν_e events, neutral-current events and ν_μ charged-current events with and without oscillations. No selection efficiencies or mis-identification rates are taken into account. The simulation used $\delta m_{32}^2 = 2.5 \times 10^{-3} \text{eV}^2$, $\sin^2(2\theta_{13}) = 0.10$. An off-axis angle of 14 milliradians and oscillation distance of 810 km was assumed [70].

When the magnetic horns are running with forward horn current (neutrino mode) Figure 3.6 shows that the wrong-sign contamination from $\bar{\nu}_\mu$ is 1.7% (1.8%) for the Far Detector (Near Detector) in the 1 to 3 GeV region. The intrinsic $(\nu_e + \bar{\nu}_e)$ background is 0.6% (0.7%)

in the Far Detector (Near Detector). When running with reverse horn current (anti-neutrino mode) Figure 3.7 shows the wrong-sign contamination from ν_μ is 11.3% (11.7%) for the Far Detector (Near Detector). The wrong-sign contamination is higher in anti-neutrino mode due to the smaller production and interaction cross-sections for anti-neutrinos.

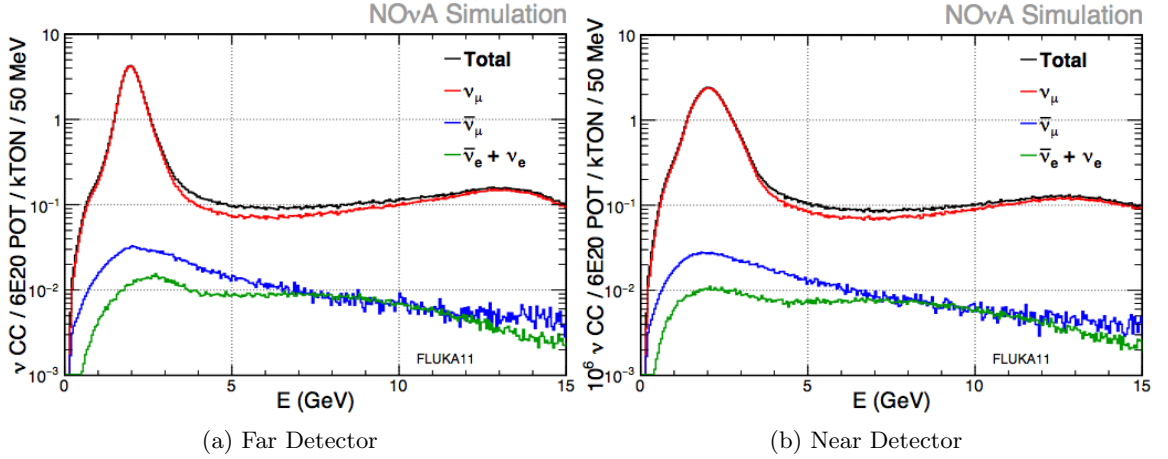


Figure 3.6: Flux times cross-section projections for the NOvA Far Detector (a) and Near Detector (b) of the ν_μ , $\bar{\nu}_\mu$, and $(\nu_e + \bar{\nu}_e)$ components of the beam in forward horn current (neutrino dominant) mode. Oscillation weights have not been applied. Figure from [72].

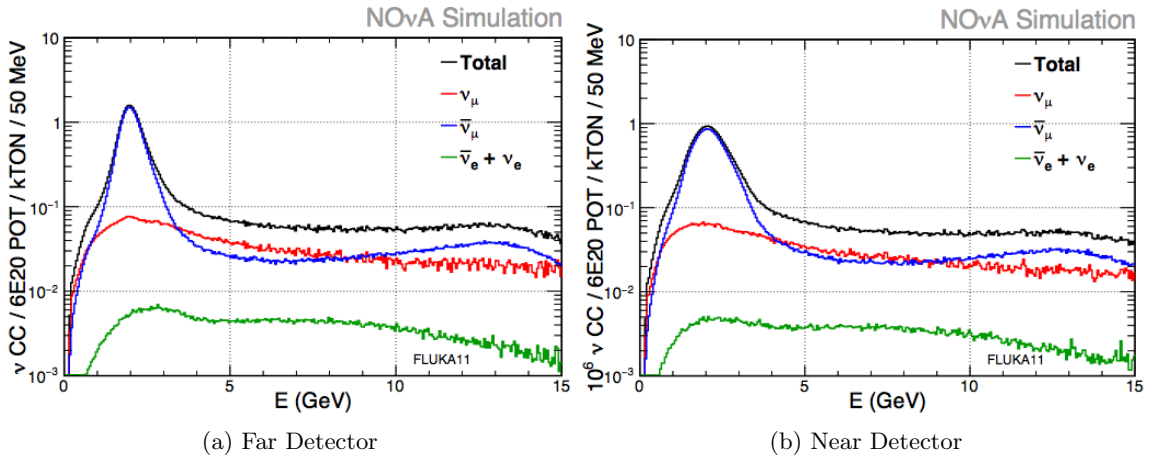


Figure 3.7: Flux times cross-section projections for the NOvA Far Detector (a) and Near Detector (b) of the ν_μ , $\bar{\nu}_\mu$, and $(\nu_e + \bar{\nu}_e)$ components of the beam in reverse horn current (anti-neutrino dominant) mode. Oscillation weights have not been applied. Figure from [72].

3.1.2 NuMI Upgrade

To achieve the goals of the NOvA physics program the NuMI beam is undergoing an upgrade from 350 kW to 700 kW of power. This required upgrades to many parts of the accelerator complex at Fermilab. Much of the work to make the system capable of handling and delivering the increased power occurred during a long accelerator shutdown between May of 2012 and September of 2013. The beam was commissioned and operated typically between 200 to 280 kW between March and September of 2014. After another accelerator shutdown the beam began operating at 300 kW in November 2014 and ramping to 400 kW by April 2015 with a peak of 520 kW, which is the world's most powerful neutrino beam. The full 700 kW beam power is expected in 2016 after upgrades to the slip-stacking and radio-frequency (RF) cavities in the Booster ring are complete. The beam exposure during the analysis period is discussed in Section 6.6.

To increase power the period between beam spills extracted to the NuMI target was reduced from 2.2 seconds to 1.33 seconds. This was done by adding two additional RF cavities to increase the acceleration rate from 204 GeV/sec to 240 GeV/sec and changing the way beam is loaded into the MI. The Recycler storage ring occupies the same tunnel as the MI and was previously used as an anti-proton storage ring when the Tevatron was operational. It has been converted to a proton ring and the Booster can slip stack 12 batches of protons into the Recycler. The beam is then extracted in a single turn from the Recycler into the MI to be ramped up to 120 GeV. This process reduces the cycle time since the Recycler can prepare the next batch while the MI is ramping as shown in Figure 3.8. The NuMI target and cooling was upgraded to handle the higher beam power. To produce a 2 GeV neutrino beam at 14 milliradians the target is positioned to begin 135 cm upstream of the first horn and end 15 cm upstream of the horn. The second focusing horn was moved from 10 meters downstream of the first horn to 23 meters. These changes result in the

on-axis neutrino beam energy shifting from an average of 4 GeV to 7.5 GeV.

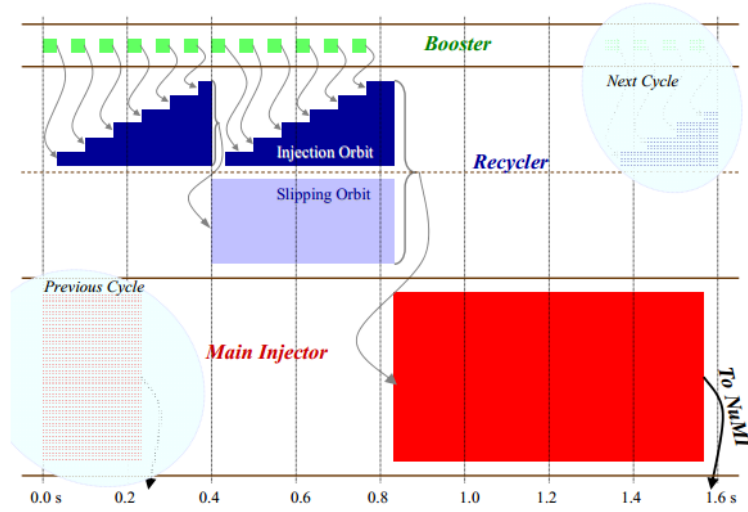


Figure 3.8: Timing schematic for creating a NuMI neutrino spill after the beamline upgrade is complete. Figure from [70].

3.2 The NOvA Detector Design

The NOvA detectors are functionally identical in order to cancel systematic uncertainties in the analysis. Both detectors consist of extruded cells of PVC plastic filled with liquid scintillator to form a three dimensional tracking calorimeter. The general detector design will be described first and then specifics for the Near and Far Detectors in the following sections.

The PVC plastic is extruded in groups of 16 cells with a cell width of 3.9 centimeters and a depth of 5.9 centimeters. These extrusions are stacked in alternating horizontal and vertical planes, shown in Figure 3.9, to provide three dimensional tracking. The cells are coated with titanium dioxide, TiO_2 , which is 90% reflective for 430 nanometer wavelength light. The PVC cells provide the structural support for the detector. In liquid scintillator the radiation length for an electromagnetic shower is 41 centimeters with a Moliere radius of 10.5 centimeters and the mean free path for photon conversion is 53 centimeters. The

low-Z, fine-grained detector allows typical electron showers in the signal region to traverse 10-80 planes and photons travel on order of 6 planes before converting which is important for background rejection.




Figure 3.9: Cutaway diagram of the orthogonal planes of cells in the detector. Not pictured is the 0.7 mm wavelength shifting fiber inside each cell. In reality the cells have rounded edges. Figure from [73].

The cells are filled with a liquid scintillator that is by mass 94.63% mineral oil, 5.23% pseudocumene (scintillator), 0.14% PPO (wavelength shifter), 0.0016% bis-MSB (wavelength shifter), 0.001% Stadis-425 (anti-static), and 0.001% Vitamin E (anti-oxidant) [74]. This solution produces scintillation light in the near ultraviolet and shifts it to the visible region of 380-450 nm. Within each cell a 0.7 millimeter double-clad Kuraray wavelength-shifting (WLS) fiber is looped down the entire cell length in a U shape. Both ends of the fiber are read out by a single photodetector which improves collection efficiency. The fiber absorbs light in the violet-blue range and emits in the blue-green (450-600 nm) range.

Pairs of 16 cell PVC extrusions are glued together side-by-side to form a 32 cell module. One end of the module is capped with a reflective plastic seal. Each end of the wavelength shifting fiber in each cell is connected to a pixel on a 32 pixel Hamamatsu avalanche photodiode (APD), pictured in Figure 3.10. The APD was chosen for its high quantum efficiency, 85%, in the green part of the spectrum. High quantum efficiency is desired in order to make long cells, 15.6 meters in the far detector, and still see a minimum signal of

20 photoelectrons as specified in the technical design criteria. A comparison of APDs to photomultiplier tubes (PMTs) for the NOvA light spectrum is shown in Figure 3.11. The APDs are operated at -15°C to reduce thermal noise. The voltage is determined individually for each channel for a gain of 100 and is approximately 400 volts. The noise thresholds are set around 10 photoelectrons. A thin and transparent paralene coating is put on the surface of each APD to isolate from humidity. A system of tubes flow dry nitrogen gas through the APD housings to keep ice from forming on the surface.




Figure 3.10: a) The ends of 32 wavelength-shifting fibers collected at the end of scintillation cells to mount to an APD. b) Front face of an APD that will be pressed against the fiber ends.

APDs are attached to a front-end board (FEB) which provides voltage and cooling and to digitize the signal. The FEB amplifies and shapes the discrete charge, q , according to

$$F(t) = q \times e^{\frac{t}{t_f}} \times \left(1 - e^{\frac{t}{t_r}}\right) \quad (3.3)$$

where t_r and t_f are the rise and fall times respectively. The FEB samples each cell at a rate of 2 MHz for the Far Detector and 8 MHz for the Near Detector and triggers on samples




Figure 3.11: The quantum efficiency of APDs (magenta) and PMTs (blue) are shown in comparison to the spectrum of light emitted from the wavelength shifting fiber at distances between 0.5 and 17 meters. Figure from [70].

above threshold based on a dual correlated sampling (DCS) algorithm. Each ADC sample is associated with the time in a register on the FEB that is incremented from an absolute master clock that keeps all electronics on the detector in sync both internally and externally with respect to the beam. This method allows the FEB to operate with 0% deadtime since a readout gate never has to be stopped or started. More information on the timing system, digitization, and subsequent timing resolution is in Section 4.6.

A group of up to 64 FEBs send hits above threshold to a data concentrator module (DCM). The DCM collects hit information from its readout region and condenses the data in 50 microsecond blocks (microslice) that the DCM uses to build a larger 5 millisecond block (millislice). The DCMs transfer data to a buffer farm at a rate of 24 Mb/s in the form of millislices [75]. One hundred percent of hits that are above threshold are digitized and stored in the buffer farm for as long as 20 minutes before being erased from memory before any trigger decision is made. Signals from the accelerator indicating the time of a beam spill arrive and start the readout of microslices from the buffer in the selected time range to create an event record that is saved for permanent processing, see Section 4.5. While data is in the buffer it is processed through a series of fast algorithms that can trigger the recording of additional blocks of time that meet selection criteria. This data-driven trigger

approach is used for exotic searches such as monopoles, super nova neutrinos and indirect dark matter.

3.3 The Far Detector

The NOvA Far Detector is located in Ash River, Minnesota 810 kilometers from the NuMI target at Fermilab at an angle of 14 milliradians with respect to the beam center. The detector is located on the surface and uses a cosmic-ray shield made of 2.5 feet of precast concrete topped with 1.5 feet of cast-in-place concrete topped with six inches of loose barite rocks (barium-loaded concrete). The barite is a high-Z material effective in shielding photons. The combined shield is effectively 14 radiation lengths thick. The detector is $15.6 \times 15.6 \times 60$ meters alternating planes of horizontal and vertical modules, each plane consisting of 12 modules 15.6 meters in length. In all there are 344,064 individual channels in 896 planes and a total detector mass of 13.968 kilotons (65% liquid scintillator) [76]. The detector was constructed in 28 blocks, each consisting of 32 planes glued together. The electronics was instrumented in 64 plane regions called diblocks. Each diblock has 12 DCMs, six on each side, with each DCM containing 64 FEBs from two rows of modules, shown in Figure 3.12. During the construction phase of the experiment the detector was instrumented modularly in diblocks, which allowed data to be collected with a partial detector. Figure 3.13 shows Far Detector performance in two diblocks in December 2013 in terms of efficiency in detecting muon hits and energy deposited at the far end (15.6 meters from the APD) of cells.

3.4 The Near Detector

The Near Detector is located one kilometer downstream of the NuMI target and 100 meters underground at Fermilab. The detector is built using the same modules as the Far Detector,




Figure 3.12: Pictured at the Far Detector are the tops of vertical 32 cell modules (A), with the end of the horizontal planes visible (B). FEBs (C) are mounted on each module with an APD housed inside. The FEBs transmit data to DCMs (D) and are powered by power supplies (E).

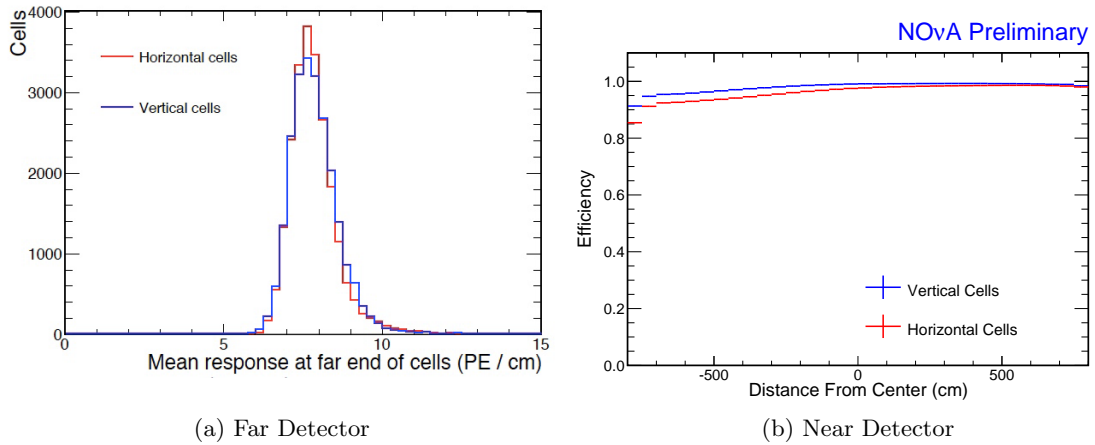


Figure 3.13: a) Average detector response to muons at the far end of cells in the Far Detector, satisfying the design requirement of 4.7 PE/cm. Figure from [77]. b) Efficiency for registering a hit on a muon track as a function of distance from the APD. Efficiency is greater than 90% for the length of cells, satisfying the design requirement. Figure from [78].

only the length is changed. The detector is 15.9 meters long, divided into a 12.8 meter active region followed by a 3.1 meter muon catcher at the downstream end as shown in Figures 3.14 and 3.15. The Near Detector totals 290 tons of which 130 tons is scintillator, 78 is steel in the muon catcher and the remainder is the PVC modules [79].

The active region consists of 192 planes, each 4.1×4.1 meters with three modules (96 cells). The active region electronics is instrumented in three 64 plane diblocks. Each diblock has two DCMs for the vertical planes and two for the horizontal planes. One DCM in each view is fully occupied with 64 FEBs and the other is half occupied with 32.




Figure 3.14: View of the upstream, 3×3 module region of the Near Detector during construction.

The muon catcher region contains 22 planes and intersperses steel plates with the scintillation planes in order to stop muons and improve containment. There are ten planes of




Figure 3.15: A) 3×3 module active region of the Near Detector. B) Muon catcher region with steel plates alternating with scintillation planes, height is two-thirds that of the active region. C) Electronics rack alcove. D) Catwalks. E) Movable access platform [80].

steel, each approximately 10 centimeters thick that are recycled from the NOvA prototype detector. The prototype had different dimensions and so the steel plates are 4.1×2.7 meters, or two-thirds the height of the active region. The horizontal modules have the same dimensions as the active region with planes consisting of two modules totaling 64 cells. The vertical modules are reduced to 2.7 meters in length with three modules per plane totaling 96 cells. The muon catcher region is a repeated sequence of horizontal plane, steel, vertical plane. The muon catcher has one DCM for the vertical modules and one for the horizontal.

3.5 Energy Calibration

The energy calibration of the NOvA detectors is divided into two phases: a relative calibration that corrects for attenuation and aging of the detector [81], and an absolute calibration that converts an energy deposition recorded by the APD into physical units of GeV [82]. Both phases of the calibration are done using a selection of cosmic ray muons. Muons are used because they can provide a source of uniform energy deposition across the detector for the relative calibration and with stopping muons the Bethe-Bloch formula can be applied to precisely calculate the energy deposited in a cell. For the calibration cell hits are chosen that also have energy deposition in the neighbor cells on both sides in the same plane. This criteria allows for more precise calculation of the path length and thus dE/dx . Statistics are accumulated for each cell in a calibration period and an attenuation curve is fit, shown in Figure 3.16, such that an energy deposition at any location in any cell can be expressed in a consistent metric of calibrated photo-electrons. Aging effects are corrected for by monitoring the mean response in a cell over time and correcting back to the calibration period. After the attenuation corrections have been applied stopping muons are selected to perform the absolute energy calibration. Cell hits are selected that are a fixed distance (between 100 and 200 cm) away from the end of a muon track where the dE/dx vs x curve

is relatively flat, shown in Figure 3.17, and the average charge per centimeter is calculated for that region. This measurement is done in data and Monte Carlo that has been tuned to match the data in units of calibrated photo-electrons. The Monte Carlo truth dE/dx information is then used to establish a scale factor connecting calibrated photo-electrons to absolute units of GeV.




Figure 3.16: A typical attenuation curve for a cell in the Far Detector where the readout end is at +750 cm and the far end of the cell is at -750 cm. Due to the looped fiber in each cell, the data is fit with a double exponential shown in red for the short and long light paths to the APD [81].

3.6 Monte Carlo Simulation

The NOvA Monte Carlo simulation begins with the production of neutrinos in the NuMI beamline. The *FLUGG* package [83] combines *FLUKA* [84] to simulate the interaction of the protons with the target and subsequent interactions of the secondaries with *GEANT4* [85] to handle the geometry description of the target and decay pipe. The output is a flux file consisting of neutrinos (flavor, energy, and momentum) at the point of creation from the parent. Flux files are created for both the forward horn current (neutrino) and reverse horn current (anti-neutrino) beam configurations. For the Far Detector oscillation weights were not applied to the files and three permutations of flux files were produced: standard flux




Figure 3.17: Energy loss with respect to distance from track end of stopping muons identified using their Michel electrons. This plot was made using cosmic data from the Far Detector. Super-imposed is the mean energy loss per cm as a function of distance from the track end. The flat region between 100 and 200 centimeters from the end of the track is used to fix the absolute energy scale [82].

consisting predominantly of ν_μ , fluxswap files where ν_e has swapped with ν_μ to enrich the appearance signal, and similar fluxswap files where ν_τ was swapped for ν_μ to explore the higher energy background. After full simulation any set of desired oscillation parameters can be applied during an analysis to produce an appropriate final weighting.

The *GENIE* package [86] determines if a neutrino has interacted with the detector and applies cross section information to output four-vectors for each particle produced in the interaction. *GEANT4* then handles propagating the products through the detector geometry. NOvA uses internal code in the *PhotonTransport* and *ReadoutSim* packages to propagate photons through the wave-length shifting fibers into the APD and produce digitized wave forms. The end product is a *ROOT* output file in the same format as data files with additional truth information ready to be used in downstream reconstruction.

Cosmic ray background is simulated in separate files (the ability to overlay in neutrino files does exist) using the *CRY* package [87] to produce the cosmic flux at the detector and then *GEANT4* takes over for particle tracking and readout simulation. In the Far Detector the neutrino flux is low enough that each event record contains one neutrino interaction with

the number of POT required to produce that interaction varying. In the Near Detector the neutrino flux is high enough that a fixed number of POT is used in each record, resulting in a varying number of neutrino interactions. Simulation files with neutrino interactions in the rock are overlaid with the fiducial Near Detector files.

CHAPTER 4

Timing System and Calibration

The NOvA experiment consists of two detectors separated by 810 km that must be precisely synchronized in time both within the detector and externally with the neutrino beam clock. This synchronization is critical for the experiment to be able to temporally separate interactions within a detector and to correlate candidate events with pulses from the neutrino beam. This chapter discusses the technical details of the design and performance of the NOvA timing system including the system layout, time synchronization, beam triggering, timing resolution and the calibration technique developed to measure delays between electronics regions of the detector. A list of frequent terms is in Table 4.1. The reader can skip this chapter without loss of continuity.

Acronym	Definition
MTDU	Master Timing Distribution Unit
STDU	Slave Timing Distribution Unit
GPS	Global Positioning System
FPGA	Field-Programmable Gate Array
ARM	Advanced RISC Machine
APD	Avalanche Photodiode
FEB	Front-end Board
DCM	Data Concentrator Module
LVDS	Low Voltage Differential Signal
DCS	Duel Correlated Sampling
ASIC	Application Specific Integration Circuit
BSYNC	Beam synchronous clock
TCLK	Tevatron Clock
ADC	Analog to Digital Converter
TDC	Time to digital Converter

Table 4.1: List of acronyms frequently used in this chapter.

4.1 Timing System Requirements

The primary physics goals of the NOvA experiment depend on recording neutrino interactions from the NuMI beam in both the Near and Far Detector. The beam is pulsed and delivers a 10 μs spill of neutrinos every 1.3 seconds. The beam spill frequency varies depending on accelerator operating conditions and which set of experiments are running. For this reason it is not possible to predict the arrival of a spill far in advance. Since the Far Detector is located on the surface with a 120 kHz background rate of cosmic rays and an expected neutrino interaction rate of 1-2 contained events per day, an activity based trigger is not capable of selecting the neutrino signal from the background. In order to collect each neutrino beam spill without being able to reliably trigger on detector activity or predict the trigger window in advance, the experiment uses a buffer farm which stores 100% of the data in memory for up to minutes before a trigger decision is made to read out data of interest into an event record. When a neutrino spill is generated at Fermilab it is time stamped by the NOvA clock and a network packet is sent to both detectors to read out data from the corresponding time window in the buffer. This trigger method requires that both detectors be precisely synchronized to an absolute wall clock that also records the spill triggers. The trigger decision is sent to both detectors, correcting for time of flight, to read out 550 μs windows centered on the beam spill. The enlarged readout window allows for side-band analysis of the background. To reconstruct and spatially separate events across the extent of the Far Detector and reduce event pileup in the Near Detector, an additional requirement was that all readout channels on the detector be in sync with the global clock within 10 ns.

4.2 Timing System Design

A diagram of the Far Detector timing layout is shown in Figure 4.1. NOvA uses a distributed system to relay timing commands to all parts of the detector [88, 89]. Both detectors have

two redundant timing chains, each controlled by an MTDU connected to an external GPS antenna. The MTDU functions to keep time for the experiment by driving time increment counters in each branch of the timing system as well as issuing commands to synchronize or reset the system among others. The MTDU interfaces with the GPS receiver through an ARM micro processor and Altera FPGA to access the data link and reference clock lines as well as provide limited control of the receiver. A Power PC single board computer is also accessible to provide a Linux platform to run software to decode accelerator triggers and access the FPGA. At the Far Detector there is one STDU for each diblock of detector connected in a daisy chain of copper cables with four LVDS lines to carry master clock, command channel, sync and sync return which carries an echo back of the sync command that is used for calibration. STDUs only have an ARM processor and do not carry a Power PC. Each STDU fans out the commands in two branches, one for the six DCMs on top of the diblock and the other for the six DCMs on the side. Each DCM chain is terminated at the end with a loop-back connector for calibration purposes. The DCM fans out to 64 FEBS. Here the sync return line is exchanged for the data link.




Figure 4.1: Schematic of the distributive timing system deployed at the Far Detector. Figure from [88].

The Near Detector uses the same timing technology and hardware as the Far Detector but the layout is modified due to the smaller size of the detector. Each diblock is its own

timing branch but the top and side views are not separated as in the Far Detector since there are only two DCMs per view, except in the Muon Catcher which has only one DCM per view. The timing link flows from the STDU to the lower DCM on the side of a diblock, then up the detector and across the top. This difference is an important distinction when calibrating the timing delays. A diagram of the Near Detector timing is in Figure 4.2.

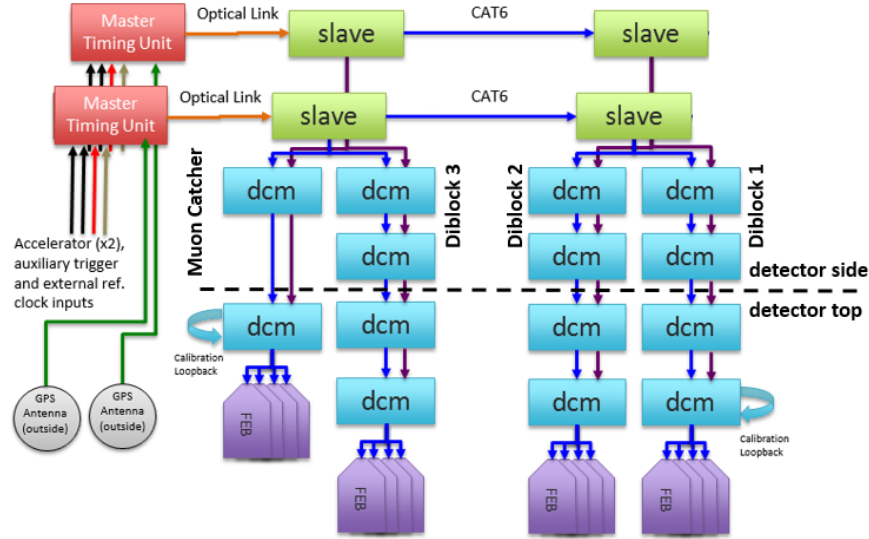


Figure 4.2: Schematic of the distributive timing system deployed at the Near Detector.

4.3 NOvA Time Specification

To establish a universal time along every step of the timing system, each TDU, DCM and FEB has a time stamp register driven by a 64 MHz clock that is derived in a phase lock loop from the 10 MHz oscillator on the GPS receiver. On each successive clock cycle the timing registers are incremented. The “NOvA Epoch” is defined as the number of 64 MHz clock ticks beginning at 00:00:00 January 1, 2010 GMT. The time stamp is encoded in a 56 bit register with the lower 32 bits containing the full 64 MHz resolution of the last 67.1 seconds, and the upper 24 bits providing lower precision, for a validity of 35.7 years. The FEBs derive a 16 MHz clock from the master. On the Far Detector the ASIC uses 8:1 multiplexing so each of the 32 readout channels on an APD is sampled at 2 MHz. At the

Near Detector, where the close proximity to the beam source creates a pileup of interactions within the 10 μ s window, the ASIC is modified to use 2:1 multiplexing and sample each channel at 8 MHz, producing an improved timing resolution discussed in Section 4.6. When the timing system performs a delay calibration a 128 MHz time stamp is used so that the detector is synchronized to a higher precision than the electronics sampling rate. During operations the readout systems use only the value in the time stamp counter to assign a time to the data packets. The system is never required to start or stop readout gates, which allows for periodic synchronization of the entire detector to keep all components within one clock cycle. If the satellite lock is lost, the 10 MHz oscillator remains stable to 2 parts per billion per day and there are monitoring tools to detect clock drift.

4.4 Timing Synchronization

To precisely synchronize the time stamp counters to the NOvA time, the timing signal propagation delays between each component of the system must be calibrated. The delay calibration is initiated by setting the “learn enable” bit in the control register of the MTDU, which starts the calibration in each STDU and DCM. The delay offset value, which is arbitrary but must be larger than the maximum travel time from the MTDU to the farthest element in the system, is loaded into each unit. Upon receipt of the next synchronization signal sent from the MTDU, every element of the chain clears its time register and initiates a counter. The counter is stopped when the unit receives a return sync signal that echoes back from the end of the timing branch. Each STDU then loads one half of the time-of-flight (TOF) value into its delay register, which corresponds to the time from that unit to the end of the timing chain. Each slave keeps an independent counter for the delay value down each DCM branch. The delay offset value is added to the delay calculated for the slave backbone and then one half the TOF for the DCM branch is subtracted to compute

the value loaded into the STDU delay register for each DCM branch. The MTDU loads its delay value, which is the total delay down the STDU backbone plus the delay offset value into the “early sync” register. This early sync value is the time prior to a GPS 1 second boundary that a sync signal is issued such that it will arrive at the hardware component of the system farthest away from the MTDU at the 1 second boundary. It is important to note that it is not possible to calibrate the cable delay from a DCM to each individual FEB so strict requirements were placed on all FEB cables being the same length with the delay set in the hardware. An example of how timing delays would be calculated and loaded is shown for a toy detector in Figure 4.3. This delay calibration procedure can be performed periodically to monitor stability and seasonal temperature variations. Figure 4.4 illustrates the stability of the calibration within one 128 MHz clock tick on the Far Detector. The timing delays can also be monitored offline using cosmic muons as described in Section 4.7.




Figure 4.3: Sample calculation of the MTDU, STDU and DCM delay values that would be loaded into the firmware registers of a toy system. Figure adapted from [90].




Figure 4.4: a) Delay measured from a DCM to the end of the timing branch in the Far Detector. DCMs 1-6 are in a branch on top of the detector with 6 being closest to the STDU backbone. DCMs 7-12 are located in a branch on the side, with 7 being closest to the backbone. b) Delay calibrated from a timing unit to the end of the backbone in the Far Detector. The MTDU is located off of the detector hall and has a longer delay, each STDU is uniformly spaced 93.75 ns apart.

The timing system uses the scheme “At the tone the time will be...” to synchronize the detector. When a time synchronization is requested the master timing unit looks at the current time and determines how close it is to the next 1 second GPS boundary. It then uses the delay loaded in the “early sync” register to calculate the next 1 second boundary sufficiently far in the future to complete transmission of the sync to all regions of the detector. This new NO ν A time is sent out and pre-loaded in to the registers of each system component. The sync pulse is sent prior to the upcoming 1 second boundary calculated so the pulse reaches all elements before the designated time. When a sync is received by an electronics component (TDU, DCM, FEB) it is placed in a delay loop buffered with the calibrated value. With proper calibration the entire detector will exit the buffer loop and begin counting from the new NO ν A time simultaneously. After the sync is completed the time stamp register in each device runs free, driven by the 10 MHz reference clock. In the DAQ setup used to take data between 02/06/2014 and 05/15/2015, hereafter defined as the first analysis period that is the subject of this thesis, the TDU delays were set within the

hardware but not the DCMs since these units were repeatedly turned off during construction and no non-volatile memory was available. Instead the DCM delays were monitored and stored in an offline database for use in event reconstruction. For future NOvA running a system has been developed to retrieve and load DCM delays at boot time so the delay correction can be made online when the data is digitized on the FEB.

4.5 NuMI Beam Triggering System

At Fermilab the MTDUs are connected to inputs from the accelerator controls network to decode and time stamp signals. These signals provide the neutrino beam spill event times in addition to reference pulses for diagnostics. The MTDU deterministically decodes and time stamps these signals in the NOvA time with the full 64 MHz resolution. This resolution is necessary to accurately identify the 10 μ s NuMI neutrino beam spills to a higher precision than the system used to log data from the accelerator system.

The TDU takes inputs from both the Fermilab beam-synchronous (BSYNC) clock, as well as the Tevatron Clock (TCLK) [91]. The Altera FPGA on the TDU decodes the events and filters out selected signals to a buffer that can be accessed from the ARM microprocessor or an attached PowerPC single-board computer. Figure 4.5 illustrates the model used to transmit spills to the detector. A spill server application runs on the PowerPC on the TDU, accesses the event queue, and publishes spills over XML/RPC to the Near Detector computing cluster at Fermilab. On the cluster a spill repeater application broadcasts messages over the internet to the Far Detector computing cluster and also to spill receivers on the Fermilab cluster. The Far Detector cluster has a second repeater application that broadcasts to receivers. A spill receiver communicates with the global trigger application when a run is in progress to trigger a readout of data from the buffer farm. Both the near and far detectors continuously write data to the buffer farm, where data is held for

20 seconds allowing time for spill messages to be received and converted to triggers. The backbone of the spill server and spill repeater remains up at all times broadcasting to up to 10 receivers simultaneously. Spill messages are broadcast asynchronously to each destination to reduce pileup if one message is delayed by a network glitch. Multiple attempts are made to send each message to minimize the chance of network traffic causing a spill to be dropped. In tests, spills were lost at a rate of less than 0.01% at the Near Detector and 0.2% at the Far Detector. More than 99% of messages arrived at the Far Detector trigger within two seconds with remaining messages coming in under ten seconds except in cases of major network interruptions. This delay is well within the limits of the buffer system which has recently been expanded to be capable of storing data for 20 minutes while waiting for a trigger.

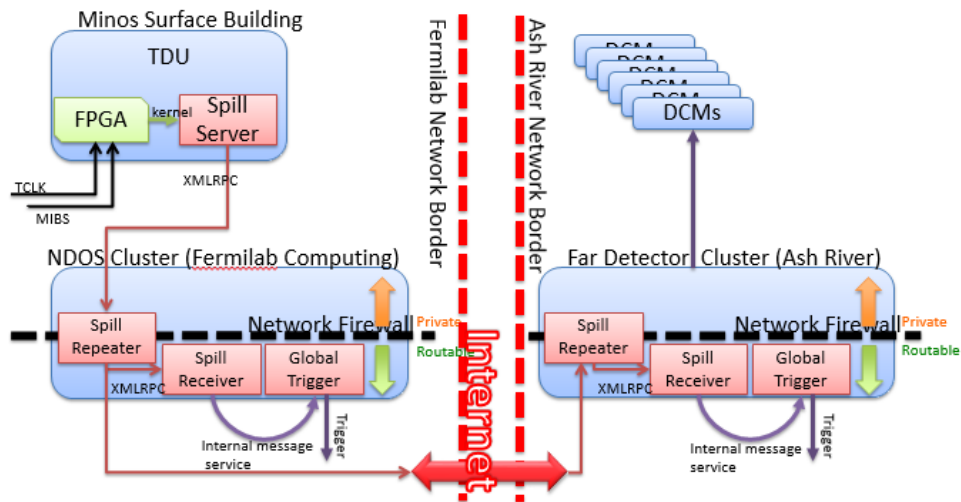


Figure 4.5: Schematic of the spill server system designed to relay accelerator event time stamps from the MTUDUs at Fermilab to applications running on the Near and Far detectors triggering data readout. Figure from [89].

A timing calibration reference Unit (TCR) is located at each detector site to monitor the MTUDUs for any glitches in synchronization or clock drift. The TCR consists of an independent GPS antenna and receiver that produces several highly stable one pulse-per-second trigger outputs. These triggers are decoded by a standalone application of the spill

server running on each MTDU which detects clock errors and send messages to the NOvA computing interface that controls data acquisition to issue a sync or abort the run.

4.6 Timing Resolution

The ASIC on each FEB shapes the pulse signal from the APD with a 460 ns rise-time and 7000 ns fall-time at the Far Detector where the signal is sampled every 500 ns. For the Near Detector, which samples four times as frequently, the rise-time is 140 ns with a fall time of 4500 ns. The FEBs at both detectors use a Duel Correlated Sampling (DCS) algorithm for determining which information to save in the readout. The DCS algorithm looks at the ADC difference between each sample, s_i , and the sample three ticks before the current sample, s_{i-3} , requiring the difference to be above the threshold value for that channel. If this criterion is satisfied the hit is recorded with the time being the TDC value for sample s_i and the charge set as the ADC difference between s_i and s_{i-3} . This method is known as single-point readout since the recorded hit is represented by a single time and charge and was used in the Far Detector between February and August of 2014 during part of the analysis period. The best case timing resolution for this method is

$$\sigma_{tsingle} = \frac{t_{sample}}{\sqrt{12}} = \begin{bmatrix} 144 \text{ ns (Far Detector)} \\ 36 \text{ ns (Near Detector)} \end{bmatrix}, \quad (4.1)$$

where t_{sample} is the time between samples, 500 or 125 ns, and $\sigma_{tsingle}$ is the resultant resolution for that single time. This method can lead to ambiguity where two different pulse-shape curves are determined to have the same charge and time because of the discrete sampling as shown in Figure 4.6. For the Near Detector this is not sufficient to be able to temporally separate neutrino interactions. Instead multi-point readout is used where in addition to the trigger sample the three preceding baseline subtracted ADC values are also recorded. This approach allows for an offline fit of the pulse shape from this set of four

samples of charge and time to be done in order to more precisely determine the ADC and TDC value [92]. The Far Detector adopted multi-point readout mode in August 2014.




Figure 4.6: An illustration of the DCS algorithm used to trigger readout on the FEB. The black and blue curves represent two different pulses and the markers s_{i-3} , s_{i-2} , s_{i-1} , and s_i are the digitization samples. In single-point mode these two pulses would have the same change and time recorded. In multi-point mode saving all four samples allows a more accurate fit for the pulse start time and peak ADC [93].

The timing resolution of both detectors was determined empirically from the data. Events were spatially and temporally separated with an algorithm into groups of hits that are known as slices, discussed in Section 5.3. The slices were then fit with a simple straight-line fitter originally designed for the first stage in rejecting cosmic muon backgrounds. A series of quality cuts detailed in Section 4.7.1 were used to select long through-going cosmic muons or muons from neutrino interaction in the rock outside the detector in the case of the Near Detector. The times of hits are corrected time of flight along the track and distance to readout in the cell as described in Section 4.7.2. After these corrections the time difference between all pairs of hits within a single DCM on the cosmic track is calculated. A two dimensional histogram is filled with these time differences and the number of photoelectrons (PE), shown in Figure 4.7. An entry is only made for a pair if both hits fall in the same 25 PE bin. Within each bin the top and bottom 1% are truncated to remove outliers and then the RMS is calculated which represents the timing resolution for that charge. The

resolution was fit with a function of the form:

$$\sigma_t = \frac{p_0}{p_1 + n_{pe}^{p_2}} + p_3 \quad (4.2)$$

which describes the data from the Far Detector (Figure 4.8) and Near Detector (Figure 4.9). Here n_{pe} is the number of photo-electrons, σ_t is the timing resolution, and p_0 , p_1 , p_2 , and p_3 are fit parameters. Further improvements to the timing resolution are limited by the sampling frequency, phasing of readout channels, the uncertainty on the detector synchronization of 7.8125 ns (one clock tick at 128 MHz), and noise. This timing resolution is used to optimize time clustering and is also necessary to separate upward and downward going muons in the Far Detector.






Figure 4.7: The time difference between pairs of cell hits in the same DCM on reconstructed muon tracks as a function of the number of photo-electrons in each hit shown for Far Detector multi-point data. The RMS of each column is the timing resolution.

4.7 Offline Timing Calibration

In order to have a crosscheck of the online TDU delay calibration and monitor the stability of the timing system, DCM delays are computed offline using cosmic rays in a technique modeled from the MINOS experiment [94]. Given the location of the Far Detector on




(a) single-point




(b) multi-point

Figure 4.8: Timing resolution determined from Far Detector data.



(a) single-point



(b) multi-point

Figure 4.9: Timing resolution determined from Near Detector data.

the surface, cosmic muons are so numerous that strict quality cuts can be placed without suffering from a lack of statistics. Since the Near Detector is located 100 meters underground the cosmic background for this calibration is reduced and a higher statistics calibration sample is obtained by using muons produced in neutrino interactions in the rock outside the detector. After a track selection is applied the relative time differences (offsets) between hits in different DCMs is calculated. Then a matrix based on these relative differences is inverted to solve for the absolute timing offsets between each DCM in the detector and a fixed reference DCM. If the synchronization described previously is performed properly all the absolute offsets should come out to zero. However, as of the time of writing this document the synchronization only accounts for the cable delay between STDUs and not for the delays between DCMs on a given STDU branch.

4.7.1 Track Selection

In order to select suitable hits on cosmic tracks for the timing calibration the following cuts are used:

1. The fraction of cell hits kept in the straight-line cosmic track fit is required to be at least 70% of the total cells in a reconstructed space-time slice (defined in Section 5.3. This cut removes poorly fit tracks when the slice is either misreconstructed or contains a neutron shower or neutrino event instead of a muon.
2. Tracks are required to pass through at least 10 planes to remove highly vertical tracks.
3. Individual hits on a track are only kept if the path length through that cell is less than 10 cm. This cut mainly removes hits in vertical cells in the X-Z planes where steeper muons traverse larger sections of these cells and produce photons across a larger range of distances within a cell that can distort the reconstructed hit time, creating an asymmetry between views.

4. Tracks are required to be anti-fiducial, having a start and end point within 20 cm of a wall in X and Y, or 10 cm in Z. This criterion is used to select through-going tracks so that the particle speed can be approximated by the speed of light along the entire path length without worry of endpoint effects.
5. The difference between the furthest upstream (south) track hit in the X-Z and Y-Z views is required to be no more than 3 planes. The same is required of the furthest downstream hit (north). This cut removes another class of track misreconstruction.
6. The asymmetry $\frac{nhit_x - nhit_y}{nhit_x + nhit_y}$ is required to be less than 0.5, where $nhit_x$ and $nhit_y$ are the number of hits in the vertical and horizontal readout views in the slice. This cut ensures that tracks sample both views and maintains balance in the combinations of DCMs used on a track.
7. Tracks are required to have at least 175 cells in the Far Detector (50 in the Near Detector) to provide an adequate profile of the track timing.
8. Only DCMs with at least 25 hits (10 in the Near Detector) are used in the calibration. A track must have at least two DCMs that meet this criteria. The hit requirement is reduced in the Near Detector since the steel in the muon catcher makes larger numbers of hits unlikely.

After this selection the Far Detector produces 200,000 tracks per hour. This number is sufficient to calibrate the detector and allows the timing system to be monitored on a fine level if the need arises or there is reason to suspect a certain period of running. For the Near Detector, which is ~ 64 times smaller, 24 hours is necessary to accumulate a sufficient sample of ~ 3500 events.

4.7.2 Calibration Procedure

After the previously described cuts are applied, the following procedure is used to compute the absolute timing offsets between DCMs on the detector:

1. The time of each cell hit in a DCM (t_i) is taken from fitting the four readout samples to a curve as shown in Figure 4.6. This time is corrected for the time of flight along the muon track path length (d_p) assuming the speed of light. The time is also adjusted based on the distance of the hit to the readout APD (d_r) assuming a fiber speed of 15.3 cm/ns resulting in

$$t'_i = t_i - \frac{d_p}{c} - \frac{d_r}{15.3} \quad (4.3)$$

where all times are in nanoseconds and distances are in centimeters. This correction does not factor in time for the scintillation light to be absorbed or any variations in fiber speed. For a discussion on fiber speed and additional corrective factors that may be necessary for a precision timing fit see Section 4.7.4. The intent is that the adjusted times t'_i represent the simultaneous readout time of each hit on a track within one DCM.

2. The weighted average and standard deviation of the time for each DCM on a track is computed from the corrected hit times by

$$\overline{DCM}_i = \frac{\sum_i t'_i \sigma_i^{-2}}{\sum_i \sigma_i^{-2}}, \quad (4.4)$$

where σ_i is taken from equation 4.2.

3. The relative timing offsets between pairs of DCMs on a track are computed as

$$\Delta_{ij} = \overline{DCM}_i - \overline{DCM}_j, \quad (4.5)$$

where by convention j always corresponds to a higher numbered DCM than index i . When done over a large collection of tracks this builds up an ensemble measurement of the relative DCM offsets and errors σ_{ij} . An example is shown in Figure 4.10. A matrix of relative offsets is then compiled from all DCM pairs that were used at least

10 times. This offset matrix is shown for the Far Detector in Figure 4.11 and for the Near Detector in Figure 4.12.




Figure 4.10: A representative example of the measured relative time offset between a pair of DCMs and the expected value based on the detector operating conditions at the time of writing this document.

4. The matrix of relative offsets can then be solved for the absolute timing offsets between DCMs. One DCM will be chosen as a reference with a fixed time and all other DCMs will be solved for the time difference between hits in it and the reference DCM. In a perfectly calibrated detector these absolute offsets would all be 0. For current operating conditions the solution will show the underlying DCM cable delays. First, a χ^2 can be written of the form

$$\chi^2 = \sum_{i=1}^n \sum_{j=1}^n \sigma_{ij}^{-2} (\Delta_{ij} - (a_i - a_j))^2 \Theta_{ij}, \quad (4.6)$$

where a_1, a_2, \dots, a_n are the absolute offsets of the n DCMs and $\Theta_{ij} = 1$ for $j > i$ and




Figure 4.11: The matrix of relative timing offsets between DCMs at the Far Detector. Dotted lines indicate the diblock boundaries and blank spaces indicate DCM pairs with insufficient statistics. The structure seen in the plot is due to the repeated pattern of cable delays between DCMs in a diblock, discussed in Figure 4.1.




Figure 4.12: The matrix of relative timing offsets for the Near Detector. Dotted lines indicate the diblock boundaries. The structure in the plot is different from that of the Far Detector shown in Figure 4.11 because the DCMs are connected in a different pattern, most notably the detector views are not on separate timing branches (Figure 4.2).

is 0 otherwise. Taking the derivative

$$\frac{\partial \chi^2}{\partial a_i} = 0 \quad (4.7)$$

produces a system of linear equations that can be solved, each of the form

$$a_i \sum_{j=1, j \neq i}^n \sigma_{ij}^{-2} - \sum_{j=1, j \neq i}^n \sigma_{ij}^{-2} a_j = \sum_{j=1}^{n, j \neq i} \sigma_{ij}^{-2} \Delta_{ij}. \quad (4.8)$$

It is then necessary to fix one DCM to a known offset. A discussion on choices for the reference DCM can be found in Section 4.7.3. In matrix form the system of equations is:

$$\begin{bmatrix} 1 & 0 & 0 & \dots \\ -\sigma_{12}^{-2} & \sum_{j=1}^{n, j \neq 2} \sigma_{2j}^{-2} & -\sigma_{32}^{-2} & \dots \\ -\sigma_{13}^{-2} & -\sigma_{23}^{-2} & \sum_{j=1}^{n, j \neq 3} \sigma_{3j}^{-2} & \dots \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & & & \\ \cdot & & & \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \cdot \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} 0 \\ \sum_{j=1}^{n, j \neq 2} \sigma_{2j}^{-2} \Delta_{2j} \\ \sum_{j=1}^{n, j \neq 3} \sigma_{3j}^{-2} \Delta_{3j} \\ \cdot \\ \cdot \\ \cdot \end{bmatrix} \quad (4.9)$$

which can be solved for a_i , the absolute timing delays of each DCM, by inverting the matrix. The solution is highly correlated, but the errors can be overestimated from the diagonal elements by taking the second derivative of the χ^2 equation and producing errors for the i^{th} DCM of the form

$$\sigma_i = \sqrt{\frac{1}{\sum_{j=1, j \neq i}^n \sigma_{ij}^{-2}}}. \quad (4.10)$$

4.7.3 Timing Calibration Results

The first timing calibration of the Far Detector used multi-point data taken between September and November 2014. During this time parts of the detector were being retrofitted which varied the active detector mass. A few periods of stable running, listed in Table 4.2 were selected during this period to monitor the stability of the calibration over time and test the result over as much detector volume as possible.

diblock configuration	runs	date
5 - 14	17212 - 17224	September 11 - 12, 2014
6 - 14	17562 - 17588	October 4 -5, 2014
3 - 14	17808 - 17826	October 18 - 20, 2014
2 - 14	17993 - 18003	November 1 - 2, 2014

Table 4.2: Far Detector run periods used for the timing calibration.

For this calibration pass DCM’s 18 and 19 were fixed with a timing offset of zero nanoseconds. These two DCMs have the shortest cable path to the MTDU for the detector volume used in the calibration with DCM 18 on top of the detector and DCM 19 on the side. The resulting absolute timing offsets are shown in Figure 4.13. It is clear that starting in diblock 7 there is a systematic drift in the delay values for the second half of the detector. Choosing a different reference DCM shifts the calibration results for the entire detector up or down, but does not change the slope. This drift was seen in all periods of data that were used for the calibration. In addition, restricting the analysis only to diblocks 7 through 14 shows the same drift, indicating the drift is not caused by a pull in relative timing offsets between DCMs separated by large distances in the front and back halves of the detector. The cause is unknown but possibly due to a change in the performance of the liquid scintillator or wavelength shifting fiber over the course of construction. This drift is not a real effect of the timing system and so it can be corrected for by adjusting the absolute offsets in diblocks 7-14 to center the errors around zero. This corrected calibration is shown in Figure 4.14 and a comparison of the timing calibration result to the expected delay values before and after




Figure 4.13: The absolute timing calibration for diblocks 2 to 14 in the Far Detector, showing a systematic drift beginning in diblock 7. The red line indicates the measured delays from the TDU and the black points are the calibration result using muons. Dotted lines indicate the diblock boundaries.

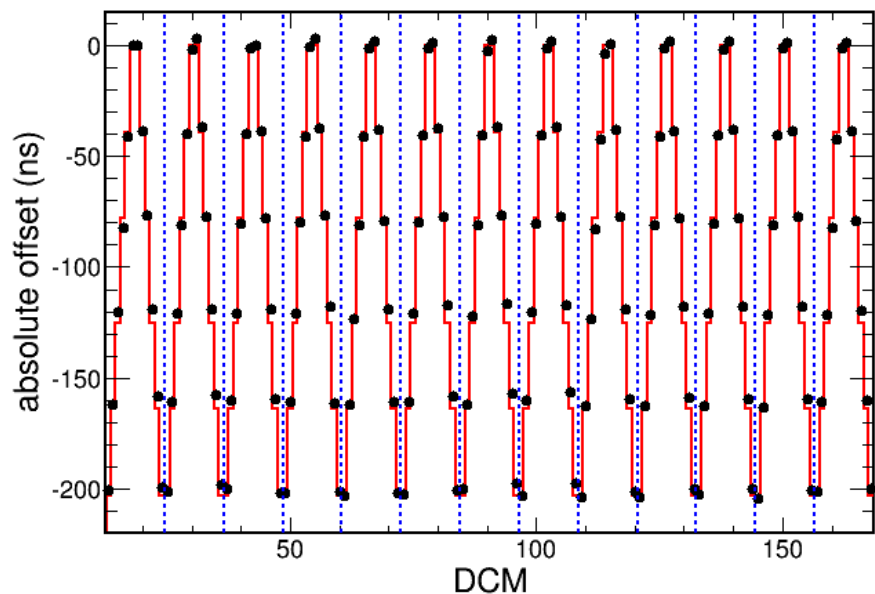


Figure 4.14: The absolute timing calibration for diblocks 2 to 14 in the Far Detector after a drift correction was applied to diblocks 7 to 14. The red line indicates the measured delays from the TDU and the black points are the calibration result using muons. Dotted lines indicate the diblock boundaries.

the drift correction is shown in Figure 4.15. This final result produces errors of less than 10 nanoseconds when comparing the muon calibration result with the timing delays calculated from the TDU. There is still a small asymmetry between the delays calculated for DCMs on top of the detector and those on the side. This asymmetry is because the vertical modules see on average slightly higher energy depositions than the horizontal modules, which leads to improved timing resolution. This effect comes from the nature of muons coming from the atmosphere tending to traverse a larger cross section of a vertical cell. The effect is mitigated to some extent by imposing a path length quality cut. A tighter cell quality cut or an additional correction could be used in the future if higher precision is required of the timing calibration. This precision is not necessary for an oscillation analysis where the current level of calibration is sufficient to temporally separate interactions, but may be required for exotic searches that involve distinguishing upward from downward going muons. The stability of the timing calibration was studied on a run-by-run basis over a span of

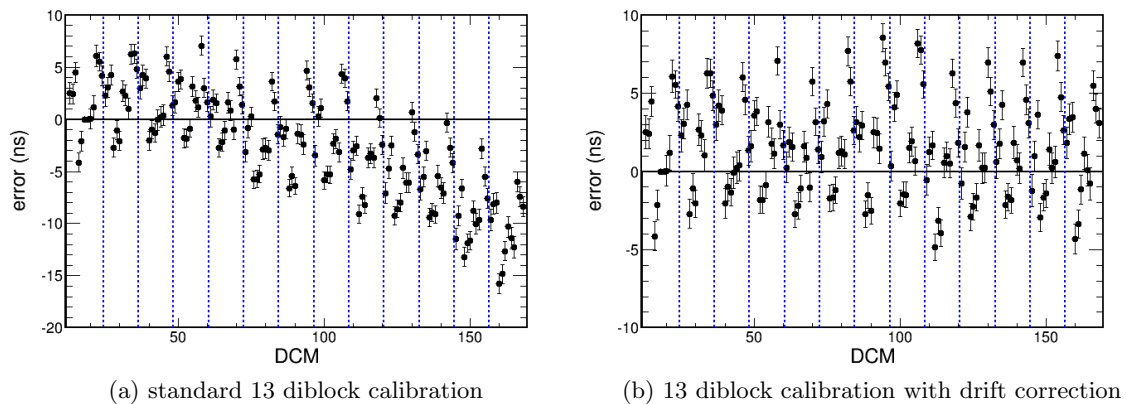


Figure 4.15: The errors in the absolute timing calibration are shown before (a) and after (b) a drift correction was applied to diblocks 7-14.

several weeks and found to be stable within a few nanoseconds as illustrated in Figure 4.16. This stability demonstrates that the timing delay calibration constants do not need to be calculated on a fine level and only need to be done every few months or when TDUs or

DCMs are replaced. The result is that tight cuts can be used if needed since statistics can be accumulated over the course of several months.


The same calibration procedure was performed on the Near Detector using through-going muons induced from neutrino interactions in the rock upstream of the detector. For this case DCM 12, located in diblock 3, and DCM 14, in the muon catcher, were fixed as the reference units and represent the closest path to the TDU backbone. The timing offsets in the first two diblocks were again systematically adjusted to remove a drift. The resulting calibration produces errors of the same size as the Far Detector, shown in Figure 4.17.

4.7.4 Additional Correction Factors


The simulation suggests that an additional correction is needed to the cell times as a function of location in the cell. The origin of this correction is the fact that the wavelength shifting fiber in each cell is looped and connected to the APD at both ends. Depending on where the energy is deposited in the cell the relative proportions of light traveling in both directions and arriving at different times can distort the pulse shape and change the reconstructed time. This effect was studied in simulation by calculating the left hand side of Equation 4.11,

$$T_{FLS} + \frac{d_r}{S_{fiber}} - T_{reco} = f(d_r), \quad (4.11)$$

where T_{FLS} is the true time of the energy deposit, d_r is the distance to the readout APD, S_{fiber} is the speed of light in the fiber, and fitting it to $f(d_r)$ which is an eighth degree polynomial. An eighth degree polynomial was necessary in order to properly interpolate the features of Figure 4.18. It is clear that the result is somewhat dependent on what choice is made for the speed of light within the fiber, with 15.3 cm/ns representing a mean value given the optical properties of the fiber. It was found that adding this correction to the calibration results provided negligible improvements. It is likely the current timing



(a) DCM 80



(b) DCM 120

Figure 4.16: The timing calibration was performed on a set of runs two weeks apart to show stability in the result on a run-by-run basis. Results from two representative DCMs are shown.




Figure 4.17: a) The absolute calibration of the Near Detector after correcting for diblock drift is shown by the black points with the expected result in red. b) The errors on the calibration result. The blue dotted lines indicate diblock boundaries.

resolution and cell-to-cell variations in fiber speed overshadow this effect. It is possible to perform a measurement of the distribution of fiber speeds in the detector by writing the time of an energy deposition on the muon path, T_p as


$$T_p = T_0 + \frac{d_p}{\beta c} \quad (4.12)$$

where T_0 is the start time of the track and d_p is the distance along the track. It will be assumed that the muon is not stopping inside the detector and $\beta \approx 1$. Stopping muons were not used since the statistics were not needed and the end of a stopping muon could introduce complications from the particle slowing down or a Michel electron included on the track. Then the reconstructed time digitized from the APD, T_{APD} can be written as


$$T_{APD} = T_p + \frac{d_r}{S_{fiber}}. \quad (4.13)$$

Then equation 4.12 and 4.13 can be combined to produce

$$T_{APD} - \frac{d_p}{c} = T_0 + \frac{d_r}{S_{fiber}}. \quad (4.14)$$



(a) Far Detector



(b) Near Detector

Figure 4.18: The correction to the reconstructed times as a function of the distance of the energy deposition to the APD for the far (a) and near (b) detectors. The three curves illustrate the effect variations in fiber speed have on the correction.

If the DCM timing delay corrections are applied based on the first calibration, then each cosmic track can be fit to a straight line using equation 4.14 where the slope will be a measure of the fiber speed. After quality cuts are applied to select long straight tracks and outlying hits are removed, measuring a large track sample would produce an estimated fiber speed in each cell. This information can then be fed back into the timing calibration in an iterative process. Taking this approach may improve performance of the timing calibration in the back half of the Far Detector where there is a drifting effect. Having a measurement of the distribution of fiber speeds within the detector and the timing correction factors as a function of distance to the APD are important future developments for attempts to use timing to determine track directionality. This measurement is of interest both as a tool to reduce some of the cosmic background in the neutrino analysis and to separate downward cosmic muons from upward going muons that are of interest for an indirect dark matter search.

CHAPTER 5

Event Reconstruction

5.1 Event Topologies

In order to perform the measurement described in this thesis a method is required to identify ν_e charged-current interactions and reject backgrounds from ν_μ charged-current (CC) and neutral-current (NC) interactions. For comparison, Figure 5.1 shows events simulated with identical four-vectors for the particles. The events contain a 0.78 GeV momentum proton and then a second 1.86 GeV momentum particle (e, μ, π^0) to represent a 2.15 GeV neutrino interacting within the detector. The figure shows one detector view for each event and the cell hits are colored by the charge deposited.

In the top panel is a ν_μ CC event featuring a track-like muon that for most of its length is a minimum-ionizing particle (MIP) depositing energy near the minimum of the Bethe-Bloch curve (1.76 MeV/cm in the NOvA geometry). The desired ν_e CC signal events are shown in the middle panel. The electron has a distinct shower shape with the energy deposition per plane rising and falling, making separation of electron and muon events relatively straightforward. The more difficult background is neutral current interactions with a single π^0 , shown in the lower panel. The π^0 decays to two photons with a 98.8% branching ratio. These photons produce electromagnetic showers that can be difficult to distinguish from electrons. The photons travel some distance before converting into an e^-/e^+ pair which produce scintillation light, where the photon conversion distance in the

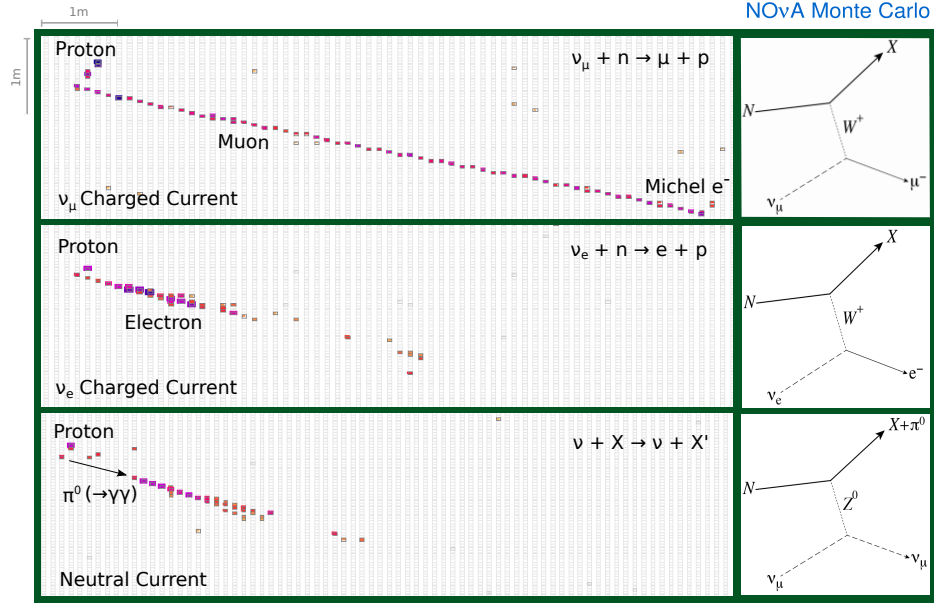


Figure 5.1: Simulated 2.15 GeV neutrino interactions showing a toy ν_μ CC interaction (top panel), ν_e CC interaction (middle panel), and NC interaction (bottom panel). Figure from [95].

NOvA detector is 53 centimeters (9 plane widths). The photons from π^0 's can be identified by their gap between the vertex and the start of the showers, which is why NOvA was designed with low-Z materials to yield a longer conversion distance. Additionally, the very start of a photon shower will have a dE/dx of approximately 2 MIP's (minimum ionizing particle) from the e^-/e^+ pair creation while the start of an electron shower will have only one MIP. In the NOvA detector a MIP deposits 1.76 MeV/cm as determined though the absolute energy calibration using stopping muons described in Section 3.5.

For the NOvA baseline of 810 km the maximum oscillation probability for ν_e appearance occurs at $E_\nu \approx 1.6$ GeV, and the off-axis NuMI spectrum peaks at 2.1 GeV. The analysis focuses on events of one to three GeV of visible energy deposited in the detector. Events below that region tend to be of low reconstruction quality and contain little signal. Higher energy events are dominated by intrinsic beam ν_e 's which represent an irreducible background to the oscillation search.

5.2 Reconstruction Philosophy

A variety of reconstruction techniques have been developed in the NOvA experiment for different purposes. For the ν_e appearance analysis it was desired to have one reconstruction chain that was successful in picking out electron showers while still performing well for the more track-like muons, protons, and pions. It was decided to take an approach which finds the global event vertex first before forming particle reconstruction objects. This approach was chosen because short tracks of only a few cells (protons, neutrons, photons) become more significant and easier to identify if they can be associated to a vertex.

The full chain of reconstruction is outlined with a toy example in Figure 5.2. Reconstruction begins by separating individual event interactions from a larger readout window into objects known in NOvA as “slices” which are intended to collect together all hits from a single neutrino interaction and serve as the foundation for all later reconstruction stages, described in Section 5.3. Next, a modified Hough transform is applied to identify prominent straight-line features in a slice (Section 5.4). Then the Hough lines are used as seeds to an algorithm to determine the global 3D vertex for the slice under the assumption that all activity in the slice has a common origin (Section 5.5). The vertex is then used as a seed to a “fuzzy k-means” algorithm that produces prongs (a collection of cell hits with a start point and direction) which contain the activity of particles in the event (Section 5.6). Finally, a variety of variables resulting from this reconstruction (dE/dx profiles, distance between vertex and prongs, length, direction, etc.) are fed into an artificial neural net to classify the degree to which the slice was a ν_e CC interaction (Section 5.7). A primary contribution of this thesis is the fuzzy k-means clustering algorithm. I also made contributions to the vertex seeding, speed optimizations, and performance evaluations of the neural network.

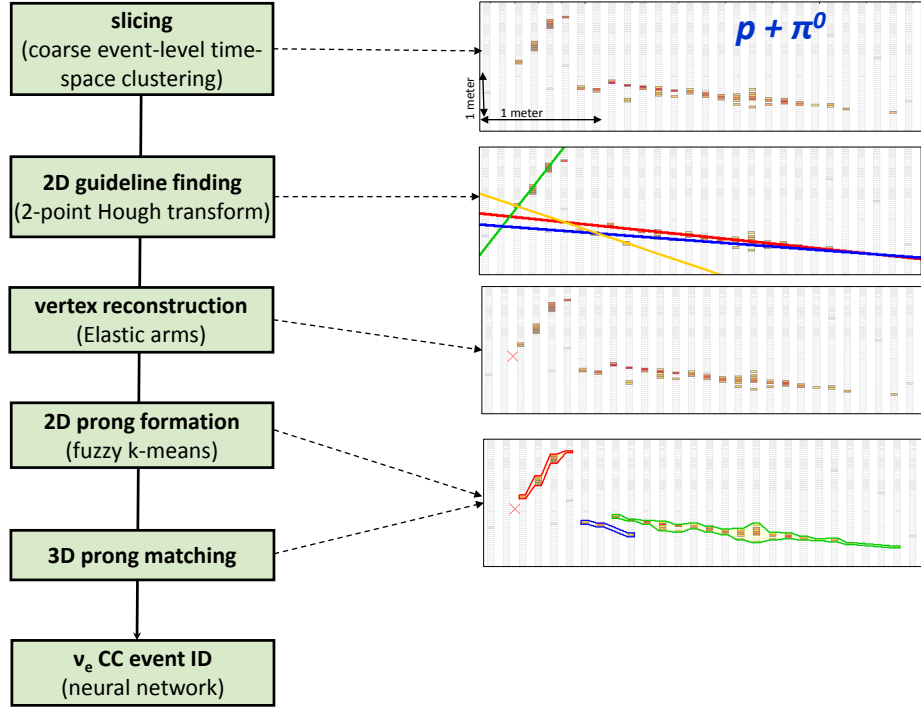


Figure 5.2: Flowchart of the reconstruction algorithms used to identify ν_e CC interactions.

5.3 Interaction Separation with DBSCAN

The first stage in reconstruction requires separating physics interactions (typically neutrinos or cosmic rays) into objects called “slices” for further processing. The spread of times recorded for cells hit in an interaction can range between tens of nanoseconds up to a microsecond depending on the length and direction of the interaction and version of the readout electronics used (refer to Section 4.6 for discussion of the time digitization design and timing resolution during the different run configurations). For the Far Detector the average standard deviation of hit times in a slice is ~ 200 ns with single-point timing and ~ 60 ns with multi-point timing. The Near Detector uses faster electronics and has a slice standard deviation of ~ 10 ns. Data readout windows can be of arbitrary size defined by the trigger software, but are currently set to $550 \mu\text{s}$ roughly centered on the $10 \mu\text{s}$ neutrino beam spill. In the Near Detector the challenge is separating 3-4 neutrino interactions that occur within each beam spill. In the Far Detector neutrino interactions are much rarer but

50-70 cosmic rays are seen in a typical readout window.

The reconstruction tool utilized to form slices is a density based clustering algorithm, DBSCAN [96], that has been tuned to perform in the NOvA environment [97]. The algorithm works by computing a distance metric between pairs of cell hits in four-dimensional space such that hits which are a light-like distance apart appear near each other. The metric is

$$D_n = \left(\frac{|\Delta T| - |\Delta \vec{r}|/c}{T_{res}} \right)^2 + \left(\frac{\Delta Z}{D_{pen}} \right)^2 + \left(\frac{\Delta XY}{D_{pen}} \right)^2. \quad (5.1)$$

where T_{res} is the timing resolution for the hits summed in quadrature, D_{pen} is a distance penalty, ΔT is the time in nanoseconds between hits, and ΔZ and ΔXY are the distances in centimeters between hits in each view. For hits in the same view $|\Delta \vec{r}| = \sqrt{\Delta Z^2 + \Delta XY^2}$ while for hits in opposite views $|\Delta \vec{r}| = \Delta Z$.

The distance metric, D_N , is computed for all pairs of points. Pairs with a value smaller than a threshold distance are classified as neighbors. A hit with at least four neighbors is a “core point” and forms the seed of a cluster. A core point and all its neighbors are put into a cluster, if any of those points are also core points the cluster expands to include the neighbors of that point and so on until all connected core points and their neighbors are grouped in a cluster. The final slice object is required to have a minimum of three hits in each view which is rare for coincident noise. A typical Far Detector event is shown before and after slicing reconstruction in Figure 5.3. To evaluate performance slice purity and completeness are measured in simulations where completeness is

$$\text{Completeness} = \frac{\text{Energy from interaction deposited in slice}}{\text{Total energy from interaction deposited in detector}} \quad (5.2)$$

and purity is

$$\text{Purity} = \frac{\text{Energy from interaction deposited in slice}}{\text{Total energy in slice}}. \quad (5.3)$$

In later stages of the analysis quality cuts are placed to remove small slices so performance evaluations focused on slices with a minimum of ten hits in each view. In Far Detector cosmic simulations slicing was found to have a completeness and purity of 99.3%, shown in Figure 5.4. For the Near Detector neutrino simulations slicing had a purity of 98.5% and completeness of 94.4%, shown in Figure 5.4, which makes concerns of pileup from multiple neutrino interactions negligible.


5.4 Guidelines with Hough Transform

After slicing has been applied the next reconstruction stage is to run a Hough transform algorithm on each slice to find the prominent lines. The output of the algorithm is a set of straight lines in each readout-view such that the directions and intersections of these lines can be used to seed a vertexing algorithm.

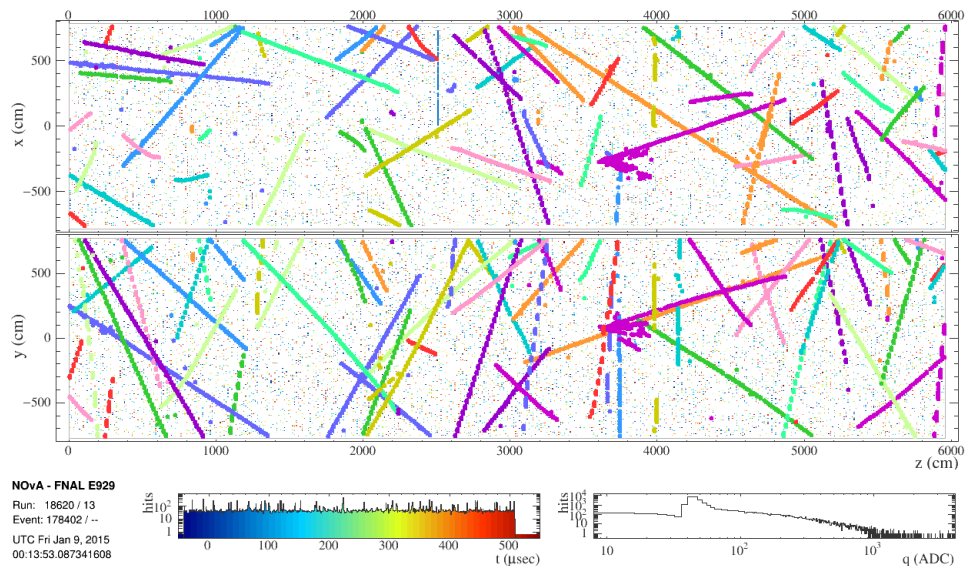
In order to be robust against noise, a modified Hough Transform is used that works on pairs of points [98]. A line segment is drawn through each pair of points and transferred into polar coordinates (ρ, θ) where ρ is the perpendicular distance from the line to the origin and θ is the angle between ρ and the x -axis.

The algorithm works in each detector view separately filling a two-dimensional Hough space in the coordinates ρ and θ with a Gaussian smeared vote for each hit pair. To reduce the number of voting pairs, a maximum distance between points is set. To prevent a tendency to make horizontal lines, points with the same xy coordinate must be more than a minimum distance apart. Peaks in the Hough space map are identified as the coordinates of a line. To separate signal from noise a threshold for peak identification is set as the average height of all ρ, θ bins in the Hough space.

To reduce the tendency to make spurious lines an iterative procedure is used. To start the highest peak in the Hough map is found and then a line is formed from the weighted




(a) before slicing




(b) after slicing


Figure 5.3: 550 μ s Far Detector readout window before (a) and after (b) slicing reconstruction. Hits are colored by time. The top panel shows the vertically oriented planes and the bottom panel the horizontal planes. Hits are colored in bold with the reconstructed slice they are a member of, un-bolded hits are not associated with a physics slice. At ~ 3600 cm in z a contained neutrino interaction is clearly seen.




(a) Far Detector



(b) Far Detector



(c) Near Detector



(d) Near Detector

Figure 5.4: The completeness and purity of reconstructed slices with at least 10 hits per view for Far Detector cosmic ray simulations (a,b) and Near Detector neutrino simulations (c,d). Figure from [97].

average ρ and θ of a 7×7 grid of bins surrounding the peak. Cell hits that fall within 6 centimeters (one cell depth) of the line are removed with the exception of the most upstream and downstream hit, which could be shared with other tracks coming from a common vertex. After hit removal, the Hough map is recalculated and a new peak is found. This process repeats until no new lines remain or the maximum number of lines (currently 10) is reached. An example of this iterative process is shown in Figure 5.5. Further information about the adaptation of the Hough algorithm for the NOvA experiment can be found in [99].

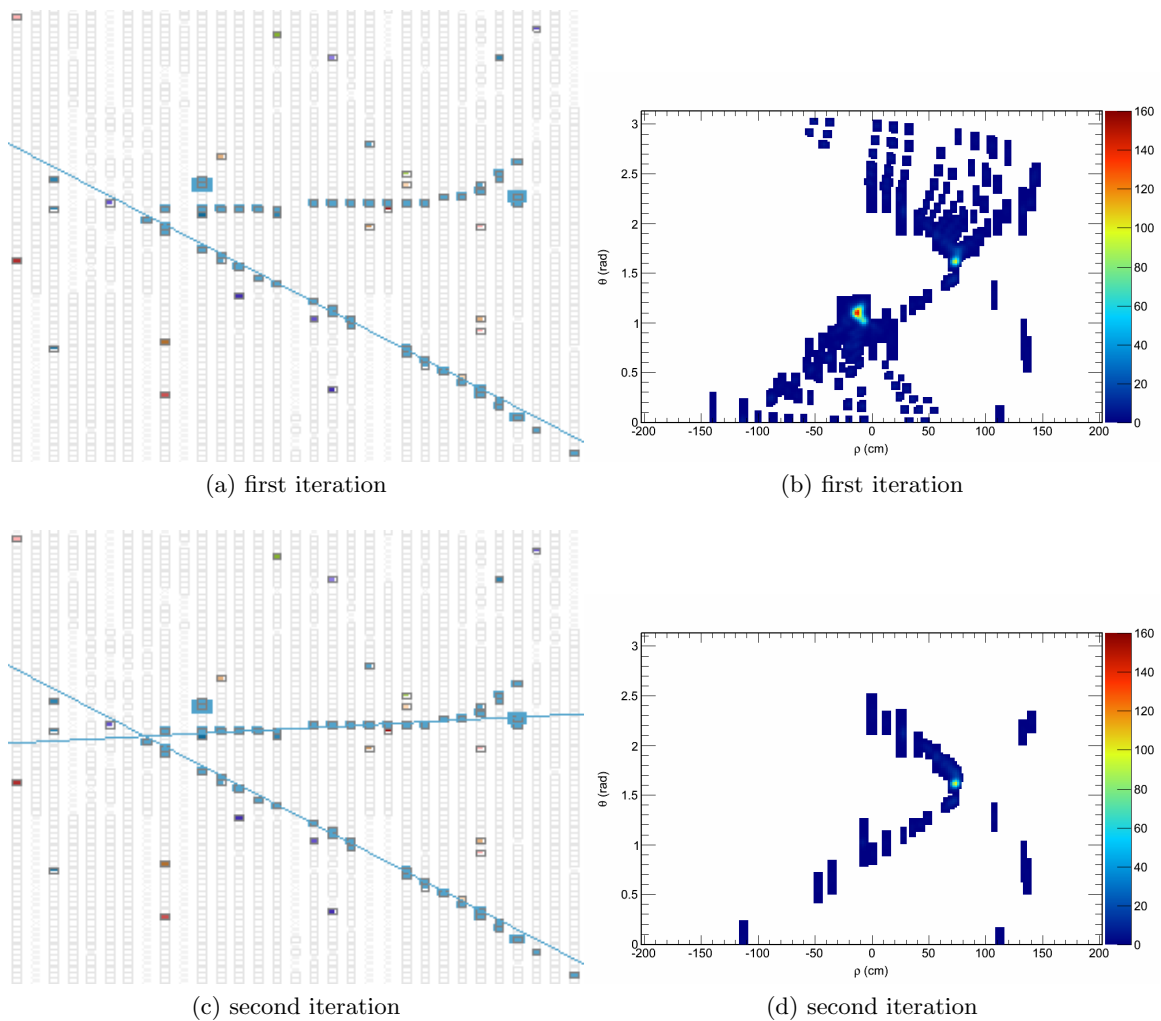


Figure 5.5: The first and second iteration of the multi-Hough transform on one view of a Far Detector neutrino simulation event. The first selected line is shown in (a) and the corresponding Hough map in (b). After hits associated with this line are removed the second line (c) and corresponding map (d) are calculated. Figure from [99].

The main performance criteria for the algorithm is for the dominant Hough lines to pass and form intersections near the primary interaction point of the slice. It was found in simulated Far Detector interactions that the primary Hough line passed within an average of 6.9 (NC), 4.1 (ν_μ CC), and 2.7 (ν_e CC) centimeters of the vertex as shown in Figure 5.6. For the secondary Hough line the average distance is 9.9 (NC), 8.2 (ν_μ CC), and 8.8 (ν_e CC) centimeters.




Figure 5.6: The perpendicular distance between the first (a) and second (b) most dominant Hough lines and the true neutrino interaction vertex for simulated ν_e CC, ν_μ CC, and NC interactions in the X-sampling readout view area normalized to one. Figure from [99].

5.5 Vertex Identification with ElasticArms

The next phase in the reconstruction is to determine the global interaction vertex. The assumption is made at this stage that all visible energy in the slice is the result of one primary interaction point, which is generally true for neutrinos. For each slice a set of “elastic arms” (also known as “deformable templates” in the literature) can be defined where each arm is a straight line with polar angle θ_a and azimuthal angle ϕ_a starting at an origin (x_0, y_0, z_0) so any point a distance s along the arm can be defined in Cartesian

coordinates by

$$\begin{aligned}x(s) &= x_0 + s \sin \theta_a \cos \phi_a \\y(s) &= y_0 + s \sin \theta_a \sin \phi_a \\z(s) &= z_0 + s \cos \theta_a.\end{aligned}\tag{5.4}$$

A classical elastic arms algorithm [100–103] works to find the parameters $(x_0, y_0, z_0, \vec{\theta}, \vec{\phi})$ that best describe the event topology. In most applications the vertex is known a priori or at least highly constrained. In the NOvA application the vertex is not known and so adaptations were made to solve for the vertex [104].

This algorithm is designed to perform on neutrino interactions and therefore is biased in seeding to prefer vertices at the upstream end (low z value) of the detector where the beam originates. This choice will naturally yield poor performance on cosmic rays where roughly 50% of the time the vertex will be located at the wrong end. There are other tracking algorithms without this bias that are used in the rejection of cosmic rays. Hits without a neighbor within 60 centimeters in a view are assumed to be noise and removed from the fit. After scrubbing, the remaining hits are sorted in z and a box is drawn two meters behind and four meters in front of the hit representing the fifth percentile from the upstream end. Only hits in this z -range are used in the vertex fit. The reasoning is that hits further downstream likely add little information to the vertex search but could contain secondary vertices or be the result of multiple scattering that may confuse the algorithm.

5.5.1 Energy Function

The algorithm searches for the optimum vertex and M arms to describe the N hits in the slice by minimizing an energy function of the form

$$E = \sum_{i=1}^N \sum_{a=1}^M V_{ia} M_{ia} + \lambda \sum_{i=1}^N \left(\sum_{a=1}^M V_{ia} - 1 \right)^2 + \frac{2}{\lambda_v} \sum_{a=1}^M D_a \quad (5.5)$$

where M_{ia} measures distance between cell hit i and arm a , V_{ia} is the strength of association between hit i and arm a , D_a is a measure of the distance between the vertex and the first hit on arm a , and λ and λ_v control the strength of the terms. The first term measures the goodness of fit between the hits and the arms while the second is a penalty term for hits not associated with any arm. The third term is not present in the literature and is a penalty for arms whose first hit is far from the vertex. This term is necessary in the NOvA application where the vertex is not known and is tuned to the distance scale of photon conversions since a common occurrence in neutral current backgrounds is a π^0 decaying into a pair of photons. The likelihood for a photon to travel a distance d before converting is proportional to $\exp(-d/\lambda_v)$, where $\lambda_v \simeq 7/9X_0$ (30 centimeters), leads to an error term

$$\chi^2 = -2 \ln \mathcal{L} = 2 \frac{d}{\lambda_v}. \quad (5.6)$$

Since there are two detector views D_a is taken as $d_a^{xz} + d_a^{yz}$ which are the distances to the closest hit in each view. Technically an arm has an association to all hits in a slice, but for this calculation only hits that have an above average association to an arm are used.

The distance between a hit and arm is computed as

$$M_{ia} = \left(\frac{d_{ia}^{\text{perp}}}{\sigma_i} \right)^2 \quad (5.7)$$

where d_{ia}^{perp} is the perpendicular distance between a hit and the 2D projection of the arm in the detector view and the spacial resolution is based on the half cell depth, $\sigma_i = 3/\sqrt{12}$ cm. For hits in the backwards direction relative to the arm the distance term has been modified as,

$$M_{ia} = \begin{cases} \left(\frac{d_i^{\text{vtx}}}{\sigma_i}\right)^2 & \text{if } \frac{d_i^{\text{vtx}}}{\sigma_i} \leq 1 \\ \left(\frac{d_i^{\text{vtx}}}{\sigma_i}\right)^4 & \text{if } \frac{d_i^{\text{vtx}}}{\sigma_i} > 1 \end{cases}$$

where d_i^{vtx} is the distance from the hit to the vertex.

It is assumed that the likelihood a hit i belongs to arm a is proportional to $\exp(-\beta M_{ia})$ and the likelihood the hit is noise is a constant $\exp(-\beta\lambda)$ so that the association strength is

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta\lambda} + \sum_{b=1}^M e^{-\beta M_{ib}}}. \quad (5.8)$$

The membership of a hit to all arms is bounded by zero and one, where any difference from one represents the noise probability. β can be equated to a temperature ($\beta = 1/T$) such that at higher temperatures arms have a larger sphere of influence in associating with hits. The parameter λ is the distance as measured by M at which a hit has a 50/50 probability of being noise or belonging to an arm.

5.5.2 Seeding and Minimization

For this algorithm to converge on a global minimum vertex solution a proper seed is required since the energy function depends on many parameters and forms a complicated surface. Seeding of the number and direction of the arms as well as the initial vertex solution relies heavily on the results of the Hough transform. The number of arms is set to the largest number of “quality” Hough lines in either readout view, with quality defined as a line with

a peak/threshold ratio above 30. Vertex seeds are produced from two sources:

- In each view hits are sorted in z . The coordinates of hits located at certain percentiles within the list (-3%, -2%, -1%, 0%, 1%, 2%, 5%) are vertex seeds. A negative percentile corresponds to reflecting the z position of a hit across the perpendicular plane associated with the most upstream hit in the slice. The third coordinate is generated from the position along the three most dominant Hough lines in the opposite view. The selected hits in each view are also combined directly to make candidates. The result is as many as 49 potential seeds near the upstream end of the slice and centered around the dominant Hough lines.
- Intersections of the five most dominant Hough lines in a view produce additional seeds. The third coordinate is derived from the position of the three dominant lines in the opposite view. This procedure produces up to 120 vertex seeds.

The arm directions, which are seeded in three dimensions, are also derived from two pools. First, directions can be picked from a minimum bias sample of 20 directions chosen from the vertices of a dodecahedron. Additionally, directions come by forming pairs of Hough lines between the two views where a pair is created if the Hough peaks are of similar size according to

$$-0.8 < \frac{H_{xz} - H_{yz}}{H_{xz} + H_{yz}} < 0.8, \quad (5.9)$$

where H_{xz} and H_{yz} are the value of the Hough peak for a line in the given view.

For each vertex candidate the arms are seeded one at a time with the direction picked from the list that minimizes the energy from Equation 5.5 for a fixed value of β . Co-linearity among the arms is prohibited. Similarly, the final seed chosen is the vertex and arm combination with the minimum energy. The seed is used with the ROOT minimizer MINUET [105] in a simulated annealing process of cooling the temperature β to converge on the global three-dimensional vertex solution. The resultant prong directions are not

taken to be accurate and in some cases end up in false or duplicate directions despite the vertex itself being well reconstructed. The next stage of the reconstruction is relied upon to find the particle directions.

5.5.3 Performance of Vertexing

Performance of the ElasticArms algorithm was evaluated on simulated Far Detector GENIE neutrino interactions. A containment cut was placed requiring the truth vertex to be within +/- 625 cm in x and y and 100 to 5500 cm in z . Resolution was evaluated by comparing the reconstructed and truth vertices in ν_e CC, ν_μ CC, and NC events of both quasielastic and non-quasielastic topology. The results are summarized in Table 5.1 and Figures 5.7, 5.8, and 5.9.

	Δx		Δy		Δz		3D resolution	
interaction	mean	FWHM	mean	FWHM	mean	FWHM	mean	RMS
ν_e CC QE	0.17	3.31	-0.04	3.75	-1.95	5.47	7.56	13.65
ν_e CC Non-QE	0.01	4.53	-0.02	4.61	1.41	7.89	11.44	17.91
All ν_e CC	0.04	4.11	-0.02	4.44	0.73	7.38	10.65	17.20
ν_μ CC QE	-0.15	4.35	-0.17	4.01	-3.83	5.78	10.07	14.59
ν_μ CC Non-QE	0.17	4.37	-0.25	4.70	0.21	8.86	11.93	17.64
All ν_μ CC	0.11	4.43	-0.24	4.61	-0.61	8.16	11.56	17.08
NC QE	1.01	7.02	-1.89	6.21	-2.74	6.28	33.77	36.05
NC Non-QE	0.30	8.22	-0.67	8.03	0.73	11.55	28.30	30.29
All NC	0.35	8.22	-0.77	7.84	0.47	9.50	28.72	30.80
All Interactions	0.13	5.00	-0.26	4.98	0.15	8.25	14.85	22.03

Table 5.1: Vertex resolution for simulated Far Detector neutrino interactions, all numbers in centimeters. The mean and full-width half-maximum (resolution) of the x , y , and z vertex coordinate and the combined 3D difference between the true and reconstructed vertex. Numbers are derived from Figures 5.7, 5.8, and 5.9.

For charged current events the x and y resolution is ~ 4 centimeters which is roughly the dimensions of a cell, with performance being a little poorer for NC events that contain more displaced tracks. In z there is a known bias for single track events where the vertex is pulled 1-2 planes downstream from the true vertex. This bias is caused because the Hough algorithm will frequently produce two lines even for single track slices. Under this two line

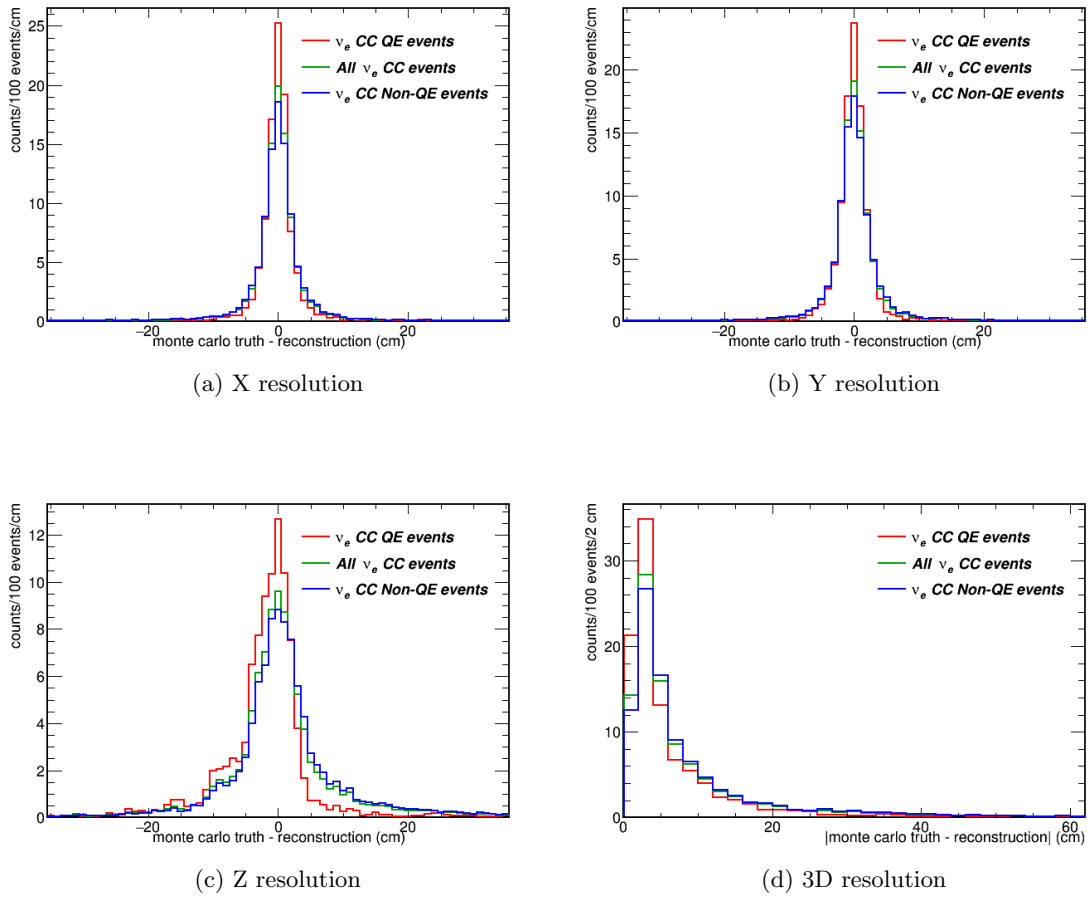





Figure 5.7: Vertex resolution for simulated ν_e CC events in the Far Detector.




(a) X resolution



(b) Y resolution




(c) Z resolution




(d) 3D resolution


Figure 5.8: Vertex resolution for simulated ν_μ CC events in the Far Detector.




(a) X resolution



(b) Y resolution



(c) Z resolution



(d) 3D resolution


Figure 5.9: Vertex resolution for simulated NC events in the Far Detector.

assumption the vertexing algorithm tends to shift the vertex forward from the upstream end in order to send one prong in the backwards direction of the fitted vertex. This bias is more common in quasi-elastic (QE) events where the proton track can be very short or co-linear with the lepton such that only one track is visible. The average 3D distance between the true and reconstructed vertex is about 12 centimeters (2 planes) for CC interactions. For NC events that frequently have activity (photons, neutrons) displaced from the vertex there is a bias to pull the reconstructed vertex downstream from the true vertex, particularly in QE events where there are fewer tracks to constrain the vertex with.


In the Near Detector vertex reconstruction was evaluated on the dominant ν_μ CC interaction mode with simulated events. A cut was placed on the true vertex at ± 140 centimeters in x and y and 100 to 700 centimeters in z . The performance achieved is on par with the Far Detector resolution, as shown in Figure 5.10. The resolution is ~ 4 centimeters in x and y , ~ 8 centimeters in z and on average within ~ 11 centimeters of the true vertex. In the Far Detector neutrino interactions are rare and so the dominant event type is cosmic backgrounds which will lead to the vertex being placed on the wrong track end half the time due to the upstream bias in vertex seeding and the isotropic nature of the cosmic background. In the Near Detector where virtually all events are beam neutrinos it is possible to compare the data and Monte Carlo agreement. This was done with 1.66×10^{20} POT of Near Detector data and Monte Carlo scaled to match. A standard set of preselection cuts were used (see Section 6.4), with results in Figure 5.11. Given the $\sim 30\%$ uncertainty in the beam flux (see Section 7.3.1) the data and Monte Carlo show good agreement.

5.5.4 Improvements to Vertexing


Currently only one 3D vertex is produced for the slice. A potential improvement for future versions of the reconstruction would be to include additional secondary vertex choices which




(a) X resolution



(b) Y resolution



(c) Z resolution



(d) 3D resolution

Figure 5.10: Vertex resolution for simulated ν_μ CC events in the Near Detector.

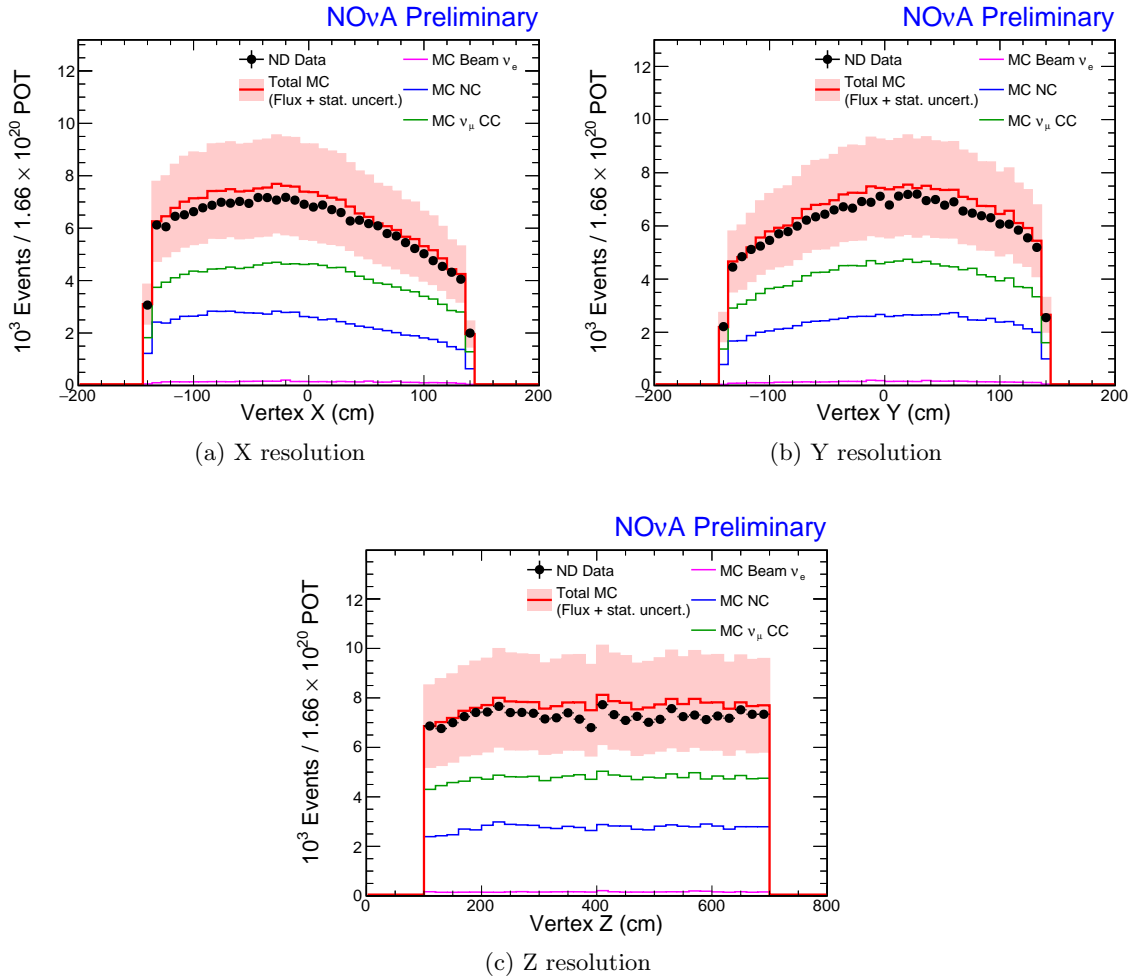


Figure 5.11: Comparison of the reconstructed vertex distribution in data and Monte Carlo for the Near Detector. Red bands show the beam flux uncertainty. Figure from [106].

could improve the track reconstruction for certain topologies, particularly in higher energy or deep inelastic scattering events. Another improvement would be to search for the vertex in both the upstream and downstream directions which can aid in the rejection of certain cosmic background events. This feature would be helpful in rejecting photons coming from the back wall of the detector that can look like electron showers and currently have the vertex placed at the wrong end. Similarly, bremsstrahlung radiation from cosmic rays can sometimes be sliced independently from the muon and then incorrectly vertexed.

5.6 Prong Formation with Fuzzy K-Means

After a vertex candidate has been produced the next stage is to produce track prototypes referred to as “prongs”. This step is done with an adapted version of the classic fuzzy k -means clustering algorithm [107], where k refers to the number of clusters and “fuzzy” allows an object to have membership in multiple clusters. The classic algorithm is described in Section 5.6.1 and then adaptations necessary to handle noise and the unknown number of prong objects follows in Section 5.6.2. Prong formation is done in each detector view separately and then 3D prongs are formed through a matching process described in Section 5.6.5.

In this application the vertex is used to recast the initial prong formation into a one dimensional problem. The assumption of the algorithm is that all hits in a slice originate from an interaction at that vertex position and are organized in prongs emanating from that vertex. To find the prongs, each cell hit is converted to an angle with respect to the vertex (in 2D) using the cell centers. The forward z direction is defined as 0 radians. The upper hemisphere then ranges from 0 to π , while the lower hemisphere represents 0 to $-\pi$. An uncertainty is then associated with each cell angular position based on the distance d of the hit from the vertex. The uncertainty function is modeled after the multiple scattering

behavior of muons and takes the form

$$\sigma = \frac{1.745}{d} + 0.0204 + 0.000173d. \quad (5.10)$$

This equation resulted from simulating 1 and 2 GeV muons and computing the distribution of differences in angle of the cell hits from the initial trajectory as a function of distance, as shown in Figure 5.12. A cutoff distance of five meters is set and the angular uncertainty




Figure 5.12: In the left panel are simulated 1 GeV muons in the Far Detector and in the right 2 GeV. The x-axis is the distance in centimeters of the hit to the vertex and the y-axis is the difference between the angle of the hit with respect to the vertex and the initial true momentum of the particle. The function drawn on the plot in red is equation 5.10.

remains constant beyond that point. This choice may cause the ends of very long muons with a hard scatter to be split into separate prongs, however performance on muons is still very strong as discussed in Section 5.6.6. As this algorithm is chiefly designed for the analysis of electron neutrinos the choice is acceptable since electrons in the signal region below 3 GeV will typically be beyond shower max in 5 meters and could pull in false hits with an expanding uncertainty. The first five meters of a muon track would still be sufficient to reject the event as background and other algorithms exist tailored to precision muon reconstruction. Other forms of the uncertainty function based on electron scattering were also considered but in testing this general form was found to be the most broadly applicable at this stage of reconstruction. It is envisioned that a redistribution of hits could

occur at a later stage of reconstruction when particle identification information could be used to refine the prong membership.

5.6.1 Classic Fuzzy K-Means

The fuzzy clustering problem is well known in the literature and the theory has been developed over several decades [107, 108]. In the NOvA application it is used to place k prong centers in the one dimensional angular space to minimize the angular distance to n hits. This is accomplished by minimizing a score function of the form

$$J_{FKM}(\mu, a) = \sum_{i=1}^c \sum_{j=1}^n \mu_{ij}^m \|x_j - a_i\|^2, m > 1 \quad (5.11)$$

where m is the degree of fuzziness of the prongs, x_j the position of the j^{th} hit, a_i the i^{th} prong center, and μ_{ij} is the membership of the i^{th} hit in the j^{th} prong. If $m = 0$ the fuzziness is removed and a hard clustering model results with each hit belonging entirely to a unique prong. The prong associations become fuzzier for larger m . The hit membership and prong center locations are updated iteratively by

$$\mu_{ij} = \left(\sum_{k=1}^c \frac{\|x_j - a_i\|^{2/(m-1)}}{\|x_j - a_k\|^{2/(m-1)}} \right)^{-1}, i = 1, \dots, c, j = 1, \dots, n \quad (5.12)$$

and

$$a_i = \frac{\sum_{j=1}^n \mu_{ij}^m x_j}{\sum_{j=1}^n \mu_{ij}^m}, i = 1, \dots, c. \quad (5.13)$$

In this model the hit memberships are always normalized

$$\sum_{i=1}^c \mu_i(x) = 1. \quad (5.14)$$

There are two main flaws when this type of cluster model was applied to the NOvA event

topologies. The first is that this model was built under the assumption that the optimal number of clusters, c , is already known, which is not generally true in this application. In fact, given the tendency for the multi-Hough and Elastic Arms algorithms to over estimate the proper number of prongs, it is best for this algorithm to optimize the prong number independently. The second failing is the fact that the normalization forces the total membership of a hit to be one. This creates scenarios like the one in Figure 5.13 where noise hits would be improperly associated with prongs.




Figure 5.13: Points 1 and 2 are both equidistant from clusters A and B . In a fuzzy k -means model this means that both points would have a membership of 50% to each cluster even though point 2 is clearly more of an outlier and is better suited to be flagged as noise or form a new cluster.

5.6.2 Possibilistic Fuzzy K-Means

In order to address the second problem a switch was made to a possibilistic clustering algorithm (PCA) [109]. This approach differs from the classic fuzzy k -means algorithm by allowing the total prong membership of a cell to float. This allows membership of noise hits to approach zero, making the algorithm more robust against outliers. While this approach to prong formation is advantageous for this application, there are some trade offs. There now exists a trivial solution to the problem in which every hit is classified as noise and thus

nothing is assigned to a prong. There are several approaches in the literature to avoid this trivial solution, with one form seeming best suited for the NOvA application [110].

The degree of membership of hits to the prongs is calculated in a two step process. First the distances between the hits x and the prong centers a are calculated

$$d_{ij} = \left(\frac{x_j - a_i}{\sigma_j} \right)^2, \quad -\pi \leq (x_j - a_i) \leq \pi \quad (5.15)$$

and then prong membership is updated with

$$\mu_{ij} = \exp - \frac{m\sqrt{a}d_{ij}}{\beta}. \quad (5.16)$$

In the above equation the “fuzziness factor” $m = 2$. The β parameter is a new term introduced in the PCA [110]. It is a normalization term that measures the degree of separation of the data set. In this application the distances between hits and prong centers can be measured in units of standard deviations about the prong center. So $\beta = 4$, represents a 2σ spread about a prong center. In the above expression a is the number of prong centers, meaning that simply adding more prongs will slightly lower all membership probabilities, even if the addition is correct.

After prong membership is calculated the next step is to update the prong centers,

$$a'_i = a_i + \frac{\sum_{j=1}^n \frac{u_{ij}^m}{\sigma_j^2} (x_j - a_i)}{\sum_{j=1}^n \frac{u_{ij}^m}{\sigma_j^2}}. \quad (5.17)$$

The new prong centers are compared to the old ones and the deviation must be less than the specified tolerance, currently 1×10^{-7} radians, or the minimization continues for another round.

At the end of the process an overall objective score for the prong formation is calculated

$$J_{PCA}(\mu, a) = \sum_{i=1}^c \sum_{j=1}^n \mu_{ij}^m \|x_j - a_i\|^2 + \frac{\beta}{m^2 \sqrt{c}} \sum_{i=1}^c \sum_{j=1}^n (\mu_{ij}^m \log \mu_{ij}^m - \mu_{ij}^m). \quad (5.18)$$

This is the analog to Equation 5.11 for the FKM algorithm. The first term is identical to the FKM function. The second term is a constraint to take into account the fact that membership probabilities are not normalized. This ensures that the objective function will be minimized when the μ_{ij} are large, avoiding the trivial solution $\mu_{ij} = 0$. This equation can be simplified by solving Equation 5.16 for $\|x_j - a_i\|^2$ in terms of μ_{ij} :

$$J'_{PCA}(\mu, a) = -\frac{\beta}{m^2 \sqrt{c}} \sum_{i=1}^c \sum_{j=1}^n \mu_{ij}^m. \quad (5.19)$$

Currently, the objective score is not used downstream in the analysis as an evaluation of the prong making. It was found that the variety in neutrino interaction topologies and energies complicate interpreting the number as “good” or “bad”. In future versions of the algorithm a proper quality assessment could prove beneficial, but as will be discussed in Section 5.7, the event classification currently performs at a sufficiently high level for this analysis without this assessment of reconstruction quality.

5.6.3 Seeding

The results of the PCA algorithm have some dependence on the seeds that are chosen. There is also a tendency for multiple prongs to end up centered at the same location. To deal with each of these issues the prongs are seeded at the angles corresponding to the densest hit activity. Since the optimum number of prongs is not known at the start, the algorithm begins with a one prong assumption. The angle with the highest hit density is calculated by populating a density matrix w consisting of 360 bins ranging from $-\pi$ to π

with

$$w_k = \sum_{i=1}^n e^{-\left(\frac{\theta_k - \theta_i}{\sigma_i}\right)^2} \quad (5.20)$$

and

$$\theta_k = -\pi + \frac{k * \pi}{180}, 0 \leq k < 360. \quad (5.21)$$

The maximum bin is used for the initial seed angle.

5.6.4 Optimizing Prongs and Cleanup

After a stable set of prong centers is found a check is performed to find and remove duplicates. There is the possibility with this algorithm that multiple seeds can converge onto the same final position. If duplicates are found the seed that produced the duplicate position is added to a list of bad seeds. Then the original density matrix calculated in seeding the prongs is consulted again. The next highest density peak in the map is chosen as a replacement seed provided it has not been used previously. If a replacement seed is found the clustering is done again. If no replacement can be found in the list, then the total number of prong centers is reduced by one and the algorithm proceeds to check for un-associated hits. This is done by looping over all hits and ensuring that each hit has a membership to at least one prong that is above the threshold, currently set to 1%. All hits that do not meet this criteria are fed into a new density matrix to find a seed. The angle with the highest density is added as a new seed, provided it is not on the list of seeds known to produce duplicates, and the overall number of prongs is increased by one. The prong formation steps above are then repeated with the new seed in addition to the final prong positions calculated previously. This entire process repeats until all hits have a minimum membership to a prong or a maximum of seven prongs has been created.

After two dimensional prongs have been produced in each view, a cleanup step is applied

to determine if any of the prongs should be merged together. Prongs within 30 degrees of each other will be merged provided a minimum of 66% of the hits in the smaller prong also share a minimum membership in the larger prong. This check is intended to help minimize the fringes of showers from being split into separate prongs.

Another check is used for the case of co-linear displaced particles. One useful tool for the later stages of the analysis is to detect a gap between the start of a prong and the vertex which can indicate photons coming from a π^0 decay. Since the prong formation is done only in angular space, slight displacement of the reconstructed vertex can cause a couple hits near the vertex that fall on the same angle as a photon several planes away to be grouped together. Additionally, there could be multiple particles that fall on the same angle with a large spacial gap between them. One concern is that there exists a non-trivial probability for a particle to pass through one or more planes without depositing any energy due to the plastic partitions between cells (the detector is $\sim 65\%$ active scintillator). If a gap between hits originates within two planes of the vertex an aggressive approach was taken to attempt to preserve gaps in π^0 's. Prongs were separated if there is a gap of at least 25 centimeters and at least one empty plane in the same view between hits. Far from the vertex splitting requires a gap of 75 centimeters (approximately two radiation lengths) and at least 4 empty planes in the same view between hits.

It should be noted that it is possible for hits, especially near the vertex, to belong to multiple prongs. Currently the task of splitting the energy of a hit among the prongs to avoid double counting is left to downstream algorithms. It was found that the membership weights of a hit within a prong were not equivalent to the weights for splitting energy. Proper energy weighting requires assumptions about the particle type of the prongs and incorporating information from the other view to account for attenuation. To keep the prong formation as generally applicable as possible it is best to distribute the energy during

particle identification and event classification when the prongs are treated with particle assumptions as discussed in Section 5.7.1.

5.6.5 Forming 3D Prongs

After a set of two dimensional prongs has been produced in each view the final step of the algorithm combines information into three dimensional prongs. The criteria for matching compares the deposition of energy along the length of a prong in each view. The matching requires the attenuation calibration so that position effects on the energy deposition can be removed. Prongs for which a suitable match cannot be found remain as 2D prongs which still provide information to the analysis.

The first step is to build a matrix of all possible prong pair combinations between views. Prongs are required to overlap endpoints by at least one plane for a pair to be considered. If a prong only has hits in a single plane a stricter requirement is made requiring a matching prong in the other view to be at most two planes and to occupy only adjacent planes. For the candidate match a temporary 3D prong is made so that information from the other view can be applied to the attenuation formula to get a calibrated energy of each hit. Then the cumulative fraction of energy as a function of the three dimensional path length is calculated for each view. The energy of a hit is distributed evenly among a 3×3 grid within the cell to smooth the function. A measure of similarity between the energy distribution in each view is calculated through the Kuiper test [111]. This test is a modified form of the Kolmogorov-Smirnov test where

$$K = D^+ + D^- \quad (5.22)$$

$$D^+ = \max (E^{XZ}(s) - E^{YZ}(s)) \quad (5.23)$$

$$D^- = \max (E^{YZ}(s) - E^{XZ}(s)) \quad (5.24)$$

with $E^{XZ}(s)$ and $E^{YZ}(s)$ being the fraction of energy accumulated in the respective view at a path length of s along the prong. By calculating the largest difference between the curves in both directions the metric is more sensitive to differences in shape at the tails than the KS test. If the maximum value of D^+ or D^- is less than zero, meaning one curve is always above the other, the value is set to zero. A lower score, K , indicates a better match. There are natural jumps that occur in the energy distributions of each view due to the alternating plane geometry of the detector. Matching performance improved if the path length in one view was allowed to shift to better align the curves and remove detector geometry effects. The curve for one view is allowed to shift between -12 and 12 centimeters in path length in steps of 0.5 centimeters, the solution with the lowest score is kept.

After scores have been computed, matching proceeds starting with the set of prongs with the lowest score, and continues until all matches have been made. A prong is only allowed to perform one match and is then removed from the list. There are situations in the detector where two particles will be overlapping in one detector view such that they form one prong, but are distinct in the other view. In the present implementation of the algorithm the merged prong would be matched to one of the distinct prongs and the other would either form an incorrect match elsewhere or be left as a 2D prong. There is not a cutoff score for determining if a match should be allowed. An example of the matching process is shown in Figures 5.14 and 5.15.

5.6.6 Performance of Fuzzy K-Means Algorithm

The performance of the prong reconstruction is evaluated with simulated Far Detector neutrino events with a simulated true visible energy deposited in the detector between 0.5 and 3.5 GeV. A containment cut was placed requiring the truth vertex to be within +/- 625 cm in x and y and 100 to 5500 cm in z . The principle performance metrics are completeness,




Figure 5.14: A simulated ν_e CC quasielastic interaction in the Far Detector with completed 3D prong reconstruction from the fuzzy-k means algorithm. The prong outlined in red in each view is the reconstructed electron and outlined in green is the reconstructed proton. The corresponding energy profile histograms used to compute the suitable 3D prong matches is shown in Figure 5.15.




Figure 5.15: Cumulative energy profile as a function of path length along a prong for perspective 3D match candidates shown in Figure 5.14. The red curves are for prongs in the XZ (vertical planes) view and the blue is for the YZ (horizontal planes) view. The upper-left and lower-right panels show the preferred matches with similar energy profiles that result in the green and red track respectively. The off-diagonal elements illustrate the difference in energy profile shape for the wrong combinations.

purity, and angular resolution. In tuning the algorithm, correct reconstruction of the most energetic prong took precedence as identification of the lepton in charged current interactions plays an important role in the event classification discussed in Section 5.7. For true lepton energies above 0.5 GeV the reconstructed prong achieves a completeness above 90%, purity of 80%, and an angular resolution better than 5 degrees for both quasielastic and non-quasielastic event topologies as shown in Figures 5.16 and 5.17. The reconstruction is similarly effective when performance of the lepton reconstruction is examined as a function of neutrino energy, shown in Figures 5.18 and 5.19. It is important that the most energetic shower be well reconstructed regardless of whether it is a lepton or not. For reconstructed energies above 0.8 GeV the primary prong has a completeness above 90%, purity of 80% and an angular resolution below 10 degrees as shown in Figure 5.20.

For the Near Detector performance was evaluated with ν_μ CC interactions. A cut was placed on the true vertex at ± 140 centimeters in x and y and 100 to 700 centimeters in z . For neutrino energies above 0.5 GeV the performance in the Near Detector is similar to that of the Far Detector as shown in Figure 5.21. There is a dip in purity above 2 GeV due to a higher rate of deep-inelastic scattering events which produce overlapping prongs. In the Near Detector, where virtually all events are beam neutrinos, it is possible to compare the data and Monte Carlo agreement. This was done with 1.66×10^{20} POT of Near Detector data and Monte Carlo scaled to match. A standard set of preselection cuts were used (see Section 6.4), with results shown in Figure 5.22. The number of hits in a prong and prong length show good agreement given the flux uncertainties. The data has a preference to make fewer 3D prongs than the Monte Carlo likely due to uncertainties in the neutron rate and hadronic energy in the simulation model. Studies have shown [112] that hadronic hits associated with neutrino interactions are seen in simulation then data. These extra hits, which can manifest as spittle at the fringes of events, can produce additional small prongs.

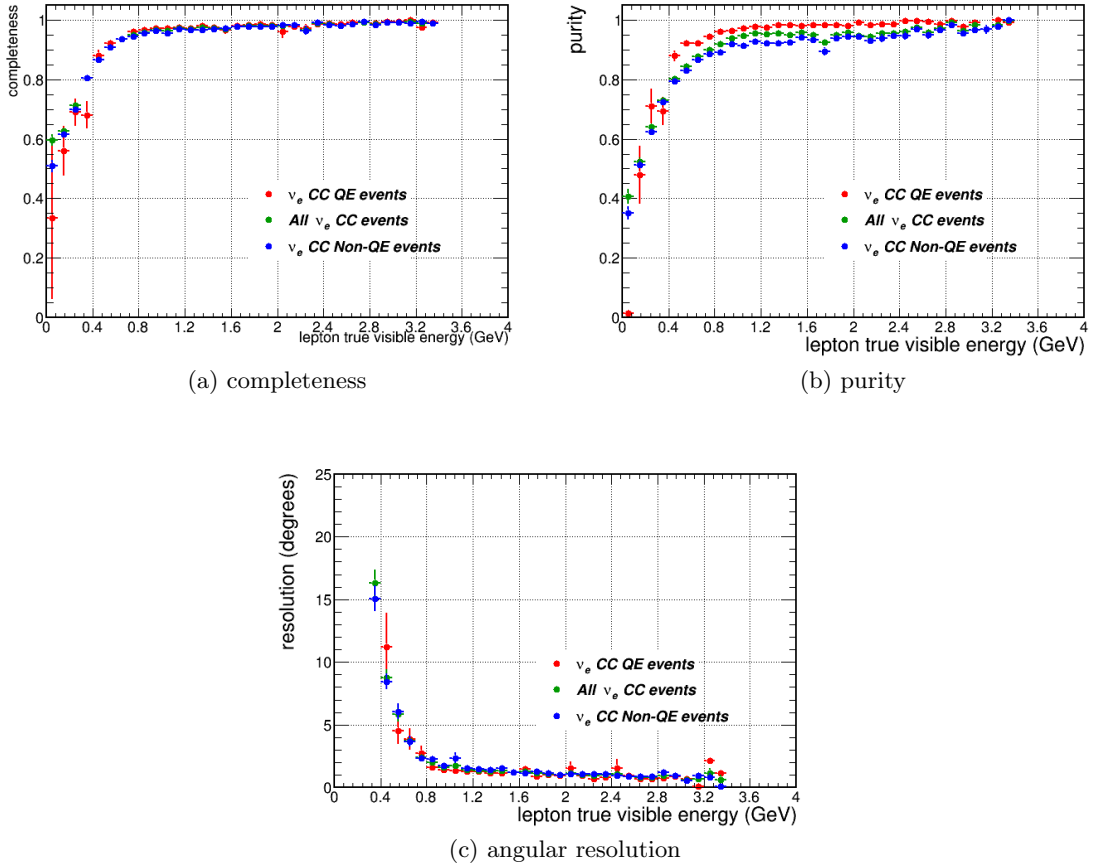


Figure 5.16: Completeness (a), purity (b), and angular resolution (c) of the reconstructed electron in simulated Far Detector ν_e CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the lepton’s visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 625$ cm and $100 < z < 5500$ cm.

5.6.7 Improvements to Prong Formation

The fuzzy-k algorithm has been shown to perform well in the one to three GeV signal energy range for an electron neutrino appearance oscillation analysis. The event classification discussed in Section 5.7 is based on the most energetic prong in the slice, which is well reconstructed. There are planned performance improvements for future versions of the algorithm geared at improving performance for secondary prongs in complex event topologies.




Figure 5.17: Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the lepton's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 625$ cm and $100 < z < 5500$ cm.




Figure 5.18: Completeness (a), purity (b), and angular resolution (c) of the reconstructed electron in simulated Far Detector ν_e CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 625$ cm and $100 < z < 5500$ cm.






Figure 5.19: Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino’s visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 625$ cm and $100 < z < 5500$ cm.



(a) completeness

(b) purity



(c) angular resolution

Figure 5.20: Completeness (a), purity (b), and angular resolution (c) of the most energetic reconstructed prong in simulated Far Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the prong's reconstructed energy. A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 625$ cm and $100 < z < 5500$ cm.







Figure 5.21: Completeness (a), purity (b), and angular resolution (c) of the reconstructed muon in simulated Near Detector ν_μ CC interactions depositing 0.5 to 3.5 GeV of visible energy, plotted as a function of the neutrino's visible energy. Distributions are shown for quasielastic events (red), non-quasielastic events (blue) and the total (green). A containment cut was placed on the selection requiring the true event vertex to be in $|x, y| < 140$ cm and $100 < z < 700$ cm.



(a) # of prongs



(b) # of prong hits



(c) prong length

Figure 5.22: Comparison of the number of 3D prongs (a), number of hits in the most energetic prong (b) and length of the longest prong in a slice (c) in data and Monte Carlo for the Near Detector.

The most significant performance gains in this algorithm are likely to come from modifications to the matching of 2D prongs into 3D prongs. Instead of calculating scores for a matrix of one to one prong matches it would be better to evaluate the matching quality of the event as a whole. In this way scenarios where two or more prongs in one view are best matched to the same prong in the other view can be properly evaluated. This improvement is critical in particular for detailed cross-section studies in the Near Detector where event classification should occur on a particle by particle basis and not just for the overall slice.

On a related note, determining the quality of a match in order to allow a prong to remain 2D rather than accept a bad match could also improve results. A quantitative assessment of reconstruction quality could be applied to the slice as a whole to inform downstream analysis stages about which slices to analyze. Normalizing the matching scores based on the number hits in a prong and considering the number of gaps (missing planes) is being evaluated.

Finally, the prong reconstruction has a complex parameter space of tuning knobs and small performance gains might be seen through a global optimization. One possible path of study is to broadly sample the parameter space with special sets of files and use a genetic algorithm to identify a promising region and “breed” new generations of samples to converge on a solution. This intensive computational effort is being delayed until a time when the computing resources of the experiment are not needed for higher priorities.

5.7 Event Classification

The final step in the reconstruction chain used for the analysis in this thesis is to classify the slice according to the degree to which it is consistent with a ν_e CC interaction which is the desired signal. This step is performed with a likelihood identification (LID) algorithm that uses reconstruction quantities derived from fuzzy-k prongs as input variables to an

artificial neural network (ANN). The primary input variables are a series of likelihoods that result from calculating dE/dx profiles for the prongs and comparing to templates for different particle hypotheses. The algorithm is described in detail in [113]. Note that the inputs to the neural network, unless specified, utilize only the most energetic prong. The classification of the slice through this algorithm operates under the assumption that the primary lepton in a CC interaction is in the most energetic prong which is correct 85% of the time in the neutrino energy region of interest between 0.5 and 3.5 GeV. For simulated Far Detector oscillated events passing the expected preselection the algorithm is $\sim 48\%$ efficient in identifying ν_e CC signal events with a background rejection of $\sim 97\%$.

5.7.1 Hit Redistribution

The first step in the algorithm is to redistribute hits and energy in the 3D fuzzy-k prongs, which will now be referred to as shower candidates. Hits are never dropped from a prong but may be added either from ungrouped hits or other prongs. In the first eight planes of a prong, all hits within two cell widths (one width is ~ 4 cm) of the prong direction are added. Outside of this range hits within 20 cell widths are added. For hits that are now members of multiple prongs the energy weights are calculated according to:

$$w_{ik}^{\text{cell}} = \frac{PE_i^{\text{shower}} e^{-D_{ik}/\lambda}}{\sum_{j=1}^N PE_j^{\text{shower}} e^{D_{jk}/\lambda}} \quad (5.25)$$

where w_{ik}^{cell} is the weight of the k^{th} cell in the i^{th} shower, D_{ik} is the distance from the hit to the core of the shower, λ is a constant derived from the simulated transverse profile of electrons and is 3.05 cm, PE^{shower} is the total number of photo-electrons in the shower, and N is the total number of showers (must be greater than one) that hit k is a member of.

5.7.2 dE/dx and Log-Likelihood Calculations

After the formation of prong candidates the longitudinal and transverse dE/dx profiles are calculated. The longitudinal component is calculated plane-by-plane as the sum of energy in that plane divided by the path length through the plane. The plane in which the start point of the shower is located is numbered as 0 and the rest are indexed accordingly in the direction of the shower. In an attempt to normalize the calculation for various shower angles an approximation is made:

$$P'_i = \frac{P_i}{D_z} \quad (5.26)$$

so that a plane P_i in a shower with a z -component D_z to the shower direction is translated into plane number P'_i . A more rigorous technique of rotating the plane axis into the reference frame of the shower and redistributing the energy of the hits into the new pseudo-planes was considered but found to offer only slight improvements since most showers are at small angles with respect to the z -direction.

The transverse dE/dx is computed through a series of steps:

1. In each plane the energy weighted average cell position is computed. This cell has a transverse cell index of 0.
2. In each plane moving outward in units of one cell in both directions from the core cell is cell index one, two, and so on capping at 20.
3. For a given transverse cell index the energy in all planes is summed and divided by the total path length to compute the average dE/dx for that cylindrical slice of the shower.

These dE/dx distributions can be used to differentiate particle types. From simulated Monte Carlo neutrino interactions reconstructed prongs that match by truth to e , γ , μ , π^0 , p , n , and π particles are used to generate dE/dx templates. The templates are divided into eleven energy bins between 0 and 5.25 GeV. The templates are further subdivided into four


quadrants based on the xy vertex position to minimize attenuation and threshold effects for events occurring at different distances from the APDs. Longitudinal templates are made for the first 200 planes in the shifted coordinate frame from Equation 5.26 and the transverse direction is sampled in 20 cylindrical planes. The top panel of Figure 5.23 compares the longitudinal dE/dx distributions for two planes in electrons and photons above 0.5 GeV. The photon distribution is much broader in the early planes than an electron shower since at the start a photon converts into an e^+e^- depositing two MIP's of energy at the shower start. The bottom panel of Figure 5.23 compares the longitudinal distributions for electrons and muons. In plane 10 it can be seen that an electron has a large spread in the energy distribution as the electromagnetic shower grows while a muon continues to make relatively uniform energy depositions of a MIP. In Figure 5.24 are transverse dE/dx distributions for electrons and π^0 's (meaning the daughter particles associated with the π^0 decay). Due to the possibility for overlapping photons in a π^0 decay to be merged into one reconstructed shower the transverse shower profile is expected to be broader.

The energy profiles provide a powerful handle on identifying the reconstructed objects. After calculating the dE/dx for each longitudinal plane and transverse cell in a reconstructed shower the probability of that value occurring is computed for each particle hypothesis:


$$P_i^\alpha = \frac{N^\alpha(x) \times n_{bin}^\alpha}{N_{tot}^\alpha}. \quad (5.27)$$

P_i^α is the probability of finding a dE/dx value of x in longitudinal or transverse plane i for particle hypothesis α where N_{tot}^α is the total number of entries in the particle histogram, n_{bin}^α is the number of bins, and $N^\alpha(x)$ is the number of entries in the corresponding dE/dx bin. This probability can be translated into a likelihood, LL_i^α , through:

$$LL_i^\alpha = \ln(P_i^\alpha). \quad (5.28)$$



(a) electron and photon longitudinal dE/dx



(b) electron and muon longitudinal dE/dx

Figure 5.23: The longitudinal dE/dx profile comparing simulated electrons (red) to other particles (blue), specifically photons (a) and muons (b). The left panels show the profile for plane 2 of the shower and the right panels show plane ten. Figure from [113].




Figure 5.24: The transverse dE/dx profile for simulated electrons (red) and π^0 's (blue). The left panel shows the profile for transverse cell index 0 and the right shows transverse cell index 3. Figure from [113].

The total longitudinal (LL_L^α) and transverse (LL_T^α) likelihoods for particle hypothesis α are defined as:

$$LL_L^\alpha = \frac{\sum_i LL_i^\alpha}{N_{\text{plane}}} \quad (5.29)$$

$$LL_T^\alpha = \frac{\sum_i LL_i^\alpha}{N_{\text{trans. cell}}} \quad (5.30)$$

where N_{plane} is the number of longitudinal planes in the shower and $N_{\text{trans. cell}}$ the number of transverse cells. Figures 5.25 and 5.26 compare the longitudinal and transverse likelihood values for true electrons to the other true particle hypotheses. Differences in these likelihoods will be used in the identification process to find slices containing electron neutrino interactions.

5.7.3 Neural Net Input Variables


In total 16 variables were found useful in building the neural net for electron neutrino classification. The variables are calculated at the shower level and so a particle identification (PID) value will be assigned to each 3D shower in a slice. The PID value for the most energetic shower will be used as the classification of the entire slice. For each candidate shower the variables are:

1. Twelve of the input variables are the log-likelihood differences between electrons and the other six particle hypotheses ($p_h = \gamma, \mu, \text{proton, neutron, pion, } \pi^0$), for both the longitudinal and transverse shower profile:


$$LL_{\text{long}}(e/p_h) = LL_{\text{long}}(e) - LL_{\text{long}}(p_h)$$

$$LL_{\text{trans}}(e/p_h) = LL_{\text{trans}}(e) - LL_{\text{trans}}(p_h)$$


2. m_{π^0} : The invariant mass is calculated between the candidate shower and each other



(a) $LL_{e^-} - LL_\gamma$




(b) $LL_{e^-} - LL_\mu$




(c) $LL_{e^-} - LL_p$


Figure 5.25: The likelihood differences between electrons and other particles: photons (a), muons (b), and protons (c). True electrons are plotted in red, the other particle in blue. The left panels show the longitudinal likelihood difference and the right panels show the transverse. Figure from [113].



(a) $LL_{e^-} - LL_n$



(b) $LL_{e^-} - LL_{\pi}$



(c) $LL_{e^-} - LL_{\pi^0}$

Figure 5.26: The likelihood differences between electrons and other particles: neutrons (a), pions (b), and π^0 's (c). True electrons are plotted in red, the other particle in blue. The left panels show the longitudinal likelihood difference and the right panels show the transverse. Figure from [113].

shower in the slice. The mass that is closest to the nominal π^0 mass is chosen for the variable.

3. $E_i^{\text{shower}} / \sum_i E_i^{\text{shower}}$: The energy ratio between the candidate shower (E^{shower}) and the summed energy of all showers in the slice.
4. E_{vertex} : The energy deposited in GeV within 8 planes of the vertex from the entire slice excluding the candidate shower.
5. Gap: The distance in centimeters between the vertex and the start of the candidate shower.

5.7.4 Artificial Neural Network and Training

The variables outlined in the previous section are used in an artificial feed-forward neural network [114, 115] (ANN) to produce the ν_e CC classification variable known as LID. The algorithm maps a set of input variables $x(x_1, x_2, \dots, x_N)$ onto an output variable $O(x)$ with three hidden layers of neurons in between. Each neuron layer only passes information forward such that the output of a layer does not effect the same layer. For a neuron j in layer k its output can be described by the equation:

$$x_j^k = A \left(w_{0j}^k + \sum_{i=1}^{M_{k-1}} w_{ij}^k x_i^{k-1} \right), \quad (5.31)$$

where M_{k-1} is the total number of neurons in layer $k - 1$, x_i^{k-1} are the input signals from the previous layer, w_{ij}^k is the synaptic weights applied for neuron j , and w_{0j}^k is a bias term from adding a new synapse to neuron j whose input is $x_{0j}^{k-1} = 1$. $A(x)$ is an activation function that handles the transfer from input to output. The activation function takes the form of the Sigmoid function:

$$S(x) = \frac{1}{1 + e^{-x}} \quad (5.32)$$

and will produce values between 0 and 1. The output value from the ANN will represent the degree to which an event is seen as ν_e CC with values near 1 representing signal and 0 for background. Other types of neural network were explored and evaluated based on the figure of merit (FOM) of $\text{signal}/\sqrt{\text{background}}$ and performance was found to be similar. In the same way the number of neurons was optimized with a decision made to use three hidden layers with 22, 12, and 6 neurons. Figure 5.27 shows a visualization of the network.

The neural network was trained on a sample of simulated Far Detector neutrino interactions between 0 and 5 GeV. Neutrino oscillation weights were applied to the sample assuming no matter effect or CP violation with $\sin^2 2\theta_{23} = 1$ and $\sin^2 2\theta_{13} = 0.1$. These weights represent an average oscillation result so as not to bias the training with assumptions about mass ordering and CP violation. Events were selected which had less than 200 cells in the slice and a longest prong length less than six meters to remove some ν_μ CC background. Signal events are further required to have the direction of the most energetic reconstructed shower match within thirty degrees of the true electron. The sample was composed of a roughly 2 to 1 background (ν_μ CC and neutral current) to signal (ν_e CC) ratio and in total was about 200,000 events for training and 200,000 for testing. Training takes place over 600 iterations to minimize the mean squared error. The importance, I_i , of the i input variables can be ranked by the sum of the weights-squared between the input variable neuron and the 22 neurons in the first hidden layer:

$$I_i = \bar{x}_i^2 \sum_{j=1}^{22} \left(w_{ij}^{(1)} \right)^2, \quad (5.33)$$

where \bar{x}_i is the sample mean of the input variable. Variable rankings from the training are shown in Table 5.2.

All ANN training is performed on Far Detector simulations. The training results are also applied to the Near Detector as it is functionally identical. Since the Near Detector front




Figure 5.27: The structure of the artificial neural network passing from 16 input variables through three hidden layers of neurons to the output. Variable names and importance shown in Table 5.2. Figure from [113].

Rank	variable	description	importance
1	enllt	electron - neutron trans. likelihood difference	0.33
2	epllt	electron - proton trans. likelihood difference	0.32
3	epillt	electron - pion trans. likelihood difference	0.29
4	eplll	electron - proton long. likelihood difference	0.27
5	egllt	electron - photon trans. likelihood difference	0.26
6	shE	fraction of slice energy in most energetic shower	0.25
7	epilll	electron - pion long. likelihood difference	0.24
8	eglll	electron - photon long. likelihood difference	0.19
9	emlll	electron - muon long. likelihood difference	0.19
10	enlll	electron - neutron long. likelihood difference	0.15
11	epi0lll	electron - π^0 long. likelihood difference	0.13
12	pi0mass	mass of π^0 candidate in slice	0.12
13	epi0llt	electron - π^0 trans. likelihood difference	0.09
14	vtxgev	energy deposited near vertex	0.06
15	emllt	electron - muon trans. likelihood difference	0.05
16	gap	distance between vertex and shower start	0.04

Table 5.2: Near Detector data (1.6×10^{20} POT) and Monte Carlo (normalized to data) event counts passing the levels of preselection cuts. Table from [113].

face is $1/16^{\text{th}}$ the size of the Far Detector the approach of separating the dE/dx histograms into quadrants based on vertex is not needed. The entire Near Detector is approximated as one quadrant of the Far Detector where both the x and y position are near the readout electronics.

5.7.5 Performance of Event Classification

During the process of constructing the LID algorithm performance was evaluated using simulated Far Detector events. A standard event selection is applied, outlined in detail in Section 6.3 and representative oscillation weights without a matter effect were applied to simulate an average result. The result is scaled to 2.8×10^{20} full detector equivalent POT (see Section 6.6). The LID performance is shown in Figure 5.28. A LID cut at 0.95 is placed in the analysis to achieve the best signal/ $\sqrt{\text{background}}$ figure of merit. With this cut 4.36 ν_e CC signal events were selected on a background of 0.05 ν_μ CC events, 0.34 neutral currents, and 0.44 intrinsic beam ν_e CC interactions for an FOM of 4.73. Before the LID cut 9.08 ν_e CC events and 23.92 background events remained after preselection. The LID algorithm has a signal selection efficiency of 48% and rejects background at 97%.




Figure 5.28: Event classification output from the LID algorithm where values closer to one represent ν_e CC signal. The plot shows simulated Far Detector events scaled to the analysis exposure of 2.8×10^{20} full detector equivalent POT with standard preselection cuts (a) and a version zoomed into the signal region (b).

Evaluation of the LID algorithm and comparisons between data and Monte Carlo are discussed in greater detail in the context of the event selection for the Near Detector (Section 6.4) and Far Detector (Section 6.3). Systematic errors in the reconstruction and event classification due to dead channels in the detector are discussed in Section 7.3.12. One cross check on LID performance in Far Detector data was done using electromagnetic showers from cosmic ray Bremsstrahlung radiation in data and Monte Carlo, the results of which are discussed in Section 7.2.1.

5.7.6 Improvements to Event Selection

One element of the algorithm with the potential for improvement is to move beyond focusing on the most energetic shower in a slice. PID information could instead be compiled for every shower so the event picture could be built on a particle by particle basis. A second level of neural networking would then fold the shower information into a classification of the slice. This approach would work well for cases where either the electron is not the most energetic shower or when a clear muon or π^0 is present among the secondary showers but the primary shower was not properly reconstructed. This approach would have an impact beyond the oscillation analysis for cross section measurements in the Near Detector where classifying an event by more than the primary lepton is desired.

Small improvements might be found in a dedicated training for the Near Detector and an optimization of the training sample. The dE/dx training histograms for longitudinal showers have 200 planes \times 11 energies \times 4 quadrants \times 7 particles for a total of 61600 histograms which might lead to some rarer shower topologies falling in underrepresented training histograms.

CHAPTER 6

Event Selection

This chapter details the analysis period (Section 6.1), selection of good data (Section 6.2), event selection (Section 6.3 and 6.4), and the calculation of effective fiducial mass and beam exposure (Section 6.5 and 6.6). The work described here is the work of several people and in some cases the subject of additional theses. I largely developed the Near Detector event selection, performed the effective fiducial mass calculation, and calculated the exposure window.

All Far Detector event selection studies on selection optimization applied a standard set of oscillation weights to the simulation. These weights assumed no matter effect or CP violation, $\sin^2 2\theta_{23} = 1$, $\Delta m_{32}^2 = 2.35 \times 10^{-3}$, and $\sin^2 2\theta_{13} = 0.1$. These weights represent an average oscillation result so as not to bias the selection with assumptions about mass ordering and CP violation.

Cosmic rejection was optimized with an independent high statistics sample of minimum bias triggers at 10 Hz and not coincident to the beam window. The cosmic background estimate uses an independent sample taken from the NuMI beam spill triggers. The beam occupies a 10 μ s period of the 500 μ s trigger window, allowing the out-of-time slices to be used for cosmic background estimation.

Event selection criteria in both detectors was tuned on the figure of merit (FOM) quantity $\text{signal}/\sqrt{\text{background}}$. This metric optimizes the selection for the discovery of a nonzero

value of the mixing angle θ_{13} which is the goal of the analysis in this thesis. In future versions of the analysis when the priority is a precision measurement the selection will be re-optimized based on the FOM $\text{signal}/\sqrt{\text{signal} + \text{background}}$.

6.1 Analysis Period

The data used in this thesis span a time period from February 6th, 2014 until May 15th, 2015 in the Far Detector. During part of this time period the detector was under construction and expanded from four diblocks to a full 14 diblocks. From September 5th, 2014 through October 24th, 2014 the NuMI beam was shutdown for upgrades and maintenance. A time line of detector size in usable data is shown in Figure 6.1. The Near Detector began collecting data on August 18th, 2014 and due to its smaller size only operated as a complete detector.




Figure 6.1: Number of consecutive analyzable diblocks in the Far Detector during the analysis period from February 6th, 2014 until May 15th, 2015.

At the Far Detector the timing window for the NuMI beam is 218.125 to 230.125 μs from the start of the 500 μs trigger window. The beam window used in the analysis is 12 μs , adding a 1 μs buffer on each side of the measured window to allow for drift. The window is determined by fitting the timing distribution of selected neutrino candidates [116]. The out-of-time region of the trigger, a factor of 35 times larger than the beam window, measures the cosmic background. A timing diagram of the selection regions is shown in Figure 6.2.

For the Near Detector which sees low cosmic background the beam peak is visible on an event-by-event basis in the cell hit times, as discussed further in Section 6.4.




Figure 6.2: Timing diagram of the beam window and out-of-time cosmic sideband.

During the pre-shutdown data acquisition period the Far Detector was susceptible to a known but undetectable failure in the timing system. When the timing system was rebooted, a frequent occurrence during that phase of detector construction, the GPS on the Master Timing Unit would occasionally initialize shifted forwards by 64 microseconds from the true one second time boundaries. Since the accelerator shutdown a monitor is in place at the Far Detector that provides an independent 1 Hz reference signal from a different GPS antenna in order to verify the time. For the pre-shutdown data, which represents 23% of the total exposure, a second timing window was opened with a $+64 \mu\text{s}$ shift since it is unknown how much of the data is affected.

6.2 Quality Run Selection

Data taken from the detector is partitioned into runs and subruns with a typical Far Detector subrun lasting 2-3 minutes and a Near Detector subrun lasting an hour. Run numbers roll over after 64 subruns or 24 hours. A series of selection criteria were developed at the run and subrun level to assess data quality and remove files from further analysis.

6.2.1 Far Detector

Due to the size of the Far Detector data quality is assessed in each diblock so that even when a portion of the detector is not performing and is marked bad other regions can be analyzed. Run quality selection for the Far Detector require [117]:

1. **Sanity:** Does the subrun make sense?
 - subrun start time < subrun end time
 - Year \geq 2013
 - File is not empty
2. **Subrun duration:** subrun > 1 second
3. **MIP Rate:** Selection on the median MIP rate (where a MIP has an ADC value between 175 and 3200) of events in the subrun: $13 \text{ Hz} < medMIPRate < 23 \text{ Hz}$
4. **Slices:** $1.2 < slices/event/10^4 \text{ channels} < 3.2$
5. **Reconstructed tracks:** Apply a simple straight line tracker to each slice, require fraction of tracks per event that are 2D < 15%. This selection is effective when DCMs are found to be out of time sync with each other.
6. **Detector size:** A minimum of four consecutive diblocks that are **good**. A diblock is considered good if all 12 DCMs within the diblock are good. A DCM is good if at least 56 of 64 FEBs are good. An FEB is good if:
 - At least 26 of 32 pixels are good (a good pixel is defined below).
 - The MIP rate for the FEB in the subrun falls in the expected range (which depends on DCM location):
 - DCM 1-6 (top of detector): $400 \text{ Hz} < mipRateFEB < 1000 \text{ Hz}$
 - DCM 7 (side of detector): $300 \text{ Hz} < mipRateFEB < 1200 \text{ Hz}$
 - DCM 8 (side of detector): $200 \text{ Hz} < mipRateFEB < 1000 \text{ Hz}$
 - DCM 9 (side of detector): $200 \text{ Hz} < mipRateFEB < 850 \text{ Hz}$
 - DCM 10 (side of detector): $100 \text{ Hz} < mipRateFEB < 700 \text{ Hz}$
 - DCM 11 (side of detector): $100 \text{ Hz} < mipRateFEB < 600 \text{ Hz}$
 - DCM 12 (side of detector): $100 \text{ Hz} < mipRateFEB < 550 \text{ Hz}$

A pixel is defined as good if:

- $10^{0.5} \text{ Hz} < \text{hitRate} < 10^{3.5} \text{ Hz}$
- The MIP rate for the pixel in the subrun falls in the expected range (which depends on DCM location) as shown in Figure 6.3:
 - DCM 1-6 (top of detector): $13 \text{ Hz} < \text{mipRate}_{FEB} < 31 \text{ Hz}$
 - DCM 7 (side of detector): $5 \text{ Hz} < \text{mipRate}_{FEB} < 45 \text{ Hz}$
 - DCM 8 (side of detector): $4 \text{ Hz} < \text{mipRate}_{FEB} < 36 \text{ Hz}$
 - DCM 9 (side of detector): $4 \text{ Hz} < \text{mipRate}_{FEB} < 30 \text{ Hz}$
 - DCM 10 (side of detector): $2 \text{ Hz} < \text{mipRate}_{FEB} < 26 \text{ Hz}$
 - DCM 11 (side of detector): $2 \text{ Hz} < \text{mipRate}_{FEB} < 23 \text{ Hz}$
 - DCM 12 (side of detector): $1.5 \text{ Hz} < \text{mipRate}_{FEB} < 20 \text{ Hz}$




Figure 6.3: Distributions of MIP rate for APD pixels in Far Detector data as a function of DCM location. Figure from [117].

In the analysis period 158189 good subruns were found out of a total of 188600 subruns for an efficiency of 84%. Of the failed runs $\sim 80\%$ had less than four usable diblocks while construction was in progress. Since January 2015 the subrun failure rate has been $\sim 0.3\%$. In this period, routine detector maintenance has decreased and the detector is operating stably with 95% to 98% uptime when the NuMI beam is on.

6.2.2 Near Detector

The Near Detector sees a high rate of neutrinos, 3-4 interactions per NuMI spill, and can thus be more selective in data quality. As the detector is small, all DCMs and diblocks must be marked as good to use a subrun. Run quality selection for the Near Detector requires [118, 119]:

1. **Subrun duration:** $\text{subrun} > 1000$ spills
2. **MIP Rate:** Selection on the median MIP rate (where a MIP has an ADC value between 175 and 3200) per spill (scaled to 2.5×10^{13} POT equivalent): $12 \text{ Hz} < \text{medMIPRate} < 20 \text{ Hz}$
3. **Empty spills:** Fraction of empty spills in subrun $< 3\%$
4. **Timing peak:** The $10 \mu\text{s}$ beam window is visible in the structure of the average slice times in a subrun. The leading and falling edge of the peak are fit (Figure 6.4) and requirements are placed on the start and end times of the peak:
 - $217.25 \mu\text{s} < \text{peakStartTime} < 219.25 \mu\text{s}$
 - $227.25 \mu\text{s} < \text{peakStopTime} < 229.25 \mu\text{s}$
5. **Slice rate:** Number of slices per trigger (scaled to 2.5×10^{13} POT equivalent): $3.5 \text{ Hz} < \text{sliceRate} < 5.5 \text{ Hz}$
6. **Detector size:** All four diblocks are required to be **good** for a subrun to be selected. A diblock is good if all DCMs in it are good. A DCM is good if at least 80% of the FEBs within are good (not all DCMs are fully instrumented so the number of FEBs varies). An FEB is good if at least 26 of 32 pixels are good. A pixel is good if the hit rate (scaled to 2.5×10^{20} POT equivalent) falls in a range: $10^{0.5} \text{ Hz} < \text{pixelRate} < 10^{3.5} \text{ Hz}$.

In the analysis period 3097 good subruns were found out of a total of 3533 subruns for an efficiency of 88%. Since January 2015 the subrun failure rate has been $\sim 6\%$ with most failures coming from short runs that occur during periods of testing.




Figure 6.4: Mean slice times for Near Detector Data fit for the rise and fall of the neutrino beam peak. The tail on falling edge is due to delayed neutron shine in the detector from neutrino interactions in the surrounding rock. Figure from [119].

6.3 Far Detector Event Selection

For NuMI data events additional requirements are placed on the position and spread of the beam on the target, current in the magnetic focusing horns, and a spill must contain at least 2×10^{12} POT for an event to be accepted (in normal operations a spill typically has $2.5 - 5 \times 10^{13}$ POT) [120]. Every DCM is required to have reported to the DAQ to remove cases where a DCM fails to send information when an event is built [121]. In each event at least 20% of track segments touching DCM boundaries are required to have a matching edge in the adjacent DCM to check for time synchronization failures [122].

In the Far Detector, the event selection is optimized to identify ν_e CC signal events in the 1 to 3 GeV region of maximum oscillation and reject cosmic ray and NuMI beam backgrounds. Cosmic ray backgrounds are largely rejected through containment and measuring the momentum of the slice with respect to the beam direction. The final selection places a requirement on LID to reject candidates not consistent with a ν_e CC topology. For more detail on the selection see [123]. The following selection stages are listed in order of application and are applied to each slice:

1. **Data Quality:**

- **Hits in slice, vertical view:** $nHitX > 5$
 - **Hits in slice, horizontal view:** $nHitY > 5$
 - **Flasher:** $\frac{\text{hits in slice}}{\text{planes in most energetic shower}} < 8$, High energy deposits in a cell can produce a ringing effect depositing energy on the other pixels in the APD with a time structure of 30 microseconds. These events which light up an FEB are called “flashers”, the delayed time structure often leads the afterglow hits from the initial energy deposit to be sliced separately from the track and can appear contained.
 - **Reconstruction:** Minimum one 3D reconstructed shower with LID in the slice
 - **Gap:** $gap < 100$ cm. A reconstruction quality check on the gap between the start of the most energetic shower in the slice and the reconstructed vertex.
 - **Colinearity:** $shw1.Dir \cdot shw2.Dir < -0.95$, Remove cases where the two most energetic showers are back-to-back indicating the reconstructed vertex was pulled forward along the shower direction from the true vertex.
 - **Shower hit asymmetry:** $\frac{|nHitShw1X - nHitShw1Y|}{nHitShw1X + nHitShw1Y} < 0.4$, Check the asymmetry of hits between the two readout views in the most energetic shower.
 - **Shower Reconstruction Fraction:** $\frac{\sum_i nHitShw_i}{nHitSlice} > 0.7$
2. **Containment:** Containment is defined based on the minimum projected distance from the endpoints of the most energetic 3D fuzzy-k prong in a slice to the active detector volume boundaries:
- **East wall:** $minDist \geq 15$ cm
 - **West wall:** $minDist \geq 10$ cm
 - **Top wall:** $minDist \geq 150$ cm
 - **Bottom wall:** $minDist \geq 10$ cm
 - **South wall:** $minDist \geq 35$ cm
 - **North wall:** $minDist \geq 200$ cm

The containment requirement on the north wall is tighter to remove photons entering the detector through the back due to holes in the overburden for the loading dock. The requirement on the top of the detector is also tighter to reduce neutron activity.

3. Cosmic Rejection:

- **Transverse momentum:** Select on the ratio of transverse momentum with respect to the beam direction for the reconstructed 3D showers to the total momentum. This requirement has two settings depending on the minimum distance of all 3D shower endpoints to the top of the detector to be more aggressive in the removal of cosmic background near the boundary:

- if ($minDistTop < 25$ cm) $\frac{P_{trans}}{P} \leq 0.4$
- else if ($minDistTop \geq 25$ cm) $\frac{P_{trans}}{P} \leq 0.65$

4. Preselection:

- **Slice Calorimetric Energy:** $1.5 \text{ GeV} < sliceEnergy < 2.7 \text{ GeV}$
- **Slice hits:** $40 < nHitSlice < 115$
- **Prong Length:** $140 \text{ cm} < prongLength < 500 \text{ cm}$ (for longest 3D Fuzzy-k prong in the slice)

5. Event Classification: $LID > 0.95$

The final performance evaluation of the event selection is shown in Figure 6.5 and Table 6.1. The selection rejects cosmic background events at a rate of 100 million to 1, while maintaining 21% of the ν_e CC signal (for signal events passing preselection 48% survive the LID selection). NuMI beam backgrounds are rejected at 99.8%. In the background selected events 92% of the NC events and 86% of the ν_μ CC events contain a π^0 in the final state. Improving π^0 reconstruction and identification can lead to significant background reduction in the next analysis.

6.4 Near Detector Event Selection

In the Near Detector events are selected that are representative of beam neutrino backgrounds expected in the Far Detector. These selected events will be used in an extrapolation technique, see Section 7.1, to predict Far Detector backgrounds and reduce systematic errors. In the Near Detector the cosmic background rate is about 15 Hz and studies of




Figure 6.5: LID spectrum of un-extrapolated Far Detector Monte Carlo scaled to 3.52×10^{20} POT. Also shown is cosmic background data from minimum bias trigger data, scaled to 165 live seconds. The selection at an LID value greater than 0.95 is drawn, which maximizes the FOM. Figure from [123].

	ν_e CC signal	Cosmic Data	total background MC	FOM
No Selection	20.63	5.98×10^6	540.40	0.0
Data Quality	15.20	3.69×10^5	102.82	0.0
Containment	12.75	1.60×10^4	83.61	0.1
Cosmic Rejection	11.97	319.33	72.17	0.6
Preselection	9.08	73.07	23.92	1.3
Event Classification	4.37	0.04	0.83	4.6

Table 6.1: Far Detector event counts and Figure of Merit (signal/ $\sqrt{\text{background}}$) for out-of-time NuMI cosmic background scaled to 165 liveseconds and NuMI Monte Carlo scaled to 2.8×10^{20} POT. Table from [123].

the out-of-time data have shown that after selection zero cosmics are expected in the sample [124]. The same beam spill quality criteria are applied in the Near Detector as the Far Detector and no missing DCMs are allowed in an event. Additionally, less than 45% of the hits in an event can be in the top horizontal row of DCMs to remove false signals caused when the lights are on in the underground detector hall. Due to the smaller size of the Near Detector tight containment is required to get well contained events and also remove events where the neutrino interaction occurs in the rock outside the detector. The final event selection is:

1. **Reconstruction:** Minimum one 3D reconstructed shower in the slice
2. **Flasher:** $\frac{\text{hits in slice}}{\text{planes in most energetic shower}} < 8$, same as the Far Detector selection
3. **Fiducial:** Apply fiducial containment based on position of *ElasticArms* vertex. In z the selection was designed to be two photon conversion lengths from the front face to reduce contamination from neutrino interactions in the rock and at least 5 meters from the start of the muon catcher to avoid shower topology differences between the detectors. In x and y containment was placed one photon conversion length from the wall which is all that was possible without a significant reduction in statistics.
 - $-140 \text{ cm} < x < 140 \text{ cm}$
 - $-140 \text{ cm} < y < 140 \text{ cm}$
 - $100 \text{ cm} < z < 700 \text{ cm}$
4. **Shower Containment:** Selection on the endpoint of all 3D reconstructed showers in the slice
 - **East wall:** $\min(\sum_i \text{shwStartX}_i, \sum_i \text{shwStopX}_i) \geq -180\text{cm}$
 - **West wall:** $\min(\sum_i \text{shwStartX}_i, \sum_i \text{shwStopX}_i) \leq 180\text{cm}$
 - **Bottom wall:** $\min(\sum_i \text{shwStartY}_i, \sum_i \text{shwStopY}_i) \geq -180\text{cm}$
 - **Top wall:** $\min(\sum_i \text{shwStartY}_i, \sum_i \text{shwStopY}_i) \leq 180\text{cm}$
 - **South wall:** $\min(\sum_i \text{shwStartZ}_i, \sum_i \text{shwStopZ}_i) \geq 25\text{cm}$
 - **North wall:** $\min(\sum_i \text{shwStartZ}_i, \sum_i \text{shwStopZ}_i) \leq 1225\text{cm}$

The fiducial and shower containment cuts are illustrated in Figure 6.6.

5. **Front Planes:** No hits in a slice in first 6 planes of the detector to remove rock events.
6. **Slice hits and Energy:** $20 < nHitSlice < 200$. The compact Near Detector results in slices with a smaller energy resolution for a given number of hits than the Far Detector, shown in Figure 6.7. In order to sufficiently sample the same energy range the selection is placed looser in the Near Detector. Additionally, a loose selection on the slice calorimetric energy was placed, requiring less than 5 GeV.
7. **Prong Length:** $140 \text{ cm} < prongLength < 500 \text{ cm}$ (for longest 3D Fuzzy-k prong in the slice), same as Far Detector.
8. **Gap:** $gap < 100 \text{ cm}$. A reconstruction quality check on the gap between the start of the most energetic shower in the slice and the reconstructed vertex, same as Far Detector.




Figure 6.6: Drawing of the fiducial (red) and shower containment boundaries (green) inside the Near Detector (black). Figure from [125].

One method of verifying that the Near Detector event selection contains a representative sample of Far Detector backgrounds is to compare the distributions of invariant kinematic variables in simulated neutrino interactions at both detectors. An illustration of lepton-nucleon scatterings is shown in Figure 6.8 to define the kinematic variables in the interaction. Two useful variables for comparing the kinematics of near and far detector selected events are the four-momentum transfer, Q^2 , and the mass of the system recoiling against




Figure 6.7: Slice calorimetric energy versus the number of hits in the slice for both detectors.

the scattered lepton, W^2 . Q^2 is defined as:

$$Q^2 = -q^2 = 2(E E' - \vec{k} \cdot \vec{k}') - m_l^2 - m_l'^2 \quad (6.1)$$

where E and E' are the initial and final lepton energies, k and k' are the initial and final lepton four-momenta, and m_l and m_l' are the initial and final lepton masses. W^2 is defined as:

$$W^2 = (P + q)^2 = M^2 + 2M(E - E') - Q^2 \quad (6.2)$$

where P is the four-momentum of the nucleon with mass M , and $q = k - k'$ is the four-momentum transfer to the nucleon.

The comparison of kinematic variables between the detectors for selected events is shown in Figure 6.9. The high energy tail comes from neutral current interactions that feed down into the selection region due to a large portion of the energy leaving with the neutrino. The comparison illustrates that the Near Detector selection contains a sample representative of the Far Detector expectation. The spike in the plots at $W^2 = 1$ is from coherent single pion neutrino-nucleus scattering.




Figure 6.8: Kinematic quantities in lepton-nucleon scattering where k and k' are the four-momenta of the incoming and outgoing leptons, P is the four-momentum of the nucleon with mass M , W is the mass of the hadronic final state, and $q = k - k'$ is the four-momentum transfer to the nucleon. Figure from [28].

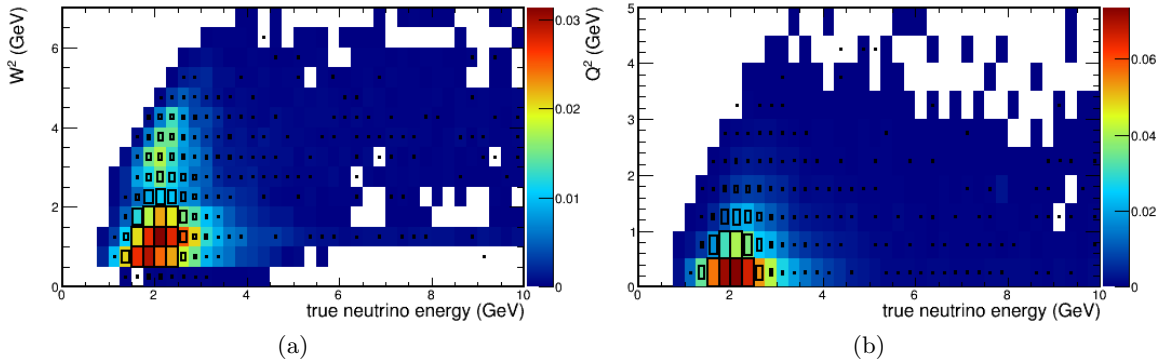


Figure 6.9: In color are the kinematic variables W^2 (a) and Q^2 (b) plotted against the true neutrino energy for events passing selection in Near Detector Monte Carlo. Overlaid with boxes are Far Detector Monte Carlo events with standard oscillation weights passing selection. Both selections are area normalized to one.

In the analysis period the Near Detector accumulated 1.66×10^{20} POT of good data. Table 6.2 shows the effect of the selection on Near Detector data and Monte Carlo. The final LID distribution after selection is shown in Figure 6.10. After selection, data and Monte Carlo agree to within 5% which is acceptable given the intrinsic uncertainty in the beam flux (Section 7.3.1). The LID distribution shows a distinct peak of beam ν_e CC events which represent “signal” in the Near Detector. The energy distribution of selected slices with LID > 0.95 is shown in Figure 6.11. There is a consistent 10% deficit of data relative to the Monte Carlo in LID bins below 0.95 with the exception of the bin at 0. This behavior is also seen in an independent electron classification algorithm [126] used in a parallel analysis branch that achieves similar performance. This discrepancy is believed to come from shifts in the LID distribution due to the modeling of hadronic energy in the backgrounds [127], [128]. In the selection region data and Monte Carlo have strong agreement as illustrated in Figure 6.11. Additional studies of data and Monte Carlo agreement in the Near Detector are documented in [125].




Figure 6.10: LID distribution of Near Detector data and Monte Carlo passing the selection criteria on a log scale (a) and zoomed in near the signal region (b). Plots are normalized to 1.66×10^{20} POT in the data.

	Total MC	ν_μ CC MC	ν_e CC MC	NC MC	Data
No Selection	30049057	25742957	364671	3941430	29802297
Flasher	29139394	25132969	339183	3667242	29015588
Reconstruction	16338569	13577267	229803	2531499	15996941
Fiducial	1139793	758778	20924	360091	1025407
Shower Containment	478835	253650	10236	214949	424518
Front Planes	461975	243569	9815	208592	413242
Slice Hits and Energy	323131	177970	5437	139724	301019
Prong Length	236834	145944	4859	86030	222046
Gap	230475	142987	4755	82732	217656
LID	2471	396	1292	783	2579

Table 6.2: Near Detector data (1.65×10^{20} POT) and Monte Carlo (normalized to data) event counts passing the levels of preselection.




Figure 6.11: Calorimetric energy of slices in Near Detector data and Monte Carlo passing selection and $LID > 0.95$.

6.5 Far Detector Effective Fiducial Mass

Due to the changing size of the Far Detector during construction fiducial mass is calculated as a function of detector size. The preselection described in Section 6.3 do not have a hard containment selection and so an effective fiducial mass is estimated by measuring the efficiency of the event selection in Monte Carlo as a function of detector size. In this work efficiency is defined as:

$$\epsilon = \frac{\text{selected } \nu_e \text{ CC signal events in detector volume}}{\text{all } \nu_e \text{ CC signal events in detector volume}}, \quad (6.3)$$

where an event is determined to be in a detector volume based on the true neutrino vertex. An event is selected based in the standard preselection. The total selection efficiency is the product of two factors:

$$\epsilon = \epsilon_0 \times \epsilon_f, \quad (6.4)$$

where ϵ_0 is the efficiency for selecting an event in a detector of infinite size, and ϵ_f is the efficiency due to the effective fiducial volume the selection criteria impose on a physical detector.

To approximate the efficiency in a detector of infinite size a narrow box was defined with $-100 < x < 100$ cm, $-100 < y < 100$ cm, and $1000 < z < 3500$ cm. In this narrow region, far from the detector edges, the selection efficiency of the selection was found to be $\epsilon_s = 0.660 \pm 0.021$. Using the MC generated with run masks from the pre-shutdown period the efficiency could be measured in a range from 4 to 14 diblocks using the dataset.




Figure 6.12: The efficiency of the Far Detector ν_e CC selection as a function of detector size with binomial errors shown. The fit is of the form $y = a + \frac{b}{x^{0.993}}$ with $a = 0.509 \pm 0.005$ and $b = -0.557 \pm 0.048$.

The resulting selection efficiency is shown in Figure 6.12. The Far Detector has a total mass of 13.902 kilotons or 0.993 kilotons per diblock [76]. The efficiency as a function of detector size can be fit to an asymptotic model of the form:

$$\epsilon = a + \frac{b}{n_{db} \times 0.993 \text{ kT}}. \quad (6.5)$$

The total efficiency measured from the run masked MC can be used in combination with

the selection efficiency measured in a small volume to extract the fiducial efficiency, which can then be converted into a fiducial mass as shown in Table 6.3. A formula for the fiducial mass, making use of Equation 6.4 and 6.5 is:

$$m_f = \frac{n_{db}\epsilon}{\epsilon_0} \times 0.993 \text{ kT.} \quad (6.6)$$

After substituting the appropriate values the result is:

$$m_f = 0.776 (n_{db} - 1.088) \quad (6.7)$$

Total mass (kT)	effective fiducial mass (kT)
3.972	2.3 ± 0.1
4.965	3.0 ± 0.1
5.958	3.8 ± 0.1
6.951	4.6 ± 0.2
7.944	5.4 ± 0.2
8.937	6.1 ± 0.2
9.930	6.9 ± 0.2
10.923	7.7 ± 0.3
11.916	8.5 ± 0.3
12.909	9.2 ± 0.4
13.902	10.0 ± 0.4

Table 6.3: Effective fiducial mass in kilotons as a function of detector size derived from selection efficiencies in the Monte Carlo.

Another way to fit the the efficiency curve in Figure 6.12 is to recast the containment efficiency in terms of effective distances from the walls,

$$\epsilon_f = \frac{(w - 2d)^2 (L - f - b)}{w^2 L} \quad (6.8)$$

where L is detector length, w detector height/width, d distance from x/y edge, f distance from front face, and b distance from back. Expanding and keeping only first order in d , f ,

and b gives

$$\epsilon \simeq \epsilon_0 \left(1 - \frac{4d}{w} - \frac{b+f}{L} \right). \quad (6.9)$$

Using $w = 1525$ cm and an L of 426 cm per diblock, and fixing $\epsilon_0 = 0.660$ for LID style containment, the data in Figure 6.12 can be fit to a solution $d \approx 88$ cm and $f + b \approx 362$ cm. This fit function reduces to an equivalent effective fiducial mass expression to Equation 6.7 and provides a useful cross check. This formula also provides physical meaning to the preselection, providing an effective containment distance.

6.6 Analysis Exposure

The final Far Detector POT and exposure accounting for the analysis period with NuMI beam quality and good runs selection applied is shown in Table 6.4. In total 3.52×10^{20} POT was collected and after factoring in the changing detector size is the equivalent of 2.8×10^{20} POT in a full detector which is $\sim 1/2$ of a nominal NOvA year when the beam reaches 700 kW. Before the accelerator shutdown the average effective fiducial mass was half of a full detector, while post-shutdown the detector has been 98% operational on average.

The accumulated exposure and detector mass are shown in Figure 6.13.

	Preshutdown	Postshutdown	Total
POT($\times 10^{20}$)	1.27	2.25	3.52
Exposure (POT \times kT $\times 10^{20}$)	6.3	21.6	27.9
Equivalent Exposure (POT \times 11.4kT $\times 10^{20}$)	0.6	2.2	2.8
Avg. Fiducial Mass (kT)	5.0	9.6	7.9
Livetime (s)	131.9	102.5	234.3
Livetime Exposure (s \times kT)	650	980	1631
Equivalent Livetime (s \times 11.4kT)	65.0	98.0	163

Table 6.4: POT, fiducial mass and livetime accounting for First Analysis period. POT exposure and livetime are also expressed in terms of a full 10.0 fiducial kiloton equivalent.

Monte Carlo is generated in the different diblock configurations proportionally to the POT accumulated in the good data runs. This generation scheme folds the varying effective




Figure 6.13: The cumulative POT exposure and fiducial mass for the analysis period. The exposure is $27.9 \text{ POT} \times kT \times 10^{20}$ which is the equivalent of 2.8×10^{20} POT for a full 10.0 kT fiducial volume detector.

fiducial volume into the Monte Carlo and allows the POT to be scaled to the data while automatically folding in varying detector configurations. The Monte Carlo generation is checked in Figure 6.14 by comparing the POT accumulated in data to the Monte Carlo for each detector configuration. This cross-check shows that the run matched Monte Carlo is an accurate representation of the data taken.




Figure 6.14: POT accumulated in each detector configuration for pre and post shutdown data (black) and real-conditions Monte Carlo (red). The Monte Carlo is area normalized to the data. Beam and data quality cuts are applied at the spill level to both data and Monte Carlo.

CHAPTER 7

Electron Neutrino Appearance Analysis

The ν_e analysis counts the number of events selected in the Far Detector in the energy window described in Section 6.3 (1.5 to 2.7 GeV). Before fitting the data for oscillation parameters, as described in Chapter 8, a prediction of the signal, background and uncertainty are made. This analysis is statistics limited, however evaluations of a number of systematic effects on the signal and background predictions are made. The comparison of Near and Far Detectors reduces, and in some cases cancels, many of the sources of systematic uncertainty. The prediction of signal and background event counts uses the Near Detector data to correct the Monte Carlo simulation and then uses the Monte Carlo to extrapolate the result to the Far Detector, described in Section 7.1. Section 7.2 shows analysis samples with modified selection criteria used to verify the performance of the event selection and LID, referred to as “sidebands”. Section 7.3 shows an evaluation of the systematic uncertainties effect on the analysis. This work was an ensemble effort with many topics the subject of their own theses. I contributed studies on the systematic uncertainties related to extrapolation and performed the high energy and low-LID sideband analyses.

7.1 Extrapolation

In the two detector design of the NOvA experiment the Near Detector data is used to inform and adjust the signal and background prediction in the Far Detector. Due to the nearly

identical construction of the detectors, systematic errors are reduced by using the Near Detector data as input to the prediction. The process of selecting Near Detector events (ν_μ CC which will oscillate into ν_e CC signal and backgrounds to the ν_e CC measurement) in data and Monte Carlo and translating that with oscillations into a Far Detector prediction is called extrapolation, which is more fully described in [129].

The extrapolation predicts the Far Detector event rate, $F_{S_e}^{Pred}(B_j^e)$; *Pred* denotes that the rate comes from extrapolation, S_e means that the ν_e event selection was applied, and B_j^e enumerate the bins of reconstructed energy (0.25 GeV in width) used in the ν_e analysis indexed by j . The extrapolation makes separate predictions for the signal channel, $\nu_\mu \rightarrow \nu_e$, and each of three major backgrounds: $\nu_\mu \rightarrow \nu_\mu$, $\nu_e \rightarrow \nu_e$, and neutral current.

For the signal channel ($\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$) events are selected in Near Detector data and Monte Carlo that pass the ν_μ style preselection [130] and event classification [131]. The data/MC ratio in each reconstructed energy bin is used to adjust the Near Detector prediction. The Monte Carlo is used to form a matrix relating reconstructed energies to true energies. The Near Detector ν_μ CC prediction, $N_{\nu_\mu}^{Pred}$, takes the form:

$$N_{\nu_\mu, S_\mu}^{Pred}(E_i^T) = \sum_k \frac{N_{\nu_\mu, S_\mu}^{Data}(B_k^\mu) N_{\nu_\mu, S_\mu}^{MC}(E_i^T, B_k^\mu)}{N_{\nu_\mu, S_\mu}^{MC}(B_k^\mu)} \quad (7.1)$$

where S_μ is the ν_μ style event selection, B_k^μ is reconstructed bins of the charged current inclusive (quasi and non-quasi elastic) ν_μ energy estimator [132] (0.1 GeV), E_i^T are energy bins in true neutrino energy, N^{Data} and N^{MC} are the rate of Near Detector events in data and Monte Carlo respectively. The ν_μ selection is applied in the Near Detector and ν_e selection in the Far Detector since the Near Detector ν_μ 's oscillate into the appearance signal. For this oscillation channel, the conversion from reconstructed to true energy is necessary because in the Far Detector the ν_e selection and energy estimator are used and these have different energy resolution. The Far Detector prediction of ν_e CC signal in true

energy bins is:

$$F_{\nu_\mu \rightarrow \nu_e, S_e}^{Pred}(E_i^T, B_j^e) = \frac{N_{\nu_\mu, S_\mu}^{Pred}(E_i^T) F_{\nu_\mu \rightarrow \nu_e, S_e}^{MC}(E_i^T, B_j^e)}{N_{\nu_\mu, S_\mu}^{MC}(E_i^T)}. \quad (7.2)$$

The Far Detector prediction is multiplied by the oscillation probability, $P_{\nu_\mu \rightarrow \nu_e}(E_i^T)$, and converted back to bins of reconstructed energy for the ν_e energy estimator (total calorimetric energy of the slice):

$$F_{S_e}^{Pred}(B_j^e) = \sum_i F_{\nu_\mu \rightarrow \nu_e, S_e}^{Pred}(E_i^T, B_j^e) P_{\nu_\mu \rightarrow \nu_e}(E_i^T) \quad (7.3)$$

The major background channels have the ν_e selection applied in both detectors so the rate of mis-identified events in the Near Detector predicts the mis-identification in the Far Detector. Since these events are mis-identified there is no reason to expect the energy estimator or other reconstructed quantities to perform well, so the event rates are reweighted in terms of reconstructed and not true energy. The $\nu_e \rightarrow \nu_e$ channel could be extrapolated in the same manor as the signal channel, but low statistics cause the reconstructed to true energy conversion matrix to be sparsely populated. The data in the Near Detector is “proportionally decomposed” into the three background components (ν_e CC, ν_μ CC, and NC) weighted according to the Monte Carlo in each energy bin. For example, if each background component comprised a third of the events in the energy bin from 1.5 to 1.75 GeV, then the selected data events will be split by the same ratio. For these background channels the re-weighting formula is:

$$F_{\alpha \rightarrow \alpha, S_e}(E_i^T, B_j^e) = \frac{N_{\alpha, S_e}^{Data}(B_j^e) F_{\alpha \rightarrow \alpha, S_e}^{MC}(E_i^T, B_j^e)}{N_{\alpha, S_e}^{MC}(B_j^e)} \quad (7.4)$$

where $\alpha \rightarrow \alpha$ denotes the background oscillation channels ($\nu_e \rightarrow \nu_e$, $\nu_\mu \rightarrow \nu_\mu$, neutral

current). After applying oscillation weights the prediction is:

$$F_{S_e, \alpha \rightarrow \alpha}^{Pred} (B_j^e) = \sum_i F_{\alpha \rightarrow \alpha, B_e}^{Pred} (E_i^T, B_j^e) P_{\alpha \rightarrow \alpha} (E_i^T). \quad (7.5)$$

There are eight remaining minor background channels: $\nu_e \rightarrow \nu_\mu$, $\nu_e \rightarrow \nu_\tau$, $\nu_\mu \rightarrow \nu_\tau$, $\bar{\nu}_e \rightarrow \bar{\nu}_e$, $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$, $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$. These background rates are small (sub percent level) and energy reweighting is not used; the Far Detector prediction is taken directly from the Monte Carlo without modification.

The extrapolated Far Detector signal and background prediction uses 1.66×10^{20} POT of Near Detector data and 9.92×10^{20} POT of Monte Carlo scaled to match the data (selection event counts are discussed in Section 6.4). In the Far Detector 8.6×10^{23} POT of run-matched Monte Carlo is used, with the result scaled to 3.52×10^{20} POT to match the exposure in the analysis period (Section 6.6). The extrapolation method decomposes the Near Detector data proportionally to the Monte Carlo backgrounds with an uncertainty assessed in Section 7.3.8. The prediction, prior to systematic errors and without cosmic backgrounds, is 4.41 oscillated ν_e CC events on a background of 0.97 neutrinos (Table 7.1). The extrapolation predicts slightly more signal and background events than the Far Detector Monte Carlo due to the slight excess of data events with respect to the Monte Carlo seen in the Near Detector. Figures 7.1 and 7.2 show the steps used for the signal and background channels to produce the prediction. The reconstructed energy of ν_μ CC selected events in the Near Detector with the ν_μ energy estimator has a shift between the data and Monte Carlo. This shift has been attributed to a discrepancy in the hadronic energy between simulated and data events [112]. Simulated events produce more hadronic hit activity, while the energy per hit, muon energy, and muon length show excellent agreement between data and Monte Carlo. This effect is now handled in the ν_μ CC energy estimator by scaling up the hadronic energy component in data events by 21%, which translates the a 6% scaling

in neutrino energy. In practice the ν_e signal extrapolation also uses this shift, although the final change to the predicted number of ν_e CC signal events is 0.5%. In future analyses where the energy spectrum will be fit instead of having a counting experiment the shift will be more important.

	signal ν_e CC	total bkg.	ν_μ CC	beam ν_e CC	NC	ν_τ CC
Extrapolated prediction	4.41	0.91	0.05	0.47	0.36	0.02
Far Detector MC	4.36	0.85	0.05	0.44	0.34	0.02

Table 7.1: Extrapolated prediction of Far Detector event counts normalized to 3.52×10^{20} POT and the standard Far Detector Monte Carlo for comparison.

7.2 Sideband Studies

7.2.1 Cosmic Bremsstrahlung Efficiency

The performance of the shower reconstruction and LID was explored using electromagnetic showers induced from cosmic muons (or muon neutrino interactions in the rock of the Near Detector). Electron showers from bremsstrahlung radiation or decay-in-flight muons are identified and the muon energy is removed cell-by-cell from the slice to leave behind the shower. Standard reconstruction and particle identification is run on the the shower remnant. The procedure is performed in data and Monte Carlo at both detectors. The comparison of LID performance in both detectors (Figure 7.3) shows strong agreement. The efficiency of the electron identification was evaluated as a function of vertex location and shows agreement at the 5% level between data and Monte Carlo and a uniform efficiency over time during the analysis period. These studies confirm that the analysis tools are operating as anticipated on real electron data. Further details on the analysis can be found in [133].




Figure 7.1: Steps to extrapolate prediction for the $\nu_\mu \rightarrow \nu_e$ signal channel. Starting with the data/MC ratio in the Near Detector (a), using a matrix (b) to convert to a true energy spectrum in the Near Detector (c), re-weighting the Far Detector Monte Carlo in true energy (d), and then applying a matrix (e) to convert the prediction back to units of reconstructed energy (f).

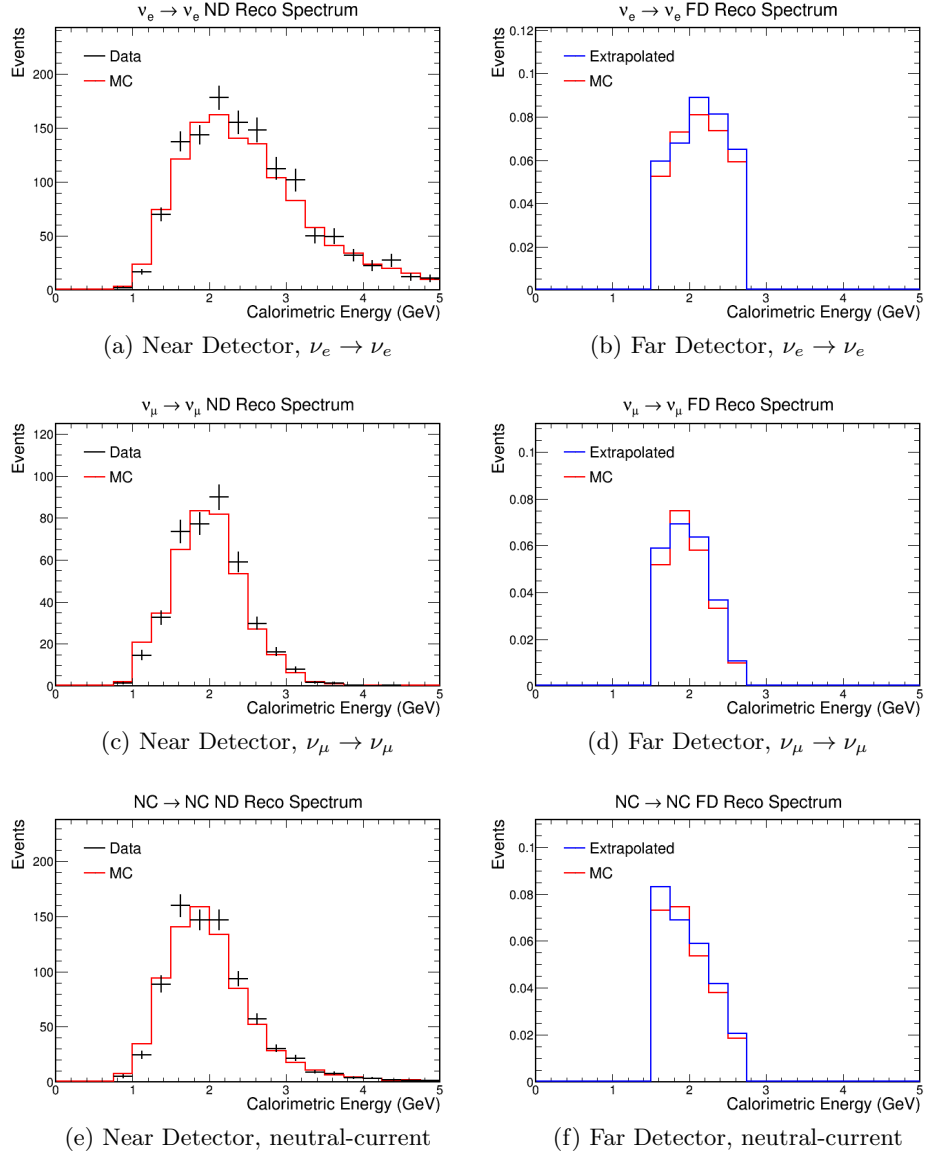


Figure 7.2: Extrapolated prediction for the major background channels: $\nu_e \rightarrow \nu_e$ (a,b), $\nu_\mu \rightarrow \nu_\mu$ (c,d), and neutral-current (e,f). Left panels show the data and Monte Carlo spectra in the Near Detector with the ν_e selection, right panels show the Far Detector Monte Carlo and extrapolated prediction.




Figure 7.3: Data and Monte Carlo comparisons of the LID distribution for cosmic muon induced electromagnetic showers in the Far Detector (a) and Near Detector (b). In the Far Detector the distribution is also compared to ν_e CC signal events. Figure from [133].

7.2.2 High Energy Neutrino Selection

The preselection and LID performance was compared in the Far Detector using the 3 to 10 GeV “high energy” side-band before un-blinding the signal region (1.5 to 2.7 GeV). All standard preselection cuts are applied in the Far Detector with the exception of removing the upper limit cut on the number of cells in the slice. In the Near Detector the upper limit on the number of cells in a slice is also removed and the extrapolation is performed between 0 and 10 GeV. Statistics are low for the extrapolation both due to the neutrino flux in the high energy region and since LID is trained on events below 5 GeV. The cosmic background in the sample was predicted using cosmic data out-of-time with the NuMI beam. The neutrino rate was predicted both from the Far Detector Monte Carlo and from an extrapolation based on Near Detector data. The predictions, in the top two panels of Table 7.2, show 32.71 cosmic rays and 22.68 neutrinos from the extrapolation (20.75 based on the un-corrected Far Detector Monte Carlo) before a LID cut.

In the beam window, 57 events are found that pass all preselection criteria (bottom panel of Table 7.2). These events were visually scanned and classified as a neutrino or

cosmic ray based on topology, detector location, and orientation. After the scan 27 events were classified as a neutrinos and 30 as a cosmic rays. The accumulation of neutrino events is overlaid with the analysis exposure in Figure 7.4, giving a Kolmogorov-Smirnov p value of 0.78 which indicates the neutrino rate is consistent with the exposure. In Figure 7.5 the data LID and energy distributions are overlaid with the prediction. Performing a Kolmogorov-Smirnov test on the LID distribution gives a p value of 0.84; the observed high energy data is consistent with the expectation. The two data events with $LID > 0.7$ are shown in Figure 7.6. The top panel event appears consistent with an electron topology while the bottom panel appears to be two co-linear photons from a π^0 decay. LID is not expected to perform at high efficiency in this energy region, still the data shows that both the preselection cuts and LID are performing in line with expectations.

	ν_e CC MC	neutrino bkg. MC	total neutrino MC	cosmic data	total
no LID cut	4.29	16.46	20.75	32.71	53.46
LID > 0.7	2.06	0.61	2.67	0.24	2.91
LID > 0.95	0.58	0.07	0.65	0.15	0.80

(a) Far Detector Monte Carlo Prediction

	ν_e CC MC	neutrino bkg. MC	total neutrino MC	cosmic data	total
no LID cut	4.57	18.11	22.68	32.71	55.39
LID > 0.7	2.15	0.67	2.82	0.24	3.06
LID > 0.95	0.59	0.08	0.67	0.15	0.82

(b) Extrapolated prediction

	Neutrino-like	Cosmic-like	total
no LID cut	27	30	57
LID > 0.7	2	0	2
LID > 0.95	0	0	0

(c) Far Detector Data Result

Table 7.2: Prediction of Far Detector neutrino background and out-of time cosmic data for the 3 to 10 GeV sideband for different LID cuts. Panel A shows the prediction based on Far Detector Monte Carlo, while panel B shows the prediction based on the extrapolation of Near Detector data. In both cases the prediction is normalized to 3.52×10^{20} POT. Panel C shows the measured result in data, where the 57 candidate events were visually hand-scanned and classified as appearing like a neutrino or cosmic ray.




Figure 7.4: The cumulative Far Detector exposure overlaid with the accumulation of 27 candidate events in the 3 to 10 GeV sideband passing selection cuts and visually identified as consistent with a neutrino. Applying a Kolmogorov-Smirnov test to the distribution gives a p value of 0.78

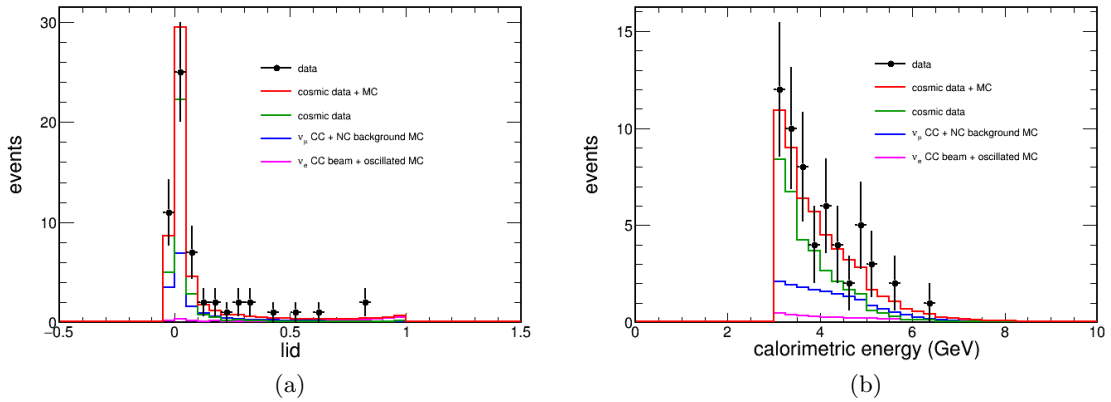

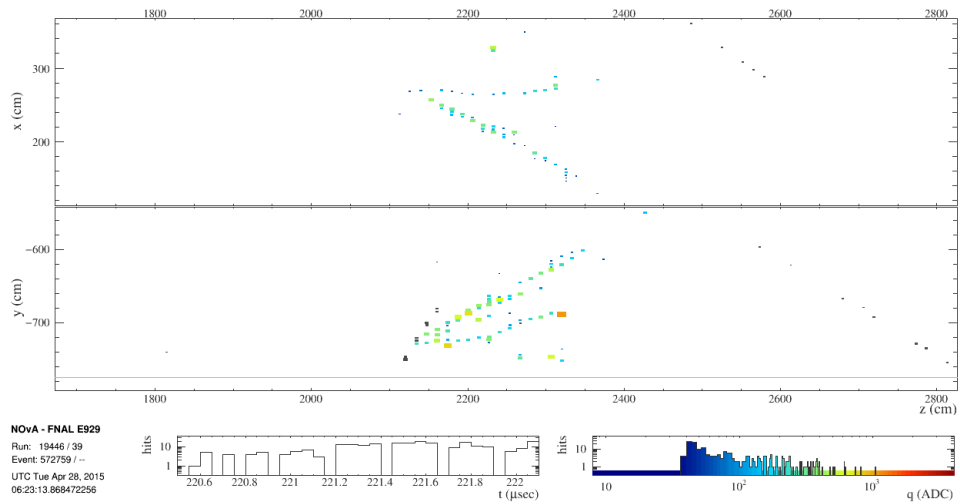


Figure 7.5: Far Detector LID spectrum (a) and calorimetric energy distribution (b) of high energy sideband events selected in data overlaid on the Monte Carlo neutrino prediction and out-of-time cosmic background. Applying a Kolmogorov-Smirnov test to the LID distribution gives a p value of 0.84.



(a)



(b)

Figure 7.6: Event displays for the two high energy neutrino events selected with an LID value greater than 0.7. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time.

7.2.3 Low LID Sideband

An additional sideband was opened in the selection energy window (1.5 to 2.7 GeV) for LID values between 0.70 and 0.95. The standard extrapolation was performed in this LID window and the cosmic background was predicted from out-of-time events in the NuMI data stream. The prediction was made assuming the extreme cases of normal mass ordering and $\delta = 3\pi/2$ and the inverted mass ordering with $\delta = \pi/2$ to bracket the expected event counts. The nominal prediction assuming no matter effect or CP violation was also used and falls in the middle. Five events were observed in the sideband, all visually consistent with a neutrino topology and in line with the predicted range.

	ν_e CC MC	bkg. MC	cosmic data	total
inverted order, $\delta = \pi/2$	1.17	1.50	0.38	3.05
no matter effect, $\delta = 0$	2.28	1.50	0.38	4.16
normal order, $\delta = 3\pi/2$	2.92	1.47	0.38	4.77

(a) Extrapolated prediction

	Neutrino-like	Cosmic-like	total
$0.7 < \text{LID} < 0.95$	5	0	5

(b) Far Detector Data Result

Table 7.3: Prediction of Far Detector neutrino background and out-of time cosmic data in the low-LID sideband. Panel A shows the prediction based on the extrapolation of Near Detector data normalized to 3.52×10^{20} POT. Panel B shows the measured result in data, where the 5 candidate events were visually hand-scanned and classified as appearing like a neutrino or cosmic ray.

7.3 Systematic Uncertainties

Studies of systematic effects on the signal and background prediction for the ν_e CC appearance analysis used one of two approaches. Beam flux (Section 7.3.1), scintillator response (Section 7.3.2), calibration (Section 7.3.3), and neutrino interaction effects (Section 7.3.5) are evaluated using modified Monte Carlo samples in the Near and Far Detector to extrapolate variations with respect to the nominal. Data driven techniques are used to study Near De-

tector containment effects (Section 7.3.6), light levels (Section 7.3.4), rock modeling (Section 7.3.7), and decomposition method (Section 7.3.8). The statistical limitations of the Monte Carlo sample (Section 7.3.9) and an overall normalization systematic due to POT accounting and mass uncertainties (Section 7.3.10) was also assessed. Additional systematics such as alignment (Section 7.3.11) and bad channel masks (Section 7.3.12) were considered and found to be negligible. For further details on the systematic uncertainties studied see [134].

7.3.1 Beam

Uncertainties in the neutrino flux seen at the Near and Far Detectors is broken into two categories: *Beam Transport*, referring to differences between the working NuMI beam and the simulation [135], and *Hadron production* meaning uncertainty in the simulated production of pions and kaons on the target [136]. The size of beam transport effects (horn current, beam spot size, beam location on target, relative position of magnetic horns, magnetic field distribution) is motivated by measurements from beam monitoring devices. Uncertainties on hadron production are obtained by benchmarking the flux Monte Carlo against data from fixed target experiments.

Variant beam Monte Carlo samples were produced for each beam parameter and neutrino fluxes were calculated at the Near and Far Detector locations. The ratios of the varied samples with respect to the nominal were turned into weights in true energy bins separated by neutrino type and detector. These weights are then used to shift the nominal Monte Carlo for both detectors and carry out the extrapolation with the standard Near Detector data. Each beam systematic component is extrapolated separately to produce an uncertainty with respect to the nominal extrapolation. These component uncertainties are added in quadrature to produce an overall uncertainty of 1.1% on the signal and 3.2% on the background. The Near Detector LID and energy spectra are shown in Figure 7.7

with the one sigma beam systematic band. In each detector the flux uncertainty is on the order of 30% but the extrapolation technique significantly reduces the effect. The largest component of the beam systematic comes from the hadron production.

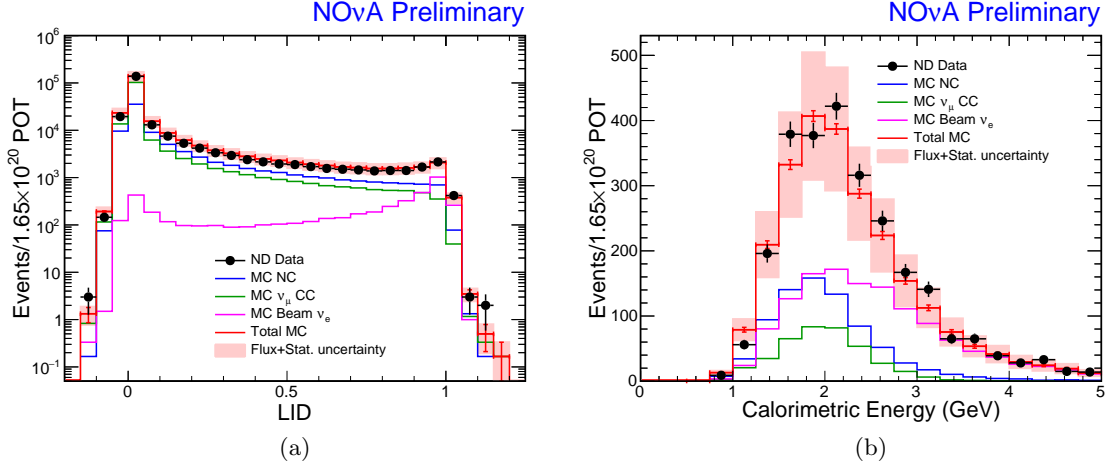


Figure 7.7: LID spectrum (a) and calorimetric energy distribution with LID > 0.95 (b) of Near Detector data and Monte Carlo with the beam flux uncertainty band drawn.

7.3.2 Birk's Law

In a scintillator-based detector the light yield is assumed to be proportional to the particle energy deposition for small energy deposits. At high energy deposition the light yield begins to quench. The two parameter Birks-Chou law [137] gives an empirical relation between energy deposition and light yield:

$$\frac{dL}{dX} = L_0 \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2}$$

where L is the luminescence, L_0 is the luminescence at low specific ionization density, and k_B and k_C are constants dependent on the scintillating medium. The parameters were measured in the NOvA Near Detector by measuring the dE/dx in planes at the end of

proton tracks in data and Monte Carlo [138, 139]. The Birks parameters were adjusted in Monte Carlo to produce the best agreement with data, giving $k_B = 0.04$ cm/MeV and $k_C = -0.0005$ cm²/MeV². These numbers are a factor of four higher than typical measurements in other experiments, however no measurement exists for the PPO based scintillator blend used in NOvA and reported results are almost all at energies below 10 MeV. A conservative uncertainty is therefore assigned to the Birks-Chou parameters. Two special Monte Carlo samples are generated for each detector, one with $k_B = 0.01$ cm/MeV and $k_C = 0$ and the other with $k_B = 0.02$ cm/MeV and $k_C = 0$. These shifted samples are used to extrapolate the standard Near Detector data to measure the variance in signal and background predictions with respect to the nominal. The result is an uncertainty of 7.2% on the signal and 5.1% on the background.

7.3.3 Calibration

Deliberate mis-calibrations of the Monte Carlo are used to assess calibration systematics. Comparisons of data and Monte Carlo show both a flat and spatially dependent disagreement in energy response after calibration has been applied. Four mis-calibration effects are accounted for: an absolute calibration shift, a relative shift between detectors, a random cell-by-cell calibration variation, and a shift as a function of position along the length of a cell. Studies of muon dE/dx using muons from rock neutrino interactions (independent source from the calibration method) [140], measuring the π^0 mass peak [141], and measuring the energy spectrum of Michel Electrons [142] which provide high statistics orthogonal to the standard calibration procedure, placed the absolute calibration systematic at 5% for both detectors. This 5% uncertainty is also taken as the relative energy scale offset between detectors. Random cell-by-cell mis-calibration is applied by shifting each cell using a gaussian distribution with $\sigma = 8\%$. The sloped mis-calibration applies a second degree

polynomial function as a function of cell length to create the largest shift at the far end of a Far Detector cell ($\sim 20\%$) [143]. Table 7.4 shows the systematic uncertainties in the signal and background predictions at the Far Detector after extrapolating the calibration shifts. Combining the systematics in quadrature gives an overall uncertainty on the signal of 7.5% and 4.5% for the background.

% diff	signal	total bkg.	ν_μ CC	NC	beam ν_e CC
Absolute	6.0	0.8	8.5	10.8	10.8
Relative	4.4	3.7	12.1	15.7	6.4
Random	-0.9	-0.4	1.5	1.5	1.9
Slope	1.3	2.3	2.9	10.5	3.9
Total	7.6	4.4	15.1	21.8	13.4

Table 7.4: Calibration systematic error relative change from nominal for the signal and background components of the LID selection. Table from [134].

7.3.4 Light Levels

There is some evidence that the Far Detector light level should be adjusted in the simulation [144], specifically that the far end of cells is producing 20% less light in the current simulation than data. To first order this effect is accounted for in the calibration, but second order effects such as hits falling above/below threshold may have an impact. The impact of this uncertainty is evaluated by examining the change in the Monte Carlo signal selection efficiency and the number of hits in the selected slice as a function of detector location [145]. A 5% change in the number of hits is equivalent to a 1% systematic error on the signal. There is negligible effect on the background.

7.3.5 Neutrino Interaction

A cocktail of 33 neutrino interaction systematics are evaluated by reweighting neutrino interactions simulated in GENIE. The effects fall generally into three categories: cross-section uncertainties, hadronization model uncertainties, and uncertainties due to final state interactions (for a complete listing see [134]). The size of the one sigma reweight parameter

for each variable comes from the GENIE authors based on a careful survey of the interaction model and inputs from various experiments including MINOS and T2K and is stored in the standard NOvA simulation [146]. The standard extrapolation procedure was run on Near Detector data with the Monte Carlo for both detectors reweighted by $\pm 1\sigma$ for each GENIE handle in turn. Summing the results in quadrature gives a total uncertainty on the signal prediction of 7.3% and 2.1% for the background. The largest single source of uncertainty in the signal comes from the axial mass in quasi-elastic events (6.9%). There was not a single dominant source of uncertainty in the background.

7.3.6 Near Detector Containment

Biases and uncertainty in the Near Detector containment are estimated by dividing the detector in half along all three axes and performing the extrapolation in each of six regions based on vertex location. Additionally, the fiducial volume of the Near Detector was divided equally into an inner and outer region. These effects were only considered for the backgrounds, the Near Detector ν_μ CC selection used to extrapolate the signal uses a different selection process [130]. The results in Table 7.5 show an average uncertainty of 1.8%.

Diff %	total bkg.	ν_μ CC	beam ν_e CC	NC	ν_τ CC
$100 < vZ < 400$ cm	-2.6	-1.5	-2.0	-3.7	0
$400 < vZ < 700$ cm	1.6	3.1	2.0	1.0	0
$0 < vX < 140$ cm	-2.7	-3.1	-2.8	-2.9	0
$-140 < vX < 0$ cm	2.1	4.6	3.1	0.5	0
$0 < vY < 140$ cm	-1.9	0	-0.9	-3.7	0
$-140 < vY < 0$ cm	0.6	1.5	0.6	0.5	0
$ vX, vY < 99$ cm	-1.9	-1.5	-1.8	-2.2	0
$ vX, Vy < 140$ cm	1.2	3.1	2.0	0	0
Absolute Average Uncertainty	1.8%	2.3%	1.9%	1.8%	0%

Table 7.5: Using the standard Near Detector decomposition technique, the extrapolation is performed with eight different geographic regions of the Near Detector data and Monte Carlo to measure containment effects. Values recorded are the relative percent difference with respect to the nominal full-detector extrapolation. The final result is an absolute average of the uncertainties.

7.3.7 Rock Contamination

Uncertainties are introduced in backgrounds from neutrino interactions in the rock surrounding the Near Detector. Due to the computing demands of simulating the large volume, a library of these events are generated and then mixed into simulated files of neutrino interactions inside the detector, with each rock event sampled ~ 350 times, causing the same rock event to be selected many times. It was found that the rock rate is too high in the Monte Carlo, but after all preselection cuts have been applied the rock rate is similar in data and Monte Carlo and small [147]. The effect on the Far Detector extrapolated prediction is estimated by performing the extrapolation with all rock interactions in the simulation removed by truth and comparing to the nominal. The result is an overall background systematic uncertainty of 0.1%

7.3.8 Decomposition

The standard extrapolation technique for the analysis involves proportionally decomposing the Near Detector data based on the Monte Carlo prediction of each component in each energy bin. The error on this method is estimated by alternatively assigning the entire excess/deficit between data and Monte Carlo in the Near Detector to one component (ν_μ CC, ν_e CC, or NC) as shown in 7.6. The variation in method results in a maximum difference of 3.9% which is taken as the uncertainty.

Diff %	total bkg.	ν_μ CC	beam ν_e CC	NC	ν_τ CC
LID ν_e decomposition	0.8	-4.62	6.41	-5.61	0
LID NC decomposition	0.86	-4.62	-6.23	11.46	0
LID ν_μ decomposition	-3.94	26.15	-6.23	-5.61	0

Table 7.6: Extrapolated prediction of Far Detector events, recorded as the relative percentage change from nominal. Different decomposition methods for the Near Detector data are shown. Any excess/deficit between data and Monte Carlo in an energy bin is either distributed proportionally among the components, or assigned entirely to one component.

7.3.9 Data and Monte Carlo Statistics

The extrapolation uses a Monte Carlo sample equal to 10 times the Near Detector data sample and 1000 times the Far Detector data. The resulting Monte Carlo statistical uncertainty is taken as a systematic and is different from the statistical uncertainty in the Far Detector measurement due to the small expectation of events. Table 7.7 contains the statistical uncertainties for the background components in each data/Monte Carlo sample. The systematic uncertainty in the background is taken as 3.5%. The calculation of uncertainty on the signal is a 0.6% effect.

Channel	FD MC	ND Data	ND MC	Total (%)
beam ν_e	4.0	2.5	1.0	4.8
ν_μ	6.2	2.5	1.1	6.8
NC	2.1	2.5	1.1	3.4
Total	2.3	2.4	1.0	3.5

Table 7.7: Statistical uncertainty on extrapolation of background components in percents. Table from [134].

7.3.10 Normalization

Several sources of uncertainty create an overall normalization uncertainty in the Far Detector event rate. A 0.5% uncertainty in POT accounting differences in the Near and Far Detector is taken from beam monitoring devices [135]. Near Detector data/Monte Carlo comparisons showed a 0.8% difference in reconstruction efficiency. A 0.7% uncertainty is assessed from differences of Near and Far Detector measurements of the mass of detector components [148]. Adding in quadrature gives a 1.2% normalization systematic on the signal and background prediction.

7.3.11 Alignment

Data and Monte Carlo in both detectors is reconstructed using an ideal Monte Carlo simulation. The Near Detector data was reprocessed and extrapolated using an alternative as-built geometry based on survey data. The difference in predicted event rates with respect to the nominal extrapolation is at the sub-percent level and is neglected.

7.3.12 Bad Channels

In both detectors a portion of channels are masked off in the analysis due to issues both with individual pixels and entire APDs (noise, cooling, data flow, ect.). This number varies as routine maintenance is performed but is kept well under 1% of the entire detector. Studies performed with random bad channel masks of varying percentage found that the figure of merit for event selection at the Far Detector remained stable at the 1% level with as much as 7% of the detector masked off. Both the reconstruction and event classification are robust against routine levels of bad channels. No systematic error is assigned to the masking.

7.3.13 Systematic Error Table

Table 7.8 shows the final systematic uncertainties on the Far Detector signal and background event rates for the the ν_e CC appearance measurement. Summed in quadrature the total uncertainty is 12.9% for the signal and 9.6% for the background. The systematic uncertainties are shown graphically in Figure 7.8

7.4 Prediction

Taking a representative set of oscillation weights that assume no matter effect or CP violation, $\sin^2 2\theta_{23} = 1$, $\Delta m_{32}^2 = 2.35 \times 10^{-3}$, and $\sin^2 2\theta_{13} = 0.1$, a Far Detector prediction is formed combining the extrapolation with systematic uncertainties and the measurement

	Total BG Syst. (%)	Signal Syst. (%)
Beam	3.2	1.1
Birks	5.1	7.2
Calibration	4.4	7.6
Light Level	n/a	1.0
Neutrino Interaction	2.14	7.28
Containment	1.84	n/a
Rock contamination	0.10	n/a
Decomposition	3.90	n/a
Data & MC Stat.	3.47	0.6
Normalization	1.2	1.2
Total	9.6	12.9

Table 7.8: Final systematic uncertainty (in percentages) on the combined background and signal in the Far Detector for LID. The last row corresponds to the sum in quadrature of all errors in this table. Table from [134].




Figure 7.8: Systematic uncertainties in percentages on background (a) and signal (b) in the Far Detector.

of the cosmic background. The ν_e CC signal expectation is $4.41 \pm 2.10(\text{stat.}) \pm 0.57(\text{syst.})$ events for 2.8×10^{20} full-detector equivalent POT. In the absence of oscillations the combined cosmic and beam background is $0.97 \pm 0.98(\text{stat.}) \pm 0.09(\text{syst.})$ events, with 0.06 events expected from cosmic rays.

CHAPTER 8

Results

This chapter presents the results of the first analysis of electron neutrino appearance in the NOvA experiment. Confidence intervals of δ vs $\sin^2 \theta_{13}$ are produced for both the normal and inverted mass orderings. The precise measurement of $\sin^2 \theta_{13}$ from the reactor experiment Daya Bay is used to constrain the fit for a preferred value of δ and the mass ordering. These results are compared to the T2K experiment which measures the ν_e appearance oscillation channel at a shorter baseline of 295 kilometers. Section 8.1 describes the statistical treatment of fitting oscillation parameters to the data, Section 8.2 presents the data seen in the NOvA experiment, and Section 8.3 shows the results of the oscillation parameter fitting. In 3.52×10^{20} POT of data 6 ν_e CC events were observed on a background of $0.97 \pm 0.09(\text{syst.})$ events which is a 3.3σ observation of electron neutrino appearance at the longest baseline in the world. The normal mass ordering is weakly favored over the inverted ordering, with a clearer picture possible in the summer of 2016 with triple the exposure.

8.1 Statistical treatment of PMNS Parameter Measurements

The ν_e appearance analysis is a “cut-and-count” analysis where the only information from the Far Detector data that enters into oscillation parameter fits is the number of events passing ν_e CC selection criteria in the 1.5 to 2.7 GeV energy window. Fits to oscillation parameters compare the observed number of events, n to the expected number of events,

m_i , where i represents a particular bin in oscillation parameter space. The expectation m depends on the values chosen for θ_{13} , θ_{23} , θ_{12} , Δm_{32}^2 , Δm_{21}^2 , δ , and the mass ordering. The expectation is the sum of signal events from the $\nu_\mu \rightarrow \nu_e$ oscillation channel, neutrino backgrounds from other oscillation channels, and cosmic-ray backgrounds predicted from out-of-time data. Some of the oscillation parameters are fixed or constrained by external data as described below. The expected event count is made from the extrapolated Far Detector event count (Section 7.1) scaled to the same exposure as the data for a given set of oscillation parameters. In this analysis a binned search of either the 2D δ vs $\sin^2 \theta_{13}$ parameter space or the 1D δ space is performed with the other oscillation parameters either fixed or marginalized. For the Poisson-distributed data the best fit for the oscillation parameters is found by minimizing the χ^2 [28]:

$$\chi^2 = 2 \sum_i \left(m_i - n + n \ln \left(\frac{n}{m_i} \right) \right) + \sum_j \frac{s_j^2}{\sigma_j^2}, \quad (8.1)$$

where the second term is a penalty for the systematic uncertainties, where s_j is the best-fit uncertainty and σ_j is the 1σ value assigned to the j^{th} source of systematic uncertainty (Table 7.8). In each bin of the oscillation parameter space MINUIT [105] is used to marginalize over the other systematic and oscillation variables to produce the minimum χ^2 . A complete overview of the analysis framework for the fitting can be found in [149].

In this analysis all systematic uncertainties have been added in quadrature to produce an overall uncertainty on both the number of expected signal and background events. The size of the systematic shifts are calculated at ± 1 , ± 2 , and $\pm 3 \sigma$ and interpolated over to describe a χ^2 function used in MINUIT to optimize the fit.

The detector locations sets a baseline of 810 kilometers [70]. The CRUST2.0 [150] model of the Earth's crust is used to calculate a density of 2.84 g/cm^3 at the average neutrino depth underground on the NOvA baseline of 9.38 kilometers. The solar oscillation

parameters are fixed at the current global best-fit points of $\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{eV}^2$ and $\sin^2 2\theta_{12} = 0.846$ [28]. The 2014 best fit of $\sin^2 2\theta_{13}$ and $0.084_{-0.005}^{+0.005}$ from the Daya Bay reactor experiment is used as a constraint in fits for the mass ordering information [49]. The atmospheric parameters, $\sin^2 \theta_{23}$ and Δm_{32}^2 , are constrained by the best fit of the T2K experiment using a publicly available χ^2 map of the atmospheric parameter space for both the normal and inverted mass ordering assumptions (Figure 2.13) [151].

Parameter	Nominal Value
Baseline	810 km
Matter Density	2.84 g/cm ³
Δm_{21}^2	$7.53 \times 10^{-5} \text{ eV}^2$
$\sin^2 2\theta_{12}$	0.846
$\sin^2 2\theta_{13}$	$0.084_{-0.005}^{+0.005}$
δ_{CP}	0
Δm_{32}^2	$2.51_{-0.10}^{+0.10} \times 10^{-3} \text{eV}^2$ (NH)
	$2.48_{-0.10}^{+0.10} \times 10^{-3} \text{eV}^2$ (NH)
$\sin^2 \theta_{23}$	$0.514_{-0.056}^{+0.055}$ (NH)
	$0.511_{-0.055}^{+0.055}$ (IH)

Table 8.1: Summary of oscillation parameters and uncertainties used in the analysis. Table adapted from [152].

When marginalizing over θ_{13} with the constraint of the Daya Bay experiment a term is added to the χ^2 of the form:

$$\chi_{\text{Daya Bay}}^2 = \left(\frac{\sin^2 \theta_{13} - (\sin^2 \theta_{13})_{\text{Daya Bay}}}{\sigma_{\text{Daya Bay}}} \right)^2. \quad (8.2)$$

Note that Daya Bay measures $\sin^2 2\theta_{13}$ and the term is converted for this analysis. When marginalizing over the atmospheric parameters using the T2K measurement the χ^2 penalty is interpolated from the likelihood maps provided by the experiment [151].

8.2 Far Detector Data Observation

In the analysis period of NOvA Far Detector data from February 6th, 2014 until May 15th, 2015, encompassing 3.52×10^{20} POT, 6 events were observed in-time with the NuMI beam

passing the LID selection criteria. Figure 8.1 shows the energy distribution of the selected events and the distribution of events in the exposure. Performing a Kolmogorov-Smirnov test on the event distribution gives a p value of 0.32. Figures 8.2 and 8.3 show event displays of the selected events. Under visual inspection all six events appear consistent with the topology of a neutrino interaction as opposed to cosmic-ray background. All candidates appear as plausible ν_e CC interactions.




Figure 8.1: a) Accumulation of ν_e candidate events as a function of exposure. b) Calorimetric energy distribution of selected events compared to the ν_e CC signal and beam neutrino background projection with the standard set of oscillation weights assuming no matter effect or CP violation.

8.3 Confidence Intervals

The observation of 6 ν_e CC candidate events in the analysis is 3.3σ evidence that θ_{13} is nonzero by itself. The significance of observing 6 or more events on a background in the absence of θ_{13} of $0.97 \pm 0.09(\text{syst.})$ ($b_0 \pm \sigma_b$) events is calculated by integrating the Poisson probability convoluted with a Gaussian for the systematic uncertainty:

$$P = \sum_{x=6}^{\infty} \int_{-\infty}^{+\infty} \frac{b^x}{x! \sigma_b^2 b_0^2 \sqrt{2\pi}} \exp\left(-b - \frac{(b - b_0)^2}{2\sigma_b^2 b_0^2}\right) db. \quad (8.3)$$

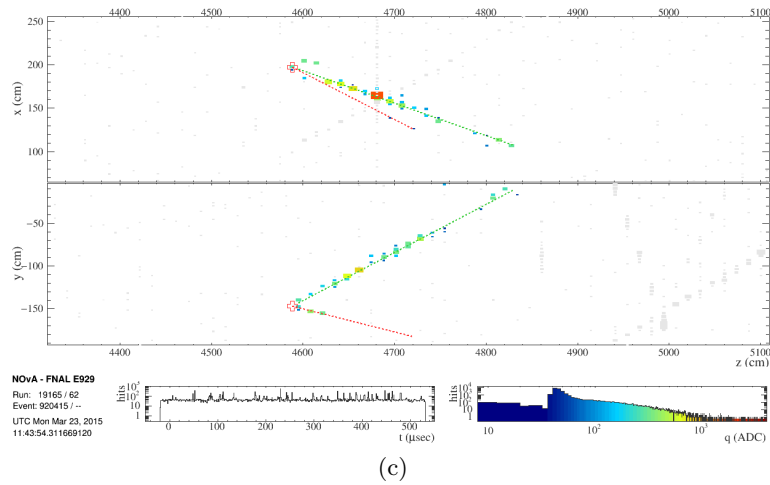
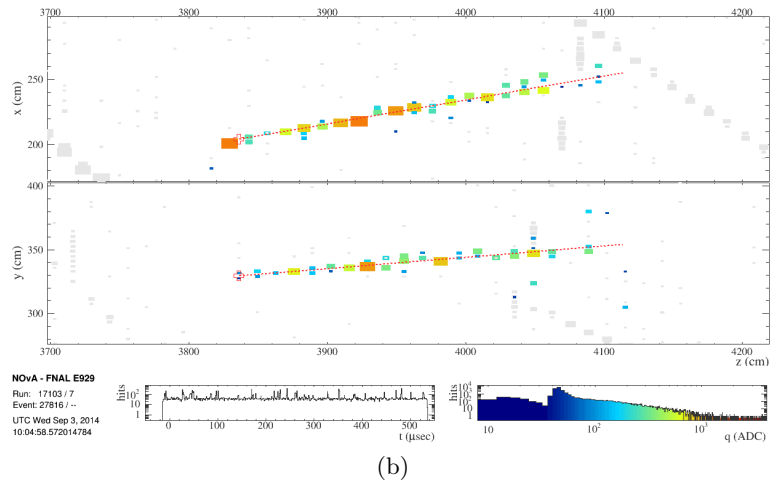
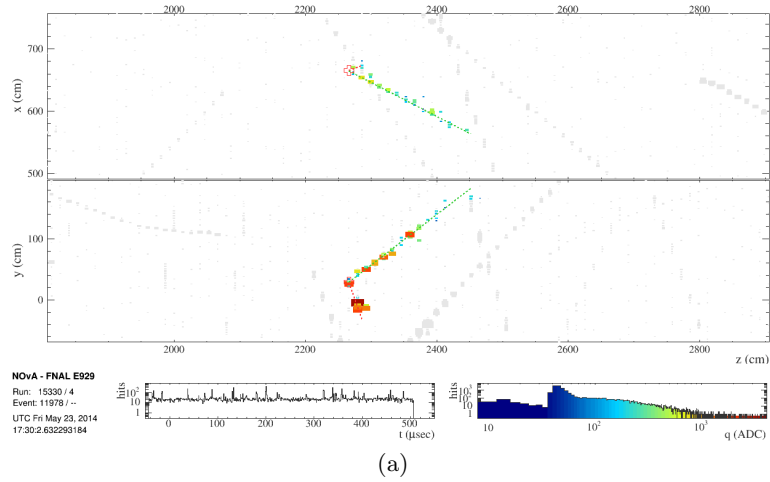


Figure 8.2: Event displays for the LID selected events in the NUMI beam window for 3.52×10^{20} POT of Far Detector data. The reconstructed vertex and 3D reconstructed prongs are drawn. Displays are zoomed on the interaction of interest but preserve the aspect ratio of the full detector. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time with the slice.

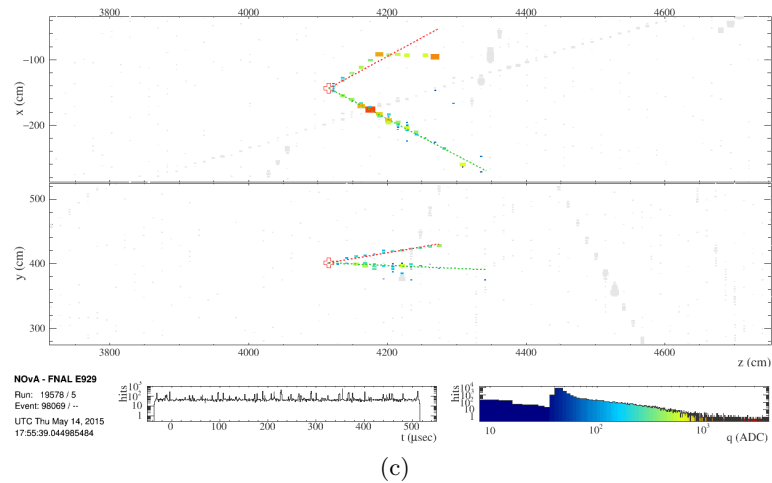
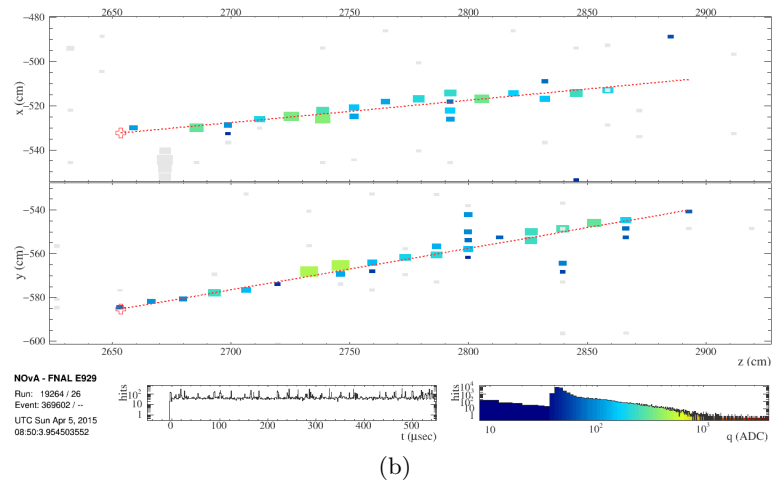
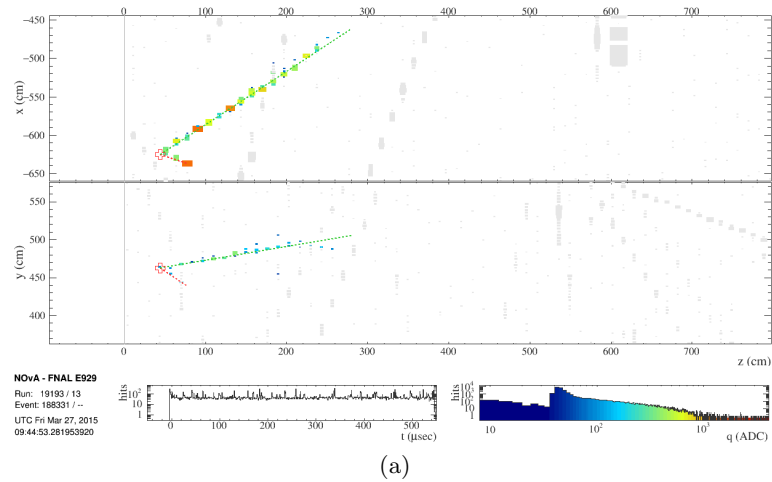


Figure 8.3: More event displays for the LID selected events in the NUMI beam window for 3.52×10^{20} POT of Far Detector data. The reconstructed vertex and 3D reconstructed prongs are drawn. Displays are zoomed on the interaction of interest but preserve the aspect ratio of the full detector. Hits are colored and scaled by the charge deposited. Grayed hits are out-of-time with the slice.

A two dimensional fit is performed in δ vs $\sin^2 \theta_{13}$ space constraining the atmospheric oscillation parameters based on the T2K results and fixing the solar parameters. The result is best fit values of

$$\sin^2 \theta_{13} = 0.021_{-0.011}^{+0.024} \text{ (NH)}$$

$$\sin^2 \theta_{13} = 0.043_{-0.027}^{+0.027} \text{ (IH)}$$

for the normal and inverted mass orderings. This value is larger than the more precisely measured value in Daya Bay and indicates that the θ_{23} octant, mass ordering, and CP violating effects are combining to enhance the ν_e appearance probability (refer to Figure 2.15). The 2D contour is shown in comparison to Daya Bay and T2K in Figure 8.4. After marginalizing over δ the 1D fit of $\sin^2 \theta_{13}$ is shown in Figure 8.5.

It is more instructive to use the precision value of θ_{13} from the Daya Bay experiment as a constraint and fit only in δ for the normal and inverted mass orderings as shown in Figure 8.6. Using the reactor information NOvA weakly favors the normal mass ordering over the inverted with a minimum difference in χ^2 between the curves of 0.2. All values of δ are valid at the 1σ level for the normal mass ordering. This fit leans in the same direction as the T2K observation. The longer baseline of the NOvA experiment (810 vs 295 km) creates a greater separation between the mass orderings in terms of event expectations. As NOvA continues to accumulate statistics this difference will give NOvA increased sensitivity to the mass ordering.

8.3.1 Discussion

Both the NOvA and T2K experiments see an excess of ν_e CC appearance events in a ν_μ beam than predicted using the θ_{13} measurement of the reactor experiment Daya Bay and assuming $\delta = 0$ and $\theta_{23} = \pi/4$. This gives a small preference to the normal mass ordering for





Figure 8.4: a) δ vs $\sin^2\theta_{13}$ contour of the NOvA result for the normal (red) and inverted (blue) mass ordering. The fit was performed without reactor experiment constraints. b) The 90% percent contours of NOvA are compared to the T2K contours from an analysis of 28 events in 6.6×10^{20} POT.






Figure 8.5: Best fit of $\sin^2 \theta_{13}$ for the normal (red) and inverted (blue) mass orderings after marginalizing over δ from Figure 8.4. The fit does not include a constraint on θ_{13} from the reactor experiments.

NOvA, and a 1σ preference for T2K. This scenario is the best-case situation for the NOvA experiment as discussed in Section 2.10. This result slightly increases the tension with the MINOS experiment which has a weak preference for the inverted mass ordering, although the MINOS detector is not optimally designed for electrons [39]. Joint fits of the experiments and additional statistics are needed to resolve this tension. Now that construction of the NOvA experiment is complete and the NuMI beam power has increased above 500 kW the NOvA exposure (POT \times kT) is expected to triple by the summer of 2016. If NOvA continues to accumulate ν_e CC events at the current rate with no changes to the analysis the preference for the normal mass ordering will increase to $\sim 1 \sigma$ on par with T2K, see Figure 8.7. There are avenues of improvement in the reconstruction, event classification, and cosmic-ray background rejection that could increase the fiducial mass and ν_e CC selection efficiency for additional gains in this analysis.



(a)



(b)

Figure 8.6: a) Best fit of δ to the NOvA data, constraining θ_{13} from the Daya Bay measurements and the atmospheric parameters from T2K. The normal mass ordering (red) is slightly preferred over the inverted ordering (blue). b) The T2K contours under the same parameter constraints with 28 candidate ν_e CC events from an analysis of 6.6×10^{20} POT [64] is shown on the same plot for comparison.




Figure 8.7: Best fit of δ to the NOvA data assuming 10.56×10^{20} POT and 18 events, which is three times the dataset used in this analysis, expected by the summer of 2016.

CHAPTER 9

Conclusions

This thesis presented the analysis of 3.52×10^{20} protons-on-target of data in the NOvA Far Detector collected between February 2014 and May 2015 for evidence of electron neutrino appearance. During this data period both the detector and the neutrino beam underwent commissioning and thus the exposure is the equivalent of about a half year of nominal operations with a completed detector and 700 kW beam. Six ν_e CC events were observed which is a 3.3σ observation of the oscillation channel $\nu_\mu \rightarrow \nu_e$ at the longest baseline for such an experiment in the world.

Neutrino oscillation physics is moving into a precision era now that all three mixing angles and both mass splittings have been well measured. Two important questions that remain are the ordering of the neutrino masses and if CP violation is non-zero in the neutrino sector. Observing CP violation is one of the requirements in the “see-saw” mechanism theory for the generation of neutrino mass which would also provide an explanation for the matter-antimatter asymmetry of the universe. Knowledge of the neutrino mass ordering contributes to narrowing the parameter space of neutrinoless double beta decay experiments attempting to confirm the Dirac or Majorana nature of the neutrino. The T2K experiment measures the same oscillation channel as NOvA at a third of the baseline and favors the normal mass ordering at 1σ with a weak preference for maximal CP violation. In this analysis a very weak preference for the normal ordering was also seen. The NOvA exposure is expected

to triple by the summer of 2016. If NOvA continues to observe electron neutrinos at the present rate, 1σ resolution of the mass ordering may be achieved at that time, comparable to T2K. The longer baseline of the NOvA experiment enhances the separation between the mass orderings. The NOvA experiment plans to collect data for six to ten years, with roughly half the data coming in anti-neutrino mode.

This thesis contributed significantly to all phases of the analysis. In particular, a general track reconstruction algorithm was developed robust enough to handle electron showers, muon tracks, displaced photons, and small energy deposits. This algorithm formed the basis for the classification algorithm used to identify electron neutrinos. The analysis strategy presented in this thesis was successful in observing electron neutrino appearance oscillations in a muon neutrino beam consistent with the global picture. The groundwork was laid for future analyses where the sensitivities to the neutrino mass ordering can set new global limits.

BIBLIOGRAPHY

- [1] Wolfgang Pauli. Letter to the physical society of Tübingen. 1930.
- [2] J. Chadwick. Distribution in intensity in the magnetic spectrum of the β -rays of radium. *Deutch. Phys. Gessel.*, 16.8:883–891, 1914.
- [3] E. Fermi. Tentativo di una teoria dei raggi β . *Ricerca Scientifica*, 1933.
- [4] J. Chadwick. Possible Existence of a Neutron. *Nature*, 129:312, 1932.
- [5] B. Pontecorvo. Chalk River Laboratory Report PD-205. 1946.
- [6] C.L. Cowan and F. Reines. Detection of the Free neutrino. *Phys. Rev.*, 92:830–831, 1953.
- [7] H.W. Kruse A.D. McGuire F. Reines, F.B. Harrison and C.L. Cowan. Detection of the Free neutrino: a Confirmation. *Science*, 124:103–104, 1956.
- [8] G. Danby et al. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys.Rev.Lett.*, 9:36–44, 1962.
- [9] K. Kodama et al. Observation of tau neutrino interactions. *Phys.Lett.*, B504:218–224, 2001.
- [10] S. Schael et al. Precision electroweak measurements on the Z resonance. *Phys.Rept.*, 427:257–454, 2006.

- [11] K.S. Hirata et al. Observation of 8B solar neutrinos in the Kamiokande-II detector. *Phys.Rev.Lett.*, 63:16, 1989.
- [12] A.I. Abazov et al. Search for neutrinos from sun using the reaction Ga-71 (electron-neutrino e^-) Ge-71. *Phys.Rev.Lett.*, 67:3332–3335, 1991.
- [13] P. Anselmann et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Phys.Lett.*, B285:376–389, 1992.
- [14] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys.Rev.Lett.*, 81:1562–1567, 1998.
- [15] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys.Rev.*, C72:055502, 2005.
- [16] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172, 1958.
- [17] M. Nakagawa Z. Maki and S. Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870, 1962.
- [18] P.A.R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. 2015.
- [19] A. de Gouvea et al. Working Group Report: Neutrinos. 2013.
- [20] Murray Gell-Mann, Pierre Ramond, and Richard Slansky. Complex Spinors and Unified Theories. *Conf.Proc.*, C790927:315–321, 1979.
- [21] Rabindra N. Mohapatra and Goran Senjanovic. Neutrino Mass and Spontaneous Parity Violation. *Phys.Rev.Lett.*, 44:912, 1980.

- [22] Enrico Nardi, Yosef Nir, Esteban Roulet, and Juan Racker. The Importance of flavor in leptogenesis. *JHEP*, 0601:164, 2006.
- [23] E. Molinaro and S.T. Petcov. A Case of Subdominant/Suppressed 'High Energy' Contribution to the Baryon Asymmetry of the Universe in Flavoured Leptogenesis. *Phys.Lett.*, B671:60–65, 2009.
- [24] S. Pascoli, S.T. Petcov, and Antonio Riotto. Leptogenesis and Low Energy CP Violation in Neutrino Physics. *Nucl.Phys.*, B774:1–52, 2007.
- [25] Boris Kayser. Two Questions About Neutrinos. 2010.
- [26] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, UK, 2007.
- [27] Hiroshi Nunokawa, Stephen J. Parke, and Jose W.F. Valle. CP Violation and Neutrino Oscillations. *Prog.Part.Nucl.Phys.*, 60:338–402, 2008.
- [28] K.A. Olive et al. Review of particle physics. *Chin.Phys.*, C38:090001, 2014.
- [29] J.L. Hewett, H. Weerts, R. Brock, J.N. Butler, B.C.K. Casey, et al. Fundamental Physics at the Intensity Frontier. 2012.
- [30] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.
- [31] S.P. Mikheev and A. Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov.J.Nucl.Phys.*, 42:913–917, 1985.
- [32] Martin Freund. Analytic approximations for three neutrino oscillation parameters and probabilities in matter. *Phys.Rev.*, D64:053003, 2001.
- [33] Y. Ashie et al. A Measurement of atmospheric neutrino oscillation parameters by SUPER-KAMIOKANDE I. *Phys.Rev.*, D71:112005, 2005.

- [34] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillation. *Phys.Rev.Lett.*, 93:101801, 2004.
- [35] K. Abe et al. Evidence for the Appearance of Atmospheric Tau Neutrinos in Super-Kamiokande. *Phys.Rev.Lett.*, 110(18):181802, 2013.
- [36] K. Abe, , et al. First muon-neutrino disappearance study with an off-axis beam. *Phys. Rev. D*, 85:031103, Feb 2012.
- [37] K. Abe et al. Measurement of Neutrino Oscillation Parameters from Muon Neutrino Disappearance with an Off-axis Beam. *Phys.Rev.Lett.*, 111(21):211803, 2013.
- [38] Yoshitaka Itow. Recent results in atmospheric neutrino oscillations in the light of large Θ_{13} . *Nucl.Phys.Proc.Suppl.*, 235-236:79–86, 2013.
- [39] P. Adamson et al. Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys.Rev.Lett.*, 110(25):251801, 2013.
- [40] M.H. Ahn et al. Measurement of Neutrino Oscillation by the K2K Experiment. *Phys.Rev.*, D74:072003, 2006.
- [41] N. Agafonova et al. Discovery of tau neutrino appearance in the CNGS neutrino beam with the OPERA experiment. 2015.
- [42] A. Renshaw et al. First Indication of Terrestrial Matter Effects on Solar Neutrino Oscillation. *Phys.Rev.Lett.*, 112(9):091805, 2014.
- [43] G. Bellini et al. Absence of day–night asymmetry of 862 keV ${}^7\text{Be}$ solar neutrino rate in Borexino and MSW oscillation parameters. *Phys. Lett.*, B707:22–26, 2012.
- [44] Abhijit Bandyopadhyay, Sandhya Choubey, Raj Gandhi, Srubabati Goswami, and D.P. Roy. The Solar neutrino problem after the first results from KamLAND. *Phys.Lett.*, B559:121–130, 2003.

- [45] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys.Rev.Lett.*, 100:221803, 2008.
- [46] M. Apollonio et al. Limits on neutrino oscillations from the CHOOZ experiment. *Phys.Lett.*, B466:415–430, 1999.
- [47] F.P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys.Rev.Lett.*, 108:171803, 2012.
- [48] F. P. An et al. Spectral measurement of electron antineutrino oscillation amplitude and frequency at daya bay. *Phys. Rev. Lett.*, 112:061801, Feb 2014.
- [49] Chao Zhang. Recent Results From The Daya Bay Experiment. 2015.
- [50] J.K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys.Rev.Lett.*, 108:191802, 2012.
- [51] Seon-Hee Seo. New Results from RENO. *PoS, Neutel2013:018*, 2014.
- [52] Seon-Hee Seo. New Results from RENO and The 5 MeV Excess. 2014.
- [53] Y. Abe et al. Indication for the disappearance of reactor electron antineutrinos in the Double Chooz experiment. *Phys.Rev.Lett.*, 108:131801, 2012.
- [54] Y. Abe et al. Reactor electron antineutrino disappearance in the Double Chooz experiment. *Phys.Rev.*, D86:052008, 2012.
- [55] Y. Abe et al. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector. *Journal of High Energy Physics*, 2014(10), 2014.
- [56] Hervé de Kerret. Results from double chooz. 26th International Conference on Neutrino Physics and Astrophysics (Neutrino 2014), 2014.

- [57] K. Abe et al. Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam. *Phys.Rev.Lett.*, 107:041801, 2011.
- [58] K. Abe et al. Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys.Rev.*, D88(3):032002, 2013.
- [59] K. Abe et al. Observation of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys.Rev.Lett.*, 112:061802, 2014.
- [60] P. Adamson et al. Improved search for muon-neutrino to electron-neutrino oscillations in MINOS. *Phys.Rev.Lett.*, 107:181802, 2011.
- [61] P. Adamson et al. Electron neutrino and antineutrino appearance in the full MINOS data sample. *Phys.Rev.Lett.*, 110(17):171801, 2013.
- [62] J. Beringer et al. Review of Particle Physics (RPP). *Phys.Rev.*, D86:010001, 2012.
- [63] P. Adamson et al. Combined analysis of ν_μ disappearance and $\nu_\mu \rightarrow \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos. *Phys.Rev.Lett.*, 112:191801, 2014.
- [64] K. Abe et al. Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with $6.6E20$ protons on target. 2015.
- [65] Alexander Himmel. Recent results from Super-Kamiokande. *AIP Conf.Proc.*, 1604:345–352, 2014.
- [66] F. Capozzi, G.L. Fogli, E. Lisi, A. Marrone, D. Montanino, et al. Status of three-neutrino oscillation parameters, circa 2013. *Phys.Rev.*, D89(9):093018, 2014.
- [67] C. Jarlskog. Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Violation. *Phys.Rev.Lett.*, 55:1039, 1985.

- [68] The numi technical design handbook. http://www-numi.fnal.gov/numwork/tdh/tdh_index.html.
- [69] S. Shukla, J. Marriner, and J. Griffin. Slip stacking in the Fermilab Main Injector. *eConf*, C960625:ACC015, 1996.
- [70] The nova technical design report. http://www-nova.fnal.gov/nova_cd2_review/tdr_oct_23/tdr.htm.
- [71] A.G. Abramov, N.A. Galyaev, V.I. Garkusha, J. Huyen, F.N. Novoskoltsev, et al. Beam optics and target conceptual designs for the NuMI project. *Nucl.Instrum.Meth.*, A485:209–227, 2002.
- [72] Raphael Schroeter. Wrong sign neutrino contamination and beam nue content for FHC/RHC beam configurations - Fluka 2011. NOvA internal docdb-13097.
- [73] Chris Backhouse and Ryan Patterson. Library event matching event classification algorithm for electron neutrino interactions in the nova detectors. NOvA internal docdb-11869.
- [74] S. Mufson et al. Liquid scintillator production for the NOvA experiment. NOvA internal docdb-13014.
- [75] A. Norman. The NOvA Data Acquisition System. *J.Phys.Conf.Ser.*, 396:012035, 2012.
- [76] John Cooper. Nova far detector mass. NOvA internal docdb-11905.
- [77] Chris Backhouse. Overview of attenuation corrections. NOvA internal docdb-10533.
- [78] Nick Raddatz. Hit Efficiencies and Cell Voids. NOvA internal docdb-10561.
- [79] John Cooper. Nova near detector mass. NOvA internal docdb-11906.

- [80] Donald Friend and Ernest Villegas. Near Detector Underground Assembly Layout. NOvA internal docdb-8282.
- [81] Chris Backhouse and Alexander Radovic. The Attenuation and Threshold Correction of the NOvA Detectors. NOvA internal docdb-13579.
- [82] Luke Vinton. Calorimetric Energy Scale Calibration of the NOvA Detectors. NOvA internal docdb-13579.
- [83] Alex Himmel. The NuMI beam simulation with Flugg. MINOS internal docdb-6316.
- [84] G. Battistoni et al. The fluka code: description and benchmarking. *AIP Conf.Proc.*, 896:31–49, 2007.
- [85] S. Agostinelli et al. GEANT4 - a simulation toolkit. *Nucl. Instrum. and Meth.*, A506:250, 2003.
- [86] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. and Meth.*, A614:87, 2009.
- [87] Jerome Verbeke Chris Haggmann, David Lange and Doug Wright. Cosmic-ray Shower Library (CRY). http://nuclear.llnl.gov/simulation/doc_cry_v1.7/cry.pdf.
- [88] A. Norman, R. Kwarciany, G. Deuerling, and N. Wilcer. The NOvA Timing System: A System for Synchronizing a Long Baseline Neutrino Experiment. *J.Phys.Conf.Ser.*, 396:012034, 2012.
- [89] E. Niner, P. Adamson, G. Deuerling, R. Kwarciany, H. Meyer, et al. Synchronization of the 14 kTon NO ν A neutrino detector with the Fermilab NuMI beam. *J.Phys.Conf.Ser.*, 513:012028, 2014.
- [90] Neal Wilcer. TDU FPGA Firmware Guide. NOvA internal docdb-5370.

- [91] D.G. Beechy and R.J. Ducar. TIME AND DATA DISTRIBUTION SYSTEMS AT THE FERMILAB ACCELERATOR. *Nucl.Instrum.Meth.*, A247:231–238, 1986.
- [92] Ryan Patterson. Progress on readout, timing, slicing. NOvA internal docdb-4275.
- [93] Chris Backhouse. Improving timing resolution. NOvA internal docdb-6567.
- [94] Brian J. Rebel. Neutrino - Induced Muons in the MINOS Far Detector. 2004.
- [95] Kanika Sachdev. Event Topologies in NOvA Detectors. NOvA internal docdb-9908.
- [96] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996.
- [97] Michael Baird. Slicing Module Comparison Tech Note. NOvA internal docdb-9195.
- [98] Leandro A.F. Fernandes and Manuel M. Oliveira. Real-time line detection through an improved hough transform voting scheme. *Pattern Recognition*, 41(1):299 – 314, 2008.
- [99] Michael Baird. Tech Note for the Multi-Hough Transform. NOvA internal docdb-8241.
- [100] M. Gyulassy and M. Harlander. *Computer Physics Communications*, 66:32–46, 1991.
- [101] M. Ohlsson and C. Peterson. *Computer Physics Communications*, 71:77–98, 1992.
- [102] M. Ohlsson. *Computer Physics Communications*, 71:19–32, 1993.
- [103] R. Fruwirth and A. Strandie. *Computer Physics Communications*, 120:197–214, 1999.
- [104] Mark Messier. Vertex reconstruction based on elastic arms. NOvA internal docdb-7530.

- [105] Fred James and Matthias Winkler. MINUIT User's Guide. <http://seal.web.cern.ch/seal/documents/minuit/mnusersguide.pdf>.
- [106] Erika Catano-Mur. Blessing plots - Nue ND Data/MC. NOvA internal docdb-13809.
- [107] J.C. Dunn. A fuzzy relative of the isodata process and its use in detecting compact, well-separated clusters. *J. Cybern.*, 3:32–57, 1974.
- [108] J.C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York, 1981.
- [109] R. Krishnapuram and J.M. Keller. A possibilistic approach to clustering. *IEEE Trans. Fuzzy Syst.*, 1:98–110, 1993.
- [110] M.S. Yang and K.L. Wu. Unsupervised possibilistic clustering. *Pattern Recognition*, 39:5–21, 2006.
- [111] N.H. Kuiper. Tests concerning random points on a circle. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen, Series A*, 36:38–47, 1960.
- [112] J. Paley et al. Executive summary for the first numu CC analysis. NOvA internal docdb-13641.
- [113] J. Bian et al. EID tech note. NOvA internal docdb-9923.
- [114] C. M. Bishop. *Neural networks for Pattern Recognition*. Clarendon, Oxford, UK, 1998.
- [115] R. Beale and T. Jackson. *Neural Computing: An Introduction*. Adam Hilger, New York, 1991.
- [116] Bruno Zamorano. Technical note: FD timing peak monitoring. NOvA internal docdb-13441.

- [117] Joao Coelho and Barnali Chowdhury. Tech note: Good data selection. NOvA internal docdb-13546.
- [118] Barnali Chowdhury. Near Detector Good Run Selection. NOvA internal docdb-12111.
- [119] Joao Coelho. ND Good Data Selection Improvement. NOvA internal docdb-13213.
- [120] Lisa Goodenough. Discussion on Beam Quality Cuts. NOvA internal docdb-12474.
- [121] Xuebing Bu and Kanika Sachdev. Spill Level Data Quality- Technical Note. NOvA internal docdb-12437.
- [122] Susan Lein. DCM Edge Metric. NOvA internal docdb-13527.
- [123] Tian Xin. Note of Selection Cuts for Nue First Analysis. NOvA internal docdb-13592.
- [124] Gavin Davies. ND cosmics. NOvA internal docdb-13259.
- [125] Chris Backhouse et al. Near Detector data/MC comparisons for the nue analysis. NOvA internal docdb-13587.
- [126] Chris Backhouse. The Library Event Matching nue PID – technote. NOvA internal docdb-13588.
- [127] Kirk Bays. Hadronic E discussion. NOvA internal docdb-13635.
- [128] Kanika Sachdev. Fitting for interaction mode scales in ND. NOvA internal docdb-13678.
- [129] Joseph Lozier. ModularExtrap Technical Note. NOvA internal docdb-12563.
- [130] Louise Suter. NuMu ND selection tech note. NOvA internal docdb-13212.
- [131] Nick Raddatz. ReMId Technical Note. NOvA internal docdb-11206.

- [132] Susan Lein. FA Numu Energy Retune. NOvA internal docdb-13327.
- [133] Sanjib Mishra Hongyue Duyang and Nitin Yadav. Cosmic and Rock Muon Induced EM Showers Technote. NOvA internal docdb-13594.
- [134] Gavin Davies. Position paper: nue systematics. NOvA internal docdb-13579.
- [135] Sarah R. Phan-Budd and Lisa Goodenough. Technical Note on the NOvA Beam Monitoring for 2015 Summer Analysis. NOvA internal docdb-13572.
- [136] Hongyue Duyang et al. A Technote Describing the Derivation and Size of NuMI Flux Uncertainties Used in the First NOvA Analyses. NOvA internal docdb-13584.
- [137] C.N. Chou. The Nature of the Saturation Effect of Fluorescent Scintillators. *Phys. Rev.*, 87:904, 1952.
- [138] Daniel Pershey. Birks-Chou Suppression and Track dE/dx . NOvA internal docdb-13080.
- [139] Daniel Pershey. Birks-Chou Parameter Selection. NOvA internal docdb-13158.
- [140] Kanika Sachdev. dE/dX Comparison in ND Data and MC. NOvA internal docdb-13322.
- [141] Gavin Davies. Pi-zero mass for First Analysis. NOvA internal docdb-13500.
- [142] Daniel Pershey. Michel Electron Calibration Checks. NOvA internal docdb-13519.
- [143] Matthew Tamsett. Calibration Systematic Uncertainties fo the Disappearance Analysis. NOvA internal docdb-13331.
- [144] Jeff Hartnell. NOvA Light/ADC Levels and Thresholds. NOvA internal docdb-13505.
- [145] Chris Backhouse. nue light level/hit efficiency systematic. NOvA internal docdb-13686.

- [146] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. As defined by GENIE in \$GENIE/src/ReWeight/GSystUncertainty.cxx.
- [147] Ji Liu. Rock muon in ND. NOvA internal docdb-13563.
- [148] Nick Raddatz. Fiducial Mass Systematic. NOvA internal docdb-13237.
- [149] Chris Backhouse. The CAFAna framework - technote. NOvA internal docdb-9222.
- [150] C. Bassin G. Laske and G. Masters. The Current Limits of Resolution for Surface Wave Tomography in North America. *EOS Trans AGU.*, F897:81, 2000.
- [151] K. Abe et al. Precise Measurement of the Neutrino Mixing Parameter θ_{23} from Muon Neutrino Disappearance in an Off-Axis Beam. *Phys.Rev.Lett.*, 112(18):181801, 2014.
- [152] Alexander Radovic. Oscillation Parameters for the First NOvA Analyses. NOvA internal docdb-13640.

Evan Niner

e-mail: eniner@indiana.edu

EDUCATION and EMPLOYMENT

Indiana University	Ph.D. , Physics	2010-2015
Indiana University	M.S. , Physics	2010-2011
University of Alabama	B.S. , Physics and Math	2006-2010

RESEARCH

NOvA Experiment (FNAL E929)

- Graduate Research Assistant under the direction of Professor Mark Messier at Indiana University, June 2010 - Present. The experiment is currently running and collecting data after construction was completed in August of 2014. First analysis results are expected to be published this winter.
- Developed tracking and shower reconstruction tool used as input to one of the primary particle identification algorithms in the electron neutrino appearance analysis. This reconstruction also serves as input to a precision momentum fitter being developed by another student for the muon neutrino disappearance analysis.
- Working on performance evaluation and validation of existing particle identification algorithms as well as potential improvements. Studying event selection in the Near Detector and systematic reconstruction errors for the first analysis.
- Served as a primary on-call expert for the timing system May 2013 to present. For the oscillation experiment precise understanding of the time of the beam in both detectors is critical. Duties included developing monitoring tools for timing system performance on a web page, writing software to calculate, store, and retrieve timing calibration constants

from the database and infrastructure to support a more automated recovery of the system from power outages. Role also included serving as one of the liaisons communicating system upgrade requests to the firmware engineers and providing testing feedback.

- Developed a calibration tool using cosmic muons to calibrate timing offsets between different regions of the detector.
- Spent time at the Far Detector in Ash River, MN learning detector construction, then stationed at Fermilab July 2013 to April 2014 to transfer knowledge and work on assembly of the Near Detector located underground at the lab. Played a leading role to install, cable, and commission the timing system and data concentrator modules used in the electronics readout. Other tasks included oil filling and installation of supporting infrastructure.

EXO Experiment (2007-2010)

- Undergraduate Research Assistant under the supervision of Professor Andreas Piepke at the University of Alabama.
- Design, construction, and testing of a mechanical support structure and computerized data acquisition system to be used to calibrate the muon veto system for the EXO neutrinoless double beta decay experiment.
- Calibration of the PMTs and scintillation panels for the system, and assembly and light sealing of completed panels.
- Design, development, assembly, testing, and implementation of calibration panels for the muon veto system after installation at the experiment site at WIPP. This included traveling to WIPP to assemble the system. Wrote a GUI to control data acquisition and HV power of the system and performed initial tests.

Research Internships in Science and Engineering (RISE) Summer Program, Summer 2008

- Intern at Dortmund Technical University with Professor Wolfgang Rhoder working on the

IceCube experiment. Work involved learning the simulation software to update and test code on the new computing grid.

TEACHING

- Fall 2010-Associate Instructor, Department of Physics, Indiana University, Bloomington, IN. Taught introductory lab for Physics 101.
- Fall 2007-Spring 2010 Undergraduate Teaching Assistant, Department of Physics, University of Alabama, Tuscaloosa, AL. Assisted in teaching introductory labs for physics.

AWARDS, LEADERSHIP, and OUTREACH

- Elected president of Young NO ν A student and postdoc organization, June 2014 to Present
- Provided tours of underground laboratory at Fermilab, January to April 2014.
- Graduate student representative on NO ν A shift committee, December 2012-June 2013.
- Attended Excellence in Detectors and Instrumentation Technologies(EDIT) School, Fermilab, IL, February 2012 .
- Department of Physics Graduate Student President, Indiana University, April 2011-April 2012
- Indiana University Graduate Fellowship, 2010-2011
- RISE Program Scholarship, Summer 2008
- Blount Presidential Scholarship, University of Alabama, 2006-2010
- Department of Physics and Astronomy E. Scott Barr Scholarship, University of Alabama, 2006-2010
- National Merit Scholarship, University of Alabama, 2006-2010

PRESENTATIONS and POSTERS

- *Roadmap to the first neutrino oscillation results of the NOvA experiment*, Invited HEP seminar at Stony Brook University, Stony Brook, NY, February 16th, 2015.
- *NOvA Electron Neutrino Appearance Analysis*, Poster presented at XXXIV Physics in Collision, Bloomington, IN, September 17th, 2014.
- *Time Synchronization and Energy Calibration in the NOvA Detector*, Poster presented at XXVI International Conference on Neutrino Physics and Astrophysics (Neutrino 2014), Boston, MA, June 2-7, 2014.
- *Synchronization of the 14 kTon NOvA Detector with the NuMI Beam*, Talk at the 20th International Conference on Computing in High Energy and Nuclear Physics (CHEP), Amsterdam, Netherlands, October 14-18, 2013.
- *Reconstruction Algorithms for Electron Neutrino Appearance in NOvA*, Talk at American Physical Society April Meeting, Denver, CO, April 13-16, 2013.
- *NOvA Electron Neutrino Appearance Measurement*, Poster presented at 45th Annual Fermilab Users Meeting, Fermilab, Batavia, IL, June 12-13, 2012.
- *NOvA: NuMI Off-Axis ν_e Appearance Experiment*, Poster presented at 44th Annual Fermilab Users Meeting, Fermilab, Batavia, IL, June 1-2, 2011.

PUBLICATIONS

- [1] E. Niner, P. Adamson, G. Deuerling, R. Kvarciany, H. Meyer, A. Norman, R. Rechenmacher, P. Shanahan and N. Wilcer. “Synchronization of the 14 kTon NOvA neutrino detector with the Fermilab NuMI beam”, J. Phys. Conf. Ser., Vol. 53, 012028, 2014.
- [2] C. Adams *et al.* [LBNE Collaboration], “The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe”, arXiv:1307.7335 [hep-ex].
- [3] D. Sinclair *et al.* [EXO Collaboration], “Prospects for barium tagging in gaseous xenon”, J. Phys. Conf. Ser., Vol 309, 012005, 2011.
- [4] A. Dobi *et al.* [EXO Collaboration], “A xenon gas purity monitor for EXO”, Nucl. Instrum. Meth. A, Vol. 659, 215, 2011 [arXiv:1106.1812 [hep-ex]].
- [5] M. Montero Diez *et al.* [EXO Collaboration], “A simple radionuclide-driven single-ion source”, Rev. Sci. Instrum., Vol. 81, 113301, 2010 [arXiv:1008.3422 [hep-ex]].
- [6] R. Neilson *et al.* [EXO Collaboration], “Characterization of large area APDs for the EXO-200 detector”, Nucl. Instrum. Meth. A, Vol. 608, 68, 2009 [arXiv:0906.2499 [hep-ex]].