

AN ANALYSIS OF MUON NEUTRINO
DISAPPEARANCE FROM THE NUMI BEAM
USING AN OPTIMAL TRACK FITTER

Michael David Baird

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

September, 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. - Committee Chair

Steven Gottlieb, Ph.D. - 2nd Reader

James Musser, Ph.D. - 3rd Reader

Richard Van Kooten, Ph.D. - 4th Reader

Defense Date: August 14, 2015

Copyright © 2015

Michael David Baird

[REDACTED]

ACKNOWLEDGMENTS

I have taken a rather circuitous route to arrive where I am and the list of people who have led me to this point is semi-infinite in length. This is but a sampling of the people that I have to thank for all of the help along the way. It is by no means a complete list and it is not in any definite order. I'll begin by thanking Cody, the Apple Store employee who helped me transfer to a new computer the contents of my laptop, which had decided to die three weeks before being finished with all of this.

Next I would like to acknowledge the members of my defense committee: Steven Gottlieb, Jim Musser, and Richard Van Kooten, who took the time to read this lengthy work and provided invaluable feedback. The professors in the I.U. NOvA group: Jim Musser (who deserves to be recognized twice), John Urheim, and Stuart Mufson provided me with continual encouragement and advice, for which I am grateful. I have also had the fortune to work with many people on NOvA outside of I.U. who have been generous with their guidance including Peter Shanahan, Ryan Patterson, Gary Feldman, Andrew Norman, Brian Rebel, Jon Paley, Greg Pawloski, Kirk Bays, and Chris Backhouse. Special thanks must also go to the graduate students I have worked closely with over the years at the University of Minnesota: Dominick Rocco, Jan Zirnstein, Susan Lein, Nick Raddatz, and Kanika Sachdev.

The post-docs on our team at I.U. Denver Whittington, Hayes Merritt, and Gavin Davies, need to be thanked for constantly listening to my complaints, helping me with the finer points of my analysis, and forcing me to take breaks for coffee and/or beer when

I needed them. One of my fellow graduate students, Fernanda Psihas, also deserves a special note for taking over my work with the NOvA Data Quality group (which was no small task) so that I could finish my analysis and write this document. And a very special acknowledgement must go to my office mate and good friend, Evan Niner. Together we pushed each other to finish this marathon that we began almost five years ago and I could not have done this without his help.

The greatest influence on my academic success comes of course, from my adviser Mark Messier. He consistently went above and beyond the call of duty to make sure that I was maximizing my potential by sacrificing his own time to help me achieve my goals. Beyond that, he also provided tremendous guidance on a wide variety of non-science topics covering everything from how to maximize my career trajectory to parenting advice. He has been an excellent mentor, adviser, and friend and I will continue to strive to be the kind of scientist that he is.

[REDACTED]

Michael David Baird

AN ANALYSIS OF MUON NEUTRINO DISAPPEARANCE FROM THE NUMI BEAM
USING AN OPTIMAL TRACK FITTER

The NOvA experiment is a long-baseline neutrino oscillation experiment based out of Fermilab National Accelerator Laboratory that uses two liquid scintillator detectors, one at Fermilab (the “near” detector) and a second 14 kton detector in northern Minnesota (the “far” detector.) The primary physics goals of the NOvA experiment are to measure neutrino mixing parameters through both the ν_μ disappearance and ν_e appearance channels using neutrinos from the newly upgraded NuMI beam line. The NOvA ν_μ disappearance analysis can significantly improve the world’s best measurement of $\sin^2 \theta_{23}$. This analysis proceeds by using the measured ν_μ charged-current energy spectrum in the near detector to predict the spectrum in the far detector, and comparing this to the measured spectrum to obtain a best fit for the oscillation parameters $\sin^2 \theta_{23}$ and Δm_{32}^2 . Since this fit is governed by the shape of the energy spectrum, the best fit will be maximized by obtaining the best possible energy resolution for the individual neutrino events.

This dissertation describes an alternate ν_μ disappearance analysis technique for the NOvA experiment, based on the idea that estimating the energy resolution of the individual events will allow them to be separated into different energy resolution samples in order to improve the final fit. This involves using an optimal tracker to reconstruct particle tracks and momenta, and multivariate methods for estimating the event energies and energy resolutions. The data used for this analysis was taken by the NOvA experiment from February 2014 to May 2015, representing approximately 3.52×10^{20} protons on target from the

NuMI beam. The best fit oscillation parameters obtained by this alternate technique are $|\Delta m_{32}^2| = 2.49_{-0.17}^{+0.19} [\times 10^{-3} \text{eV}^2]$ and $\sin^2 \theta_{23} = 0.51 \pm 0.08$ which is consistent with the hypothesis of maximal mixing, and with the results from T2K and MINOS+ published in 2015.

CONTENTS

1	A Brief History of Neutrinos and Neutrino Oscillations	1
1.1	The Absence of Evidence	1
1.2	The Evidence of Absence	2
1.3	First Detection of the Neutrino	4
1.4	Missing Solar Neutrinos	5
1.5	Confirmation of Neutrino Oscillations	6
1.6	The Importance of Neutrino Physics	8
2	Neutrino Oscillations and Summary of the Experimental Status of the Field	11
2.1	Neutrino Eigenstates and the PMNS Mixing Matrix	11
2.2	Three-Flavor Neutrino Oscillations in a Vacuum	13
2.3	Matter Effects	16
2.4	Experimental Results	18
2.4.1	Current Global Best Fits	18
2.4.2	Remaining Unknowns	19
2.5	Measurements of $ \Delta m_{32}^2 $ and $\sin^2 2\theta_{23}$	20
2.5.1	Results From Super-K	21
2.5.2	Recent Results from T2K and MINOS+	23
2.5.3	Results from OPERA	25

2.6	Physics Goals of the NOvA Experiment	25
2.6.1	The ν_μ Disappearance Measurement	26
2.6.2	The ν_e Appearance Measurement	29
2.6.3	Improving the Measurement of θ_{23}	31
3	Design of the NOvA Experiment	33
3.1	The NuMI Beam	33
3.2	The NOvA Detectors	37
3.2.1	Physical Design	37
3.2.2	The Far Detector	39
3.2.3	The Near Detector	40
3.3	The Smallest Unit of NOvA Readout	42
3.3.1	The NOvA Cell	42
3.3.2	The Avalanche Photodiode	44
3.4	The Data Acquisition System	46
3.4.1	Overview of the DAQ	46
3.4.2	The NuMI Triggers	47
3.4.3	The Timing System	47
4	Performance of the NOvA Detectors	49
4.1	Online Monitoring	50
4.1.1	The OnMon Producer	50
4.1.2	The OnMon Viewer	53
4.1.3	OnMon Metrics	57
4.2	Nearline Monitoring	59
4.2.1	Nearline Processing	59

4.2.2	Nearline Metrics	62
4.3	Data Quality Selection	66
4.3.1	Subrun Level Cuts	66
4.3.2	Event Level Cuts	70
4.4	Gallery of OnMon and Nearline Plots	72
5	Simulating NOvA Data	82
5.1	Simulating the NuMI Beam	82
5.2	Simulating Detector Events	83
5.3	Simulating the NOvA Readout Response	85
5.4	Fine Tuning the Simulations With Data	86
5.4.1	Tuning the Birks-Chou Suppression Parameters	86
5.4.2	Modeling the APD Sag Effect	87
5.5	Event Displays	88
6	Calibration and Event Reconstruction	94
6.1	Calibration	94
6.1.1	Attenuation Calibration	95
6.1.2	Absolute Energy Calibration	97
6.1.3	Drift Calibration	98
6.1.4	Timing Calibration	98
6.2	NOvA Reconstruction Challenges	101
6.3	The Slicer4D Clustering Algorithm	102
6.3.1	The Expanding Density-Based Clustering Method	103
6.3.2	The Neighbor Score Function	104
6.3.3	Algorithm Speed Optimization	105

6.3.4	Tuning the Parameters	105
6.3.5	Results and Performance	112
6.4	The Multi-Hough Transform	115
6.4.1	Filling the Hough Space Parameter Map	116
6.4.2	Calculating the Threshold	117
6.4.3	Refining the Peak Value	118
6.4.4	The Multi-Hough Loop	118
6.4.5	Results and Performance	119
6.5	The Elastic Arms Method	122
6.5.1	The Algorithm	123
6.5.2	Results and Performance	125
6.6	Particle Clustering with Fuzzy-K	126
6.6.1	The Algorithm	127
6.6.2	Results and Performance	129
6.7	Particle Tracking with Break Point Fitter	129
6.7.1	The Optimal Track Fit Expression	130
6.7.2	Adapting the Optimal Track Fit Expression for NOvA	132
6.7.3	Results and Performance	136
7	The Break Point Fitter Analysis Suite	145
7.1	Summary of Motivations and Methods	145
7.2	Muon Identification	146
7.2.1	The Muon Track dE/dx Log-Likelihood Calculation	147
7.2.2	Input Variables to the Muon ID	155
7.2.3	The Muon Identifier	158
7.3	The BPF Energy Estimator	162

7.3.1	The Energy Estimation Method	163
7.3.2	Input Variables to the Energy Estimator	165
7.3.3	Examining Bias in the Energy Estimator	169
7.3.4	Energy Estimator Performance	172
7.4	Using Energy Resolution to Improve the Final Contours	179
7.4.1	Determining the Energy Resolution Binning	179
7.4.2	Generating the Sensitivity Contours	180
7.4.3	Exploring a Set of Toy Experiments	185
8	Results of The ν_μ Charged-Current Disappearance Analysis Using the Break Point Fitter Method	192
8.1	Event Selection	192
8.1.1	Contained Event Selection	193
8.1.2	Signal Event Selection	195
8.1.3	Rejection of Cosmic Rays in the Far Detector	196
8.1.4	Break Point Fitter Quality Cut	198
8.1.5	Event Selection Summary	199
8.2	Summary of the Selected Data Sets	200
8.2.1	Far Detector Data	200
8.2.2	Near Detector Data	200
8.2.3	Simulated Data Sets	201
8.3	Data Simulation Comparisons	202
8.3.1	Results from the Kalman Tracker	202
8.3.2	Results from the Break Point Fitter Tracker	205
8.4	A Discussion on Hadronic Energy	207
8.5	The NOvA Analysis Framework	212

8.5.1	Background Estimates	212
8.5.2	Extrapolation	214
8.5.3	Systematic Uncertainties	217
8.5.4	Sensitivity Contours	220
8.6	Final Results	220
8.6.1	The Selected Data Events	220
8.6.2	Oscillation Parameter Fits	223
8.6.3	Concluding Remarks	228
	Bibliography	232
	Curriculum Vitae	

LIST OF FIGURES

2.1	Feynman diagram depicting the coherent forward scattering of an electron neutrino off of an electron via a charged current interaction. This interaction sets electron neutrinos apart from the other two flavors for neutrinos traveling through a macroscopic object such as the Earth, producing the so called “matter effect.”	16
2.2	Top: The expected distributions for sub-GeV and multi-GeV single-ring electron- and muon-like events. The boxes indicate a prediction based on a model without oscillations and the solid line indicates the best fit to a two flavor oscillation model. Bottom: Measurement contours for 68, 90, and 99% confidence intervals showing the best fit for $ \Delta m^2 $ and $\sin^2 2\theta$ with a two flavor oscillation model for both the θ_{zenith} and L/E analyses. These plots were taken from reference [22].	22
2.3	Left: Sensitivity contours taken from [24] showing the current T2K measurements of $\sin^2 \theta_{23}$ and $\sin^2 \theta_{13}$ vs. Δm_{32}^2 . Right: Sensitivity contours taken from [25] showing the most recent MINOS+ measurement of $\sin^2 \theta_{23}$ vs. Δm_{32}^2 as compared to an earlier result from T2K published in 2014. . .	24
2.4	Simulated true energy spectra for contained ν_μ CC events in the NOvA far detector. The exposure is assuming 3 years of nominal data taking with a 14 kiloton detector while running the NuMI beam in ν mode.	27

2.5	<p>Top: An example of the expected true energy spectra for contained ν_μ CC events in the NOvA far detector under two different assumptions for $\sin^2 \theta_{23}$. The exposure is assuming 3 years of nominal data taking with a 14 kiloton detector while running the NuMI beam in ν mode. Bottom: Example of the expected sensitivity contours for the full NOvA exposure (3 years running the NuMI beam in ν mode plus 3 years running in $\bar{\nu}$ mode) for the two values of $\sin^2 \theta_{23}$ shown in the top plot.</p>	28
2.6	<p>Examples of bi-probability plots for different values of $\sin^2 2\theta_{23}$. The top plot assumes θ_{23} is maximal and the bottom plot assumes it is non-maximal, with two sets of ellipses for each of the two possible octants. The star on each plot indicates the “best possible scenario” for making the strongest statements about δ and the matter hierarchy.</p>	30
2.7	<p>Examples sensitivities to the octant of θ_{23} as functions of δ, assuming $\sin^2 2\theta_{23} = 0.95$ and $\theta_{23} > 45^\circ$.</p>	32
3.1	<p>The distribution of times for all hits seen with the NOvA near detector with respect to a trigger time issued 218 μsec before the start of the beam spill. This plot includes data from 525,899 NuMI spills taken between Feb. 14, 2015 and Feb. 17, 2015.</p>	34
3.2	<p>A cartoon representation of the NuMI beam.</p>	35
3.3	<p>Left: Neutrino flux as a function of the energy of the parent pion for different off-axis angles. Right: Neutrino energies as a function of the parent pion energy for different angles with respect to the original pion direction.</p>	36
3.4	<p>Expected charged current ν_μ event spectra for the low-energy (left) and medium-energy (right) horn configurations, calculated prior to oscillations for a detector 810 km from the beam source.</p>	37

3.5	Both NOvA detectors shown to scale with each other (and with a cartoon person.) The inset shows the alternating, orthogonal plane structure. The far and near detectors are approximately 60 m and 16 m long respectively. .	38
3.6	Three simulated events showing NOvA's ability to distinguish muons, electrons, and π^0 s. The distinguishing feature of the electron over the muon is the shower-like nature of its track, and the distinguishing feature of the π^0 over the electron is the gap between the event vertex and the start of the shower.	39
3.7	Overhead view of the location of the NOvA near detector in the MINOS underground area.	41
3.8	Technical drawing of the near detector showing the smaller size of the muon catcher (right end.) For picture clarity, only some planes in the main body of the detector have been drawn.	42
3.9	Left: Schematic of the PVC cell filled with liquid scintillator. A charged particle induces scintillation light (blue line) that bounces around until it is absorbed by the transmission fiber (green line.) The light is then transmitted through the fiber to the APD. Right: Cross section of an extruded PVC group of 16 cells. The dimensions shown are the exterior dimensions of the cell.	43
3.10	Left: The custom designed NOvA APD showing the 32 pixel array. Right: Close up of the end of the module that interfaces with the APD in which all 32 fibers have been bundled together. In this picture, an individual fiber has been illuminated.	44
4.1	The OnMon viewer.	54

4.2	Example distribution of far detector subruns in terms of the median MIP hit rate variable. “Other” refers to the file time stamp cut.	69
4.3	Example distribution of far detector subruns in terms of the 2D track fractions variable. “Other” refers to the file time stamp cut.	69
4.4	OnMon Plots - Hit map for the NOvA near detector displaying the average pixel hit rate for each FEB (top.) Hit map representing DCM 2 in diblock (DB) 2 showing the hit rates for each individual pixel (bottom.) The large boxes represent a single DCM and the small boxes are an individual FEB. White spaces are either non-reporting or uninstrumented hardware.	72
4.5	OnMon Plots - FEB rates vs. time (top) and total number of microslices per data event vs. time (bottom.)	73
4.6	OnMon Plots - All alerts reported vs. time (top) and a DCM-level hit map for the “FEB buffer empty” alert (bottom.)	74
4.7	Nearline Plots - FEB hit rate spectra vs. time (top) and average number of noise hits per subrun vs. time (bottom.)	75
4.8	Nearline Plots - Number of active FEBs per subrun vs. time (top) and FEB dropout counts (bottom.)	76
4.9	Nearline Plots - Distribution of the number of slices per spill (top) and average number of slices per spill vs. time (bottom.)	77
4.10	Nearline Plots - Calibration and NuMI trigger rates vs. time (top) and track fractions vs. time (bottom.)	78
4.11	Nearline Plot - Track start positions.	79
4.12	Nearline Plots - NuMI timing peak (top) and NuMI timing peak vs. time (bottom.)	80
4.13	Nearline Plot - PE distribution for slices and tracks.	81

4.14	Nearline examples of bad detector performance. They are (clockwise from the top left) FEB hit rate spectra vs. time, track fractions vs. time (3D tracks in red, 2D in blue, contained 3D in green), percent empty spills, and timing peak distribution (from OnMon.)	81
5.1	The GENIE default cross section for ν_μ CC interactions on an isoscalar target, taken from [44]. The shaded band indicates an estimated uncertainty on the free nucleon cross section.	84
5.2	A comparison of the ADC distribution for hits on cosmic tracks that cross more than 20 planes between simulations made with CRY and far detector data. Cosmic tracks are produced by the tracking algorithm described in section 4.2.1.	86
5.3	Results showing the readout for all 32 channels on a single APD when light was injected into just channel 11. Channel response was measured in units of ADC as a function of time in units of 125 ns. The plot on the right zooms in on the y -axis showing the sag effect on the other 31 channels. This plot was taken from [56].	87
5.4	An event display showing cosmic ray data taken with the far detector. Hits are colored and scaled by charge.	90
5.5	An event display showing simulated cosmic ray data with CRY for the far detector. Hits are colored and scaled by charge.	91
5.6	An event display showing NuMI beam data taken with the near detector. Hits are scaled by charge and colored by time.	92
5.7	An event display showing simulated NuMI beam data for the near detector. Hits are scaled by charge and colored by time.	93

6.1	<p>Left: Attenuation fit for a specific cell with cosmic-ray data taken using the near detector. The ends of the cell are indicated by the vertical dashed lines. Data outside of these regions comes from various track reconstruction failures and is not used in the fit. The attenuation fit is shown in blue. Right: Example plot of PECorr/cm vs. distance from the track end, generated from far detector cosmic-ray data used to compute the absolute energy scale.</p>	96
6.2	<p>An example of the drift calibration as applied to far detector data for Diblock 3, DCM 6, FEB 24. The mean channel response and error on the mean is shown in black with the fit for the drift correction shown in red.</p>	99
6.3	<p>Timing resolution plots determined from far detector (top two plots) and near detector (bottom two plots) data. The two plots on the left show the single-point readout mode and the two plots on the right show the multi-point readout mode.</p>	100
6.4	<p>Distributions of the FOM from the Monte Carlo method for tuning Slicer4D vs. two of the Slicer4D configurable parameters, ϵ (top) and D_{pen} (bottom). The plot colors correspond to the number of parameter set combinations with a specific value of ϵ or D_{pen} that resulted in that particular FOM. Both plots were made with near detector simulations that included rock events and were used in the initial grid-search method for tuning the Slicer4D parameters.</p>	108
6.5	<p>Distance to the third closest neighbor for hits within the same view used for determining the Slicer4D parameter D_{pen} (left) and for determining OppViewPIPen (right). Lines are superimposed on each plot illustrating the “kink” used to choose the values for D_{pen} and OppViewPIPen.</p>	109
6.6	<p>Event display (top) for the data event used to generate the 3-dist plot shown on the bottom.</p>	111

6.7	Event display from a far detector data event prior to (first picture) and after (second picture) slicing has been applied.	113
6.8	Completeness and purity distributions for near detector GENIE events with at least 10 hits in each view.	114
6.9	Completeness and purity distributions for far detector GENIE events with at least 10 hits in each view.	115
6.10	Completeness and purity distributions for far detector CRY events with at least 10 hits in each view.	115
6.11	An example of the Multi-Hough algorithm. Shown are the results after four iterations through the peak finding loop. The figure on the right is the corresponding Hough space map for each event on the left.	120
6.12	Perpendicular distance from the first (left) and second (right) Hough lines to the true vertex in the X view (area normalized to 1).	121
6.13	Dot Product with the first (left) and second (right) Hough lines and the best matched MC particle trajectory in the X view (area normalized to 1.) . . .	122
6.14	Reconstructed far detector neutrino candidate event showing the Hough lines and the Elastic Arms vertex.	125
6.15	3D Event vertex resolution with the Elastic Arms algorithm.	126
6.16	Reconstructed far detector neutrino candidate event showing the matched, three-dimensional clusters produced by the Fuzzy-K method.	129
6.17	Distribution of completeness values as a function of visible energy for the clusters formed by the primary lepton in a CC event produced by the Fuzzy-K method.	130

6.18	Schematic representation of the track model used by BPF. The measurements x_i and their uncertainties σ_i are shown as dots with error bars. The track is allowed to scatter through the angles α_J at the scattering planes indicated by the dashed vertical lines.	131
6.19	This plot shows the location of the scattering planes for the muon fit assumption (tall blue lines) and the proton fit assumption (short purple lines) for a simulated 2 GeV/ c muon with a track length of approximately 850 cm. The scattering planes locations are measured along the track length s	136
6.20	Top: Reconstructed far detector neutrino candidate event showing the tracks produced by the BPF algorithm. Only the tracks produced by the muon fit assumption are shown. Bottom: Zoom in on the track trajectory shown in the gray box in the top figure. Cells borders are outlined in light gray and cell hits are color coded by the charge deposited according to the scale shown on the bottom right of the top figure.	138
6.21	An example of the event vertex being pulled forward along the track by the Elastic Arms algorithm, justifying the “half-p” cut. The event shown is a simulated single muon with a momentum of 2 GeV/ c traveling from left to right. The Elastic Arms vertex is drawn as a blue “X” with 2 back-to-back Fuzzy-K 3D prongs. The blue prong contains most of the muon and the short green prong points back towards the true muon origin.	139
6.22	Momentum fractional bias plots for four simulated single muon samples. The mean and RMS values for these distributions are listed in table 6.6.	140
6.23	Momentum fractional bias plots for the 1.0 GeV/ c simulated single pion and single proton samples.	142

6.24	Examples of the types of pion reconstruction failures that led to poor estimates of pion momenta. In the top figure, the pion (pink track) has a strong interaction knocking a proton (purple track) out of a nucleus which deposits the majority of the energy in the Fuzzy-K 3D prong (blue outline). In the bottom figure, the pion has had several hard scatters causing the Elastic Arms vertex (blue “X”) to be placed in the wrong spot.	143
6.25	Distribution of reconstructed track angles for five samples of single muons all generated at 5 degree angles to the detector z -axis. The mean and RMS values for these distributions are listed in table 6.7.	144
6.26	Distribution of reconstructed track angles for three samples of single 2 GeV/ c muons, generated at 30, 60, and 80 degree angles to the detector z -axis. The mean and RMS values for these distributions are listed in table 6.7.	144
7.1	Stopping power for μ^+ in copper. Equation 7.1 is accurate to within a few percent in the region labeled “Bethe.” This plot was taken from reference [8].	148
7.2	Computed values of dE/dx vs. muon momentum (instead of $\ln(\beta\gamma)$) for simulated 2 GeV/ c muons in the upper-West quadrant of the far detector (top left) and in the bottom-East quadrant (top right.) The bottom plot shows the bin by bin asymmetry computed as (bottom-East - top-West)/(bottom-East + top-West).	151
7.3	Computed values of dE/dx vs. muon momentum (instead of $\ln(\beta\gamma)$) for simulated 2 GeV/ c muons in the far detector with shallow tracks (top left) and steep tracks (top right.) The bottom plot shows the bin by bin asymmetry computed as (steep - shallow)/(steep + shallow).	152

7.4	Distribution of computed dx values for the steep track sample (red) and a sample of tracks from 100 reconstructed GENIE events. The dividing line for the two dx ranges was chosen to be 8 cm since this broke the steep track sample up into two distinct populations and most GENIE tracks had cell path length less than this value. Both plots are normalized to unit area. . .	153
7.5	Top: Plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, generated from 10 GeV/ c muon tracks reconstructed with the BPF muon fit assumption. Bottom: Plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, generated from 1 GeV/ c proton tracks reconstructed with the BPF proton fit assumption.	154
7.6	Normalized plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, created by combining the muon and proton results together.	155
7.7	Left: Distribution of dE/dx log-likelihood values for muon (signal) and non-muon (background) tracks. Right: Distribution of BPF muon track fit total χ^2 per degree of freedom values. All distributions are normalized to unit area.	156
7.8	Left: Distribution of track lengths for muon (signal) and non-muon (background) tracks. Right: Distribution of track/prong hit ratios. All distributions are normalized to unit area.	157
7.9	An example of the phase space density-map for a kNN with 2 input variables, x_1 and x_2 (taken from [79].) The filled circles are the signal events and the open circles are the background events. The star represents an unclassified event with the “nearest neighborhood” drawn in a circle around it. In this example, the number of signal and background events within the circle is roughly the same. Therefore $P_s \approx 0.5$	160

7.10	<p>Top: Distribution of BPF muon ID values for signal and background tracks (normalized to unit area.) Bottom: Distribution of BPF muon ID values for background tracks only, broken down into the particle that contributed the most energy to the track.</p>	161
7.11	<p>Distributions for the six input variables to the BPF energy estimator. Comparisons of these variables to data are discussed in chapter 8.</p>	166
7.12	<p>Left: Example of dE/dx values calculated by the <code>TrackCleanUpAlg</code> algorithm along the muon track used to determine the vertex region. This plot was taken from [81]. Right: Distribution of hadronic energy contamination computed by the <code>TrackCleanUpAlg</code> algorithm for muon tracks produced by BPF and identified by the BPF muon ID. This plot was made from a small sample of contained, true ν_μ CC events, simulated for the far detector with GENIE.</p>	169
7.13	<p>Reconstructed energy spectra for the three energy estimators from an unoscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.</p>	173
7.14	<p>Reconstructed energy spectra for the three energy estimators from a “minimally” oscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.</p>	174

7.15	Reconstructed energy spectra for the three energy estimators from a “maximally” oscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.	175
7.16	Top: Energy fractional bias (E_{FB}) for simulated far detector events oscillated with the “minimal” parameters and reconstructed with kNN-2. The blue curve is for all events and the green curve is only the NC sample. Bottom: Distribution of $\Delta E/\sigma_E$ for the same events. The means and RMSs for both plots and the Gaussian fits (red curves) are shown in table 7.3.	176
7.17	Top: Energy fractional bias (E_{FB}) for simulated far detector events oscillated with the “maximal” parameters and reconstructed with kNN-2. The blue curve is for all events and the green curve is only the NC sample. Bottom: Distribution of $\Delta E/\sigma_E$ for the same events. The means and RMSs for both plots and the Gaussian fits (red curves) are shown in table 7.3.	177
7.18	Left: Distribution of values for σ_E for a sample of simulated far detector events. Right: Distribution of values for σ_E/E_{reco} for a sample of simulated far detector events. All events passed a containment and basic ν_μ CC event selector cuts and had E_{reco} between 0.5 and 3.5 GeV. Both plots are normalized to unit area.	180
7.19	2D distributions of the reconstructed event energies using 1 bin, 3 bins, and 5 bins for σ_E/E_{reco} . The cuts on σ_E/E_{reco} to separate these samples are listed in table 7.4. All events passed a containment and basic ν_μ CC event selector cuts.	182

7.20	Example sensitivity contours for the true “minimal” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the $1 \sigma_E/E_{reco}$ bin (black), $3 \sigma_E/E_{reco}$ bins (blue), and $5 \sigma_E/E_{reco}$ bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.	183
7.21	Example sensitivity contours for the true “maximal” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the $1 \sigma_E/E_{reco}$ bin (black), $3 \sigma_E/E_{reco}$ bins (blue), and $5 \sigma_E/E_{reco}$ bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.	184
7.22	Example sensitivity contours generated for the true “alternate” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the $1 \sigma_E/E_{reco}$ bin (black), $3 \sigma_E/E_{reco}$ bins (blue), and $5 \sigma_E/E_{reco}$ bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.	185
7.23	Examples of the 90% sensitivity contours from generating 1000 toy experiments assuming true maximal oscillation parameters. For each plot, the black curve/dot represents the $1 \sigma_E/E_{reco}$ bin fit, and the red curve/dot represents the $3 \sigma_E/E_{reco}$ bin fit. The blue dot is the true oscillation parameters.	188
7.24	The two metrics “ $\Delta \sin^2 2\theta_{23}$ Errors” and “ $\Delta \sin^2 2\theta_{23}$ True” for the 1000 toy experiments generated assuming true “maximal” oscillation parameters.	189

7.25	Examples of the 90% sensitivity contours from generating 1000 toy experiments assuming true minimal oscillation parameters. For each plot, the black curve/dot represents the $1 \sigma_E/E_{reco}$ bin fit, and the red curve/dot represents the $3 \sigma_E/E_{reco}$ bin fit. The blue dot is the true oscillation parameters. . . .	190
7.26	The two metrics “ $\Delta \sin^2 2\theta_{23}$ Errors” and “ $\Delta \sin^2 2\theta_{23}$ True” for the 1000 toy experiments generated assuming true “minimal” oscillation parameters. . .	191
8.1	Distribution of slice event IDs computed by the ReMId algorithm for GENIE simulated events in the far (top plot) and near (bottom plot) detectors. For the far detector, the ν_μ CC events were oscillated using $ \Delta m_{32}^2 = 2.4 \times 10^{-3}$ [eV ²] and $\sin^2 \theta_{23} = 0.5$. Both distributions were normalized to the expected far and near exposures for this analysis of 3.52×10^{20} and 1.65×10^{20} POT. . . .	197
8.2	Figure of merit, defined as S/\sqrt{B} , as a function of the cut value applied to the cosmic rejection BDT output, for different cuts on the ReMId event ID value.	198
8.3	Far detector mass and integrated exposure (mass \times POT) as a function of time. Periods of running with different configurations are evident as is the scheduled accelerator shutdown September/October 2014. This is scaled to an approximate fiducial mass of 11.4 kilotons [91].	201
8.4	Distributions of track variables from near detector data and simulations for the track in the slice with the highest ReMId value. Only tracks from slices that passed the near detector containment cuts are shown. Plots were taken from [85].	203

8.5	<p>Top Left: Track length distributions for the Kalman track with the highest ReMId value for near detector data and simulation. Top Right: Distribution of the number of hits on the track with the highest ReMId value in near detector data and simulation. Only tracks from slices that passed the near detector ν_μ CC event selection cuts are shown. Plots were taken from [93]. Bottom: Distribution of ReMId values for all contained slices in near detector data and simulation. Plot taken from [85].</p>	204
8.6	<p>Distributions of variables associated with hadronic energy. Hadronic hits are defined as hits in a slice not on the primary muon track. Only tracks from slices that passed the near detector containment and ν_μ CC event selection cuts are shown. Plots were taken from [85].</p>	205
8.7	<p>Distributions of reconstructed neutrino energies for slices passing the ν_μ CC event selection cuts in near detector data and simulation. Plot was taken from [93].</p>	206
8.8	<p>Distributions of the six input variables to the BPF energy estimator in near detector data and simulations.</p>	207
8.9	<p>Zoomed versions of three of the input variables to the BPF energy estimator showing comparisons between data and simulations. The top two plots are both distributions for the variable E_{3D}.</p>	208
8.10	<p>Top: Energy spectra generated by the BPF energy estimator for near detector data and simulation. Bottom Left: BPF “hadronic energy” taken as the sum of the E_{3D} and E_{remain} variables. Bottom Right: Distribution of BPF muon ID values for the BPF track in each slice with the highest BPF muon ID value. All plots include only the slices that pass the ν_μ CC selection criteria.</p>	209

8.11	<p>Top: Distributions of the E_{had} variable in data and simulations, showing the effect of scaling this variable up by 21% in the data. Bottom: Showing the reconstructed neutrino energies for data events with E_{had} scaled up by 21% as compared to the unaltered simulated spectra. All plots include only the slices that pass the ν_μ CC selection criteria.</p>	210
8.12	<p>Showing the reconstructed neutrino energies with the BPF energy estimator for data events scaled by the ratio of the shifted to the unshifted reconstructed energy computed by equation 8.1, as compared to the unaltered simulated spectrum. Only slices that pass the ν_μ CC selection criteria have been included.</p>	212
8.13	<p>Top: Reconstructed neutrino energy spectra from the near detector simulations for all ν_μ CC selected events (black) and background events (blue) in the ν_μ CC sample as determined by truth. Bottom: Two-dimensional energy vs. energy resolution plot for only the background events in the ν_μ CC sample, showing that most background events are put into the “worst” energy resolution bin.</p>	214
8.14	<p>Top: Reconstructed neutrino energy spectra for the selected ν_μ CC events in the far detector. Black and red represent the expected true ν_μ CC and NC events determined from the simulations, and blue represents the expected cosmic ray spectrum determined from out of time data. The true ν_μ CC events were oscillated using $\Delta m_{32}^2 = 2.4 \times 10^{-3} \text{ [eV}^2\text{]}$ and $\sin^2 \theta_{23} = 0.5$. Bottom: Unscaled two-dimensional energy vs. energy resolution plot for only the 51 background cosmic ray events determined from the out of time data.</p>	215

8.15	Examples of 2 of the 51 out of time events selected by the ν_μ CC event selection criteria from the far detector data. These events were used to estimate the energy spectrum of the cosmic ray background in the ν_μ CC analysis sample. Both events occur near the top of the detector and are likely caused by neutrons entering the detector from above.	216
8.16	Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 2.62$ GeV, $\sigma_E = 0.11$ GeV (4.2%), $P_\mu = 2.33$ GeV/ c , and a BPF event ID = 0.98. . . .	221
8.17	Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 4.99$ GeV, $\sigma_E = 0.4$ GeV (8.0%), $P_\mu = 1.83$ GeV/ c , and a BPF event ID = 1.0. . . .	222
8.18	Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 2.35$ GeV, $\sigma_E = 0.11$ GeV (4.7%), $P_\mu = 2.11$ GeV/ c , and a BPF event ID = 1.0. . . .	223
8.19	Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 0.91$ GeV, $\sigma_E = 0.08$ GeV (8.7%), $P_\mu = 0.82$ GeV/ c , and a BPF event ID = 0.98. . . .	224

8.20	Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 7.68$ GeV, $\sigma_E = 0.66$ GeV (8.6%), $P_\mu = 2.38$ GeV/ c , and a BPF event ID = 0.99. This event, while well reconstructed, was not used in the final fit since $E_\nu > 5.0$ GeV.	225
8.21	Selected far detector data events from the ν_μ CC analysis sample likely to be cosmic rays.	226
8.22	Top: The reconstructed energies for the 35 far detector events with the predicted spectrum in red from the near detector extrapolation using the best fit values for the oscillation parameters. Bottom: Two dimensional reconstructed energy spectra for the 35 selected far detector events.	227
8.23	Final measurement contours for the BPF ν_μ CC disappearance analysis. The blue contours were made including the systematic uncertainties and the red contours without them. The dashed (solid) lines are the 1- σ (90%) confidence intervals.	228
8.24	Final measurement contours for the BPF ν_μ CC disappearance analysis shown with the most recent results from T2K [24] and MINOS+ [25]. All contours represent 90% confidence intervals.	229
8.25	χ^2 distributions for the two oscillation fit parameters $ \Delta m_{32}^2 $ (left) and $\sin^2 \theta_{23}$ (right.)	230
8.26	Final measurement contours using each of the three BPF energy estimators described in section 7.3.3. The green / blue / red contours were created using the unoscillated / “minimal” / “maximal” training assumption. The dashed (solid) lines are the 1- σ (90%) confidence intervals. These contours do not include systematic uncertainties.	231

LIST OF TABLES

2.1	The best fit values for the neutrino oscillation parameters within a three flavor model. Shown (when appropriate) are the fits assuming normal hierarchy (NH) and inverted hierarchy (IH). Within the OMSD assumption, measurements for $ \Delta m_{31}^2 $ and $ \Delta m_{32}^2 $ have been combined into one measurement of $ \Delta m^2 $. These values have been taken from [8, 21].	19
4.1	Percentage of subruns for both detectors that failed each of the subrun-level data quality cuts. The percentages do not add up to the totals due to subruns failing multiple cuts.	70
6.1	Input parameters used by Slicer4D.	106
6.2	Parameter values used by Slicer4D for both the NOvA near and far detectors.	112
6.3	Completeness and purity numbers (expressed as percents) for near and far detector simulated neutrino interactions (GENIE) and far detector simulated cosmic rays (CRY.) Shown are all slices and all slices passing the requirement of at least 10 hits per view.	114
6.4	Perpendicular distance from the Hough lines to the true vertex in the X-view.	121
6.5	Dot product between the Hough lines and the best matched MC particle trajectory in the X-view. The numbers in the underflow bin represent unmatched lines.	123

6.6	The mean and RMS values for the fractional bias distributions shown in figures 6.22 and 6.23.	140
6.7	The mean and RMS values for the reconstructed track angle distributions shown in figures 6.25 and 6.26.	142
7.1	Percentages for the particle that contributed the most energy to the track selected by the BPF muon ID to be the most “muon-like” track within a slice. Roughly one third of all events are NC. The ReMId selection algorithm is discussed in section 8.1.2.	162
7.2	χ^2 values computed between the true and reconstructed energy spectra for the plots shown in figures 7.13, 7.14, and 7.15. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively.	172
7.3	Shown are the numbers corresponding to the plots in figures 7.16 and 7.17. The variables “plot mean/RMS” were computed from all data in the entire plot, and “fit mean/RMS” are the parameters from the Gaussian fit to the central region. For the plots of E_{FB} , the fit range was from -0.2 to 0.2 and for the $\Delta E/\sigma_E$ plots, the fit range was from -2.0 to 2.0.	178
7.4	Showing the energy fractional bias RMS values for the reconstructed events separated into 1, 3, and 5 samples by σ_E/E_{reco}	181

CHAPTER 1

A Brief History of Neutrinos and Neutrino Oscillations

1.1 The Absence of Evidence

The physicist views the universe as a very large collection of particles that do nothing but exchange other particles back and forth while traveling through space-time. Specifically, all of the matter around us is composed of or can be built from a list of 12 fermions (6 quarks and 6 leptons) that interact through fields mediated by a short list of bosons: the photon, the W and Z particles, the gluons, and of course, the Higgs. Together, this list of fermions and the descriptions of how they interact through these bosons is what we (somewhat boringly) refer to as “The Standard Model” [1]. It is known that the Standard Model is an incomplete theory since it does not include a description of things such as dark matter, dark energy, gravity, or in fact the phenomenon of neutrino oscillations (the subject of this thesis) in which neutrinos are observed to change back and forth between their three types. However, it has achieved great success by accurately explaining how particles interact and correctly predicting the existence of many particles that have been discovered in the past decades.

The neutrino (a member of the lepton family) has proven itself to be one of the most interesting Standard Model particles. As described below, we have discovered that there are (at least) three species of them and that they interact very weakly with other matter. When we say that they interact weakly with matter, we mean that they interact ONLY

through the so called “weak nuclear” force (and of course gravity as all particles do) which is responsible for certain reactions that allow one particle species to turn into another through a decay process. The weak force is so weak (roughly 100 million times weaker than the electromagnetic force) that neutrinos can easily pass through vast distances of solid matter without leaving a trace. For example, several trillion neutrinos created by our Sun pass through you every second, day and night. Yet over the course of your entire lifetime, perhaps only ten of them will interact with an atom in your body. So in the absence of any reasonably measurable evidence for their existence, it is a small miracle that we were ever able to discover the neutrino in the first place.

1.2 The Evidence of Absence

The story of how we first discovered the neutrino despite its almost undetectable interactions with matter begins in the late 1920’s. Back then, physicists were only aware of three “particles” the photon, the electron, and the proton. Thanks to the work of Rutherford, Becquerel, and the Curies, it was known by then that the nucleus of an atom was a densely packed concentration of protons and that certain unstable nuclei could “transmute” into entirely different elements through radioactive decay. This process of course typically involves the ejection of one or more particles from the nucleus. It was through studying radioactive decay that hints of the existence of the neutrino were first found.

At the time, nuclear β decay was thought to involve the emission of an electron from the nucleus of an atom through the following reaction:



where N is the parent nucleus and N' is the daughter. In this decay, there are only two bodies in the final state. This leaves us with two unknown kinematic variables, the final

velocities of each of the two final particles. Fortunately there are exactly two equations of constraint that can provide us with a solution. The first of which is derived from the conservation of energy, and the second is derived from the conservation of momentum. In this situation, provided this reaction is energetically allowed, we will have a unique solution. This means that the electrons emitted from nuclear β decay should produce a monoenergetic spectrum.

Of course, the electrons given off in β decay are not emitted monoenergetically. If they were, I wouldn't be taking the time to describe this to you! The electrons are emitted with a continuous distribution of energies from zero up to (but not exceeding) the expected value from the two-body decay assumption. So some energy has gone missing in this reaction. As an interesting side note, Niels Bohr advocated for abandoning the principle of the conservation of energy as a possible explanation. The correct explanation however, came from Pauli.

Pauli proposed the existence of a third, undetected particle given off during β decay. The presence of a third particle in the final state would add one additional unknown kinematic variable (the final velocity of the third particle) without adding an additional equation of constraint. This would explain the observed continuous distribution for the energy spectra of the emitted electrons, making the correct form of equation 1.1:



By conservation of electric charge, this particle must be electrically neutral. The fact that this particle had yet to be observed in conjunction with β decay demonstrated that it interacted only weakly with matter. Lastly, we can deduce that this particle must have a very small mass because the spectrum for emitted electrons terminates at or very close to the predicted energy from a two-body decay. Pauli originally proposed calling this particle

the “neutron” but that name did not stick, and was repurposed for the particle discovered by Chadwick in the early 1930’s. The name that did stick was the one proposed by Fermi, “neutrino”, which translated from Italian means “the little neutral one.”

1.3 First Detection of the Neutrino

The direct observation of the neutrino did not occur for almost another 20 years through an experiment conceived by Fred Reines and Clyde Cowan [2, 3]. This experiment was based on the detection of the neutrino by inverse β decay through the reaction

$$p + \bar{\nu} \rightarrow n + e^+, \tag{1.3}$$

which was carried out in Hanford, Washington in the early 1950’s. Their neutrino source was a small nuclear reactor that gives off a reasonably high neutrino flux from the β decay of its fission fragments. Their detector consisted of a volume of cadmium-loaded scintillator instrumented with photomultiplier tubes for light detection. The detector itself was surrounded by thick layers of paraffin and lead to help ensure that the neutrinos were the only particles entering from outside. Since the reaction in equation 1.3 results in both a neutron and a positron, the characteristic signal that they looked for was a prompt pair of photons from positron-electron annihilation followed by a photon delayed by a few μsec emitted during neutron capture.

The initial results demonstrated an increase in the number of neutrino-like signals seen when the reactor was on, but on top of a much higher background rate than expected. The experiment was therefore redesigned so that information about the spatial origin of the event within the detector could be obtained thus giving more power to reject background events entering from outside. The experiment was repeated using a different nuclear reactor in Savannah River with much greater success producing an interaction rate consistent with the

expected results, allowing Reines and Cowan to definitively demonstrate the first detection of the neutrino. Their results were published in 1956 resulting in Reines being awarded the 1995 Nobel Prize. Cowan unfortunately, passed away in 1974.

1.4 Missing Solar Neutrinos

The first indication that something strange was afoot regarding neutrinos came from Ray Davis' attempt to observe them coming from the Sun. Prior to the first detection of the neutrino described above, it was known that the Sun was fueled by nuclear fusion reactions and should therefore be a source of neutrinos. The net result of these reactions is the combination of four protons to make one helium nucleus, two positrons, two neutrinos, and lots of energy in the form of photons. Because they are so easily absorbed and re-emitted, these photons can take thousands of years to work their way out to the surface of the Sun and by then, all information about what is happening in the interior is effectively lost. The neutrinos however, escape very easily thanks to their incredibly weak interactions with matter and can therefore serve as a probe to the interior of the Sun.

Ray Davis' experiment [4] was an attempt to make use of this fact to verify the flux of Solar neutrinos first predicted by John Bahcall in reference [5]. Davis' detector consisted of a 390,000 liter tank filled with liquid tetrachloroethylene (basically cleaning fluid) located deep underground in the Homestake mine in South Dakota. Conducting the experiment deep underground helped to reduce backgrounds from radiation and cosmic rays that are found in copious amounts on the surface. The neutrino detection mechanism was an inverse β decay similar to that used by Reines and Cowan through the reaction



However, unlike Reines and Cowan's experiment which detected signals immediately, Davis

used an ultra-pure chlorine solution and simply counted the number of argon atoms accumulated within the tank after a month or more had passed. Davis succeeded in making this rather difficult measurement and found a rate of roughly one third of the predicted rate, publishing his first results in 1968 [4].

The large uncertainties on the predicted solar neutrino flux combined with the difficulty in believing the claim of being able to count tens of argon atoms within a very large tank led to an initial reluctance to take Davis' results seriously. However, his measurement turned out to be accurate and was confirmed many years later by the SNO experiment [6, 7] with supporting results from many other experiments such as Homestake, GALLEX, GNO, SAGE, Super-K, and Borexino [8]. In retrospect, we can interpret Davis' results as the first indication that neutrinos oscillate between their flavor states. For this work, Davis was finally awarded the Nobel Prize in 2002.

1.5 Confirmation of Neutrino Oscillations

The reduced rate of expected neutrino interactions first observed by Ray Davis can be explained through the phenomenon of neutrino oscillations. The idea that Solar neutrinos could be changing from one type to another in mid-flight was first proposed by Bruno Pontecorvo in 1968 [9]. We now know that there are (at least) three flavors of neutrino, one paired with each of the three charged leptons: electron, muon, and tauon, with each having been experimentally observed in 1953 [2], 1962 [10], and 1999 [11] respectively. These flavor states are defined dynamically in the sense that a neutrino created in conjunction with an electron or positron is considered to be of the electron flavor. Similarly, a neutrino that interacts in your detector and produces an electron or a positron is also considered to be of the electron flavor. The neutrinos created by the Sun through a reaction such as that shown in equation 1.2 should all therefore be of the electron flavor, which is the only flavor to

which Davis' experiment was sensitive. So if electron neutrinos created in the Sun oscillate between the three flavor states as they travel, then the observed rate would be expected to be smaller by the time these oscillated neutrinos arrived at Earth.

The complete mathematical formalism governing neutrino oscillations is explained in detail in Chapter 2, but the oscillation probability for a neutrino created in one flavor and observed in another is roughly proportional to $\sin^2(CL/E_\nu)$ where C is a handful of factors (derived in Chapter 2), L is the distance that the neutrino has traveled, and E_ν is the neutrino energy. Any experiments aiming to confirm neutrino oscillations should measure neutrinos from the same or similar sources that correspond to different values of L or different values of E_ν . Variations in either of these quantities should produce a plot with the expected oscillatory behavior and would confirm the oscillation model.

The first experiment credited with “officially” observing neutrino oscillations is the Super-Kamiokande (Super-K) [12] experiment based in Japan. This experiment used a 50-kiloton water Cherenkov detector outfitted with photomultiplier tubes for light detection. Much like Davis' setup, the detector was situated deep underground to minimize the contamination from background events. The Super-K experiment was capable of observing what are known as “atmospheric” neutrinos. These neutrinos are created when high-energy cosmic rays strike atoms in Earth's upper atmosphere, creating muons and muon-flavored neutrinos. Since cosmic rays are bombarding the Earth from all directions at all times, the Super-K detector was able to observe neutrinos entering the detector from above and below. Of course, the neutrinos entering from below have crossed some reasonable fraction of the Earth's diameter and are therefore susceptible to oscillations as compared to the neutrinos entering from above, which have only traveled a “short” distance through the atmosphere. By measuring the total number of muon neutrino events seen as a function of the angle with which they entered the detector, the Super-K collaboration was able to demonstrate a clear

depletion of the number of events entering from below compared to the number entering from above that was consistent with a neutrino flavor oscillation model.

Another experiment credited with early confirmation of neutrino oscillations is KamLAND [13]. Much like the first experiments of Reines and Cowan, KamLAND used neutrinos produced in nuclear reactors as a source and detected them through the reaction shown in equation 1.3. The KamLAND detectors used liquid scintillator and photomultiplier tubes to detect the prompt electron-positron annihilation and delayed neutron capture signals and was positioned underground roughly 180 km from the source reactors. With an energy resolution of better than 10%, KamLAND was able to construct a neutrino energy spectrum that clearly demonstrated the expected energy-dependent oscillation behavior. Taken together, the observations made by these two experiments are consistent with the mathematical models that predict neutrino oscillations with changes in L and E_ν .

1.6 The Importance of Neutrino Physics

A great deal of progress has been made over the last 60 years towards understanding the nature of the neutrino; however, many questions still remain. We know that neutrinos have mass (or at least two of them do) because the oscillation framework as described in chapter 2 shows that if neutrinos were massless, they would not oscillate. However, we have yet to measure their absolute masses. Any attempts to do so have so far only yielded an upper limit and demonstrated that they are at least six orders of magnitude smaller than the electron mass. We also do not know the complete ordering of the three neutrino masses m_1 , m_2 , and m_3 . From measurements of Solar neutrinos [8] it is known that m_1 is smaller than m_2 but whether m_3 is smaller or larger than these two has yet to be determined. Answering this question would resolve the so-called neutrino “mass ordering.”

Another unanswered question is whether or not neutrinos are Majorana fermions (the

neutrino and the anti-neutrino are the same particle) or Dirac fermions (the neutrino and anti-neutrino are distinct particles) [14]. It is also not known if neutrinos exhibit CP violation or not, which could be observed through a non-zero measurement of the CP-violating phase factor δ . Lastly, while the three neutrino flavor mixing framework has had great success in explaining almost all experimental results, it has yet to be fully validated and verified. It is quite possible that more precise measurements of the parameters that govern neutrino oscillations could reveal phenomena beyond this framework, possibly even the existence of additional flavor states as some experiments such as LSND have suggested [15].

A question that the astute reader may be wondering at this point is “why should we bother studying a particle whose presence in our universe is barely noticeable?” This of course is a valid question, and one that we should often ask. Neutrinos are not only interesting, but related to many of the current big questions in particle physics. For example, CP violation within the neutrino sector could be linked to the problem of the baryon asymmetry in the universe. Neutrinos therefore have the potential to explain why the universe is filled with far more particles than anti-particles. As electrically neutral objects that only interact weakly with matter, neutrinos could also be a component of dark matter. Of course, their very small masses, and therefore ultra-relativistic nature, means that they cannot be the dominant component of dark matter since we know that most dark matter must be “cold” [8]. However they could still be part of the solution to the dark matter question and therefore contribute to structure formation in the early universe. Neutrinos are also thought to play a critical role in supernova explosions and are therefore responsible for helping to seed the galaxy with carbon, nitrogen, oxygen, and other essential ingredients for life as we know it [16].

Clearly neutrinos are tied to many phenomena that strongly suggest there is more to particle physics beyond what we currently know. In fact, since neutrinos are massless within

the Standard Model, one could argue that the existence of mass-induced neutrino oscillations is currently our best connection to possible new physics beyond the Standard Model. But apart from that, by being linked to the baryon asymmetry, dark matter, and the mechanisms that spread heavy elements into the galaxy, neutrinos have demonstrated themselves to be directly related to the origins of this planet and the life that clings delicately to its surface. Having a better understanding of neutrino physics can therefore help us strive towards the ultimate philosophical question of how did we get here in the first place.

CHAPTER 2

Neutrino Oscillations and Summary of the Experimental Status of the Field

This chapter gives a basis for understanding neutrino oscillations and provides a simple derivation of the appropriate expressions. As is the way of the particle physicist, I will work entirely with natural units in which $\hbar = c = 1$. For the purposes of these derivations, it is assumed that there are three flavors of neutrino. A more in depth derivation of the oscillation equations can be found in [1, 8, 17, 18]. I highly recommend reference [1] for its simplicity. Neutrino physics presents a rich landscape of questions, many of which are currently unanswered. This chapter will only focus on the topics directly relevant to the NOvA experiment. The equations and parameters related to many other important and fascinating topics such as neutrino-less double beta decay and the possibility of sterile neutrino flavors are not discussed.

2.1 Neutrino Eigenstates and the PMNS Mixing Matrix

We can describe the neutrino states using two different bases. The first is composed of the set of eigenstates that diagonalize the weak interaction Hamiltonian, the so-called flavor eigenstates, and are written as $|\nu_\alpha\rangle$ where $\alpha = e, \mu, \tau$. The second basis is formed from the eigenstates of the free particle Hamiltonian. These are referred to as the mass eigenstates and are written as $|\nu_j\rangle$ with $j = 1, 2, 3$. Given these two bases, we can construct a unitary rotation matrix that takes one from an expression in one of them to an expression in the

other. This means that a neutrino of a given flavor can be expressed as a linear combination of the mass states:

$$|\nu_\alpha\rangle = \sum_j U_{\alpha j}^* |\nu_j\rangle, \quad (2.1)$$

where $U_{\alpha j}^*$ is an element of this unitary rotation matrix, commonly known as the PMNS mixing matrix after Pontecorvo, Maki, Nakagawa, and Sakata [8].

It can be shown that with three neutrino flavors, the PMNS mixing matrix can be described by three real rotation angles θ_{12} , θ_{13} , and θ_{23} , and one complex CP-violating phase factor δ [8]. The standard way to parameterize the PMNS matrix is to split it into three matrices, each of which characterizes the mixing between two of the three mass states through one of the real rotation angles. There is then some freedom to choose where to place the complex phase factor δ and the convention is to associate it with the second matrix. In this factored form, the PMNS matrix appears as

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.2)$$

which when multiplied out becomes

$$U = \begin{bmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{13}s_{23}c_{12}e^{i\delta} & c_{23}c_{12} - s_{13}s_{23}s_{12}e^{i\delta} & c_{13}s_{23} \\ s_{23}s_{12} - s_{13}c_{23}c_{12}e^{i\delta} & -s_{23}c_{12} - s_{13}c_{23}s_{12}e^{i\delta} & c_{13}c_{23} \end{bmatrix}. \quad (2.3)$$

Here I have used the shorthand notation $c_{ij} \equiv \cos \theta_{ij}$ and $s_{ij} \equiv \sin \theta_{ij}$.

For completeness, I should note that if neutrinos are Majorana fermions, then the PMNS matrix shown in equation 2.2 includes a fourth matrix P that contains the Majorana phases of the form $P = \text{diag}(1, e^{i\alpha_1/2}, e^{i\alpha_2/2})$ where $\alpha_{1,2}$ are the Majorana phases. Because of the diagonal nature of this matrix, these phases are unobservable when considering flavor oscil-

lation experiments that measure the transition probabilities between flavors. The Majorana matrix P can therefore be safely ignored for the derivations that follow.

2.2 Three-Flavor Neutrino Oscillations in a Vacuum

Neutrino mass states propagating through vacuum are governed by the free particle Hamiltonian, with the time-dependent piece of the solution to the Schrödinger equation given by $\exp(-iEt)$. From equation 2.1, the evolution in time (t) of a neutrino flavor state can be written as:

$$|\nu_\alpha(t)\rangle = \sum_j e^{-iE_j t} U_{\alpha j}^* |\nu_j\rangle, \quad (2.4)$$

with E_j given by the relativistic free particle energy $E_j^2 = p^2 + m_j^2$. Since neutrino masses are very small with respect to their energies (which are typically \geq a few MeV) we can approximate the neutrino energy E_j as

$$E_j \approx |p| + \frac{m_j^2}{2|p|}. \quad (2.5)$$

Here we are making an assumption about the 3-momenta, namely $p_i = p_j \equiv p$ which is justified in references [9, 19]. The probability of oscillation from flavor α to flavor β is then given by

$$P_{\alpha \rightarrow \beta}(t) = |\langle \nu_\beta(t) | \nu_\alpha(t) \rangle|^2 = \left| \sum_j U_{\alpha j}^* U_{\beta j} e^{-i \frac{m_j^2}{2|p|} t} \right|^2. \quad (2.6)$$

Note that the phase $i|p|t$ which is common to all terms, has been dropped. Under the assumption that the neutrinos are ultra-relativistic, we can use $t \approx L$ to write $t/|p|$ as L/E where L is the distance traveled by the neutrino.

Taking advantage of the unitary nature of U , equation 2.6 can be expanded into

$$\begin{aligned}
P_{\alpha \rightarrow \beta} = & \delta_{\alpha\beta} - 4 \sum_{i>j} \Re [U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{ij}^2}{4E} L \right) \\
& + 2 \sum_{i>j} \Im [U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{ij}^2}{2E} L \right),
\end{aligned} \tag{2.7}$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. Neutrino oscillations are therefore governed not only by the four angles in the matrix U , but also by the ‘‘squared mass splittings’’ Δm_{21}^2 , Δm_{31}^2 , and Δm_{32}^2 . Experimentally, it has been shown [8] that $|\Delta m_{21}^2|$, is very small with respect to either $|\Delta m_{31}^2|$ or $|\Delta m_{32}^2|$ which allows for expressions of flavor oscillation probability to be simplified assuming $|\Delta m_{32}^2| \approx |\Delta m_{31}^2|$. This is often referred to as the ‘‘one-mass-scale-dominance’’ approximation or OMSD.

Two specific examples of equation 2.7 relevant to the NOvA experiment are the ν_μ survival probability, $P_{\mu \rightarrow \mu}$ and the ν_e appearance probability, $P_{\mu \rightarrow e}$. The ν_μ survival probability is given by

$$P_{\mu \rightarrow \mu} \approx 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{31} + 4 \sin^2 \theta_{13} \sin^2 \theta_{23} \cos 2\theta_{23} \sin^2 \Delta_{31}, \tag{2.8}$$

where $\Delta_{ij} \equiv \Delta m_{ij}^2 L / 2E$. For this derivation, a few reasonable approximations such as dropping terms proportional to $\sin^2 \Delta_{12}$ due to the much longer oscillation frequency of Δ_{12} relative to Δ_{32} have been used [20]. Given the known small value of θ_{13} , the ν_μ survival probability is dominated by the term proportional to $\sin^2 2\theta_{23}$ and is more commonly quoted using the OMSD assumption as

$$P_{\mu \rightarrow \mu} \approx 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{32}. \tag{2.9}$$

The ν_e appearance probability is a much richer expression as it involves δ as well as all

three mixing angles. For vacuum oscillations, $P_{\mu \rightarrow e}$ is given by

$$P_{\mu \rightarrow e} \approx P_{atm} + P_{sol} + 2\sqrt{P_{atm}P_{sol}} [\cos \Delta_{32} \cos \delta \mp \sin \Delta_{32} \sin \delta], \quad (2.10)$$

with

$$P_{atm} \equiv \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \Delta_{31}, \quad (2.11)$$

and

$$P_{sol} \equiv \cos^2 \theta_{23} \cos^2 \theta_{13} \sin^2 2\theta_{12} \sin^2 \Delta_{21} \approx \cos^2 \theta_{23} \cos^2 \theta_{13} \sin^2 2\theta_{12} \Delta_{21}^2, \quad (2.12)$$

where the \mp in equation 2.10 is $-$ for neutrinos and $+$ for anti-neutrinos [17]. However, we will see in section 2.3 that this expression will need to be modified due to the effects on electron neutrinos passing through matter.

Armed with equation 2.7, a qualitative interpretation of neutrino oscillations now becomes clear. Neutrinos produced in charged-current (CC) interactions (interactions involving the exchange of a W^\pm boson) always begin their life in a definite flavor state, which behind the scenes is a superposition of the three mass states. Since $m_1 \neq m_2 \neq m_3$, as the neutrino travels through space, these three components of the wave function will “beat” against each other like sound waves of slightly different frequencies traveling together. The “beating” of these quantum-mechanical amplitudes corresponds to the oscillating probability to detect the neutrino in one of the three flavor states, with $\Delta m_{ij}^2/E$ playing a role analogous to the “beat frequency.” Neutrinos should not be thought of as having a definite flavor any more than a coin should be thought of as being either heads or tails. Like the coin, the neutrino may start out in a definite flavor state and be detected in a definite flavor state, but the initial and final states need not be the same.

2.3 Matter Effects

Even though neutrinos only weakly interact with other particles, they can be affected by passing through large quantities of matter such as the Earth or the Sun in a way that is not the same for the three flavors. This is caused by the fact that macroscopic objects contain electrons and not muons, or tauons. Therefore when passing through matter, there is a coherent forward scattering amplitude shown in figure 2.1 only associated with the electron neutrino. Of course when passing through matter, all three neutrino flavors will participate in neutral-current (NC) interactions (interactions involving the exchange of a Z^0 boson) but these scattering amplitudes are flavor independent and therefore do not affect oscillations (excluding cases of extremely high density [18].) The net result of this additional scattering amplitude is a change in the effective index of refraction for electron neutrinos when passing through matter. Any oscillation probabilities involving ν_e or $\bar{\nu}_e$ will therefore be measurably different from the vacuum oscillation case.




Figure 2.1: Feynman diagram depicting the coherent forward scattering of an electron neutrino off of an electron via a charged current interaction. This interaction sets electron neutrinos apart from the other two flavors for neutrinos traveling through a macroscopic object such as the Earth, producing the so called “matter effect.”

The mathematics of the matter effect can be treated by adding a potential term to the Hamiltonian that accounts for this additional electron neutrino scattering. In the flavor

basis, this new Hamiltonian can be written as

$$H = U \begin{bmatrix} \frac{m_1^2}{2E} & 0 & 0 \\ 0 & \frac{m_2^2}{2E} & 0 \\ 0 & 0 & \frac{m_3^2}{2E} \end{bmatrix} U^\dagger + \begin{bmatrix} \sqrt{2}G_F N_e & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (2.13)$$

where G_F is the Fermi constant and N_e is the matter electron number density. For simplicity, N_e is assumed here to be constant, which is a reasonable assumption for the NOvA beam neutrinos that only pass through the Earth's crust, but not for neutrinos escaping from the Sun's core. For anti-neutrinos, equation 2.13 stays the same with the changes $\sqrt{2}G_F N_e \rightarrow -\sqrt{2}G_F N_e$ and $U \rightarrow U^*$. This Hamiltonian can then be re-diagonalized to obtain a new set of mass eigenstates for neutrinos propagating through matter, and the flavor oscillation probabilities described in section 2.2 can be rewritten in terms of these states.

After many hours (and several chalkboards) one can then make the appropriate modifications to the expression for $P_{\mu \rightarrow e}$ described above, which has been worked out nicely in [17]. Equation 2.10 remains unchanged, and equations 2.11 and 2.12 become

$$P_{atm} = \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2 (\Delta_{31} \mp aL)}{(\Delta_{31} \mp aL)^2} \Delta_{31}^2, \quad (2.14)$$

and

$$P_{sol} = \cos^2 \theta_{23} \sin^2 2\theta_{12} \frac{\sin^2 (\mp aL)}{(\mp aL)^2} \Delta_{21}^2, \quad (2.15)$$

with $a \equiv G_F N_e / \sqrt{2}$ and the $- (+)$ again referring to neutrinos (anti-neutrinos.) For the 810 km NOvA baseline, $aL \approx 0.23$ which makes the matter effect significant for the ν_e appearance channel.

Lastly, it should be noted that since the addition of the potential term in the Hamiltonian in equation 2.13 produces a new set of mass eigenstates, the ν_μ disappearance probability will not be completely unchanged by the matter effect. However, an expansion in $\sin^2 \theta_{13}$

and the term $\alpha \equiv \Delta m_{21}^2 / \Delta m_{31}^2$ shows that these types of corrections to the ν_μ disappearance probability are typically $< 1\%$ for experiments with neutrino energies and baselines similar to NOvA or T2K [18]. Therefore, equations 2.8 and 2.9 can safely be used for measurements dominated by limited statistics.

2.4 Experimental Results

2.4.1 Current Global Best Fits

Neutrino oscillation experiments are typically divided into three categories based on the types of neutrino source, energy scales, and oscillation parameter sensitivities involved. “Solar experiments” are those that strive to measure neutrinos created in the Sun’s core probing **very** long baselines at typical energies of a few MeV. These experiments are most sensitive to Δm_{21}^2 and $\sin^2 \theta_{12}$, and include experiments such as Homestake, GALLEX, GNO, SAGE, Super-K, SNO, and Borexino. “Reactor experiments” follow in the spirit of Reines and Cowan’s original 1956 detection of the first neutrinos. These experiments are primarily designed to measure θ_{13} through the disappearance of $\bar{\nu}_e$ created in the β^- decay of reactor fission fragments, with energies of a few MeV over baselines of a kilometer or so. The most prominent reactor results come from Double Chooz, Daya Bay, and RENO.

The last category of neutrino experiments (to which NOvA belongs) is the “accelerator/atmospheric experiments.” Through measuring neutrinos with higher energies (GeV and up) over longer baselines (hundreds to thousands of kilometers) these experiments have sensitivity to Δm_{31}^2 or Δm_{32}^2 , and $\sin^2 2\theta_{23}$ in the ν_μ disappearance channel. Through the ν_e appearance channel, the accelerator experiments also have some sensitivity to $\sin^2 2\theta_{13}$, $\sin^2 \theta_{23}$, and δ , potentially giving them access to the unanswered questions of the matter hierarchy and the octant of θ_{23} (described in section 2.4.2). Prominent experiments in this last category besides NOvA include Super-K, T2K, and MINOS. More information about

any of the experiments listed in this section and their results can be found in [8].

As described above, the three flavor neutrino oscillation parameters have all been measured by many different experiments, covering both appearance and disappearance measurements, over a variety of different baselines and different energy scales. Combining their results therefore represents a reasonable validation of the PMNS mixing matrix framework. A global fit to these experiments using the most recent data (as of 2014) has been performed and is summarized in [8, 21]. For this analysis, the “one-mass-scale-dominance” (OMSD) assumption, $\Delta m_{31}^2 \approx \Delta m_{32}^2$, has been used. The results from this global fit are summarized in table 2.1.

parameter	best fit ($\pm 1\sigma$) NH	best fit ($\pm 1\sigma$) IH
Δm_{21}^2 [10^{-5}eV^2]	$7.54^{+0.26}_{-0.22}$	-
$ \Delta m^2 $ [10^{-3}eV^2]	2.43 ± 0.06	2.38 ± 0.06
$\sin^2 \theta_{12}$	0.308 ± 0.017	-
$\sin^2 \theta_{13}$	$0.0234^{+0.0020}_{-0.0019}$	$0.0240^{+0.0019}_{-0.0022}$
$\sin^2 \theta_{23}$	$0.437^{+0.33}_{-0.023}$	$0.455^{+0.039}_{-0.031}$
δ/π (2σ range)	$1.39^{+0.38}_{-0.27}$	$1.31^{+0.29}_{-0.33}$

Table 2.1: The best fit values for the neutrino oscillation parameters within a three flavor model. Shown (when appropriate) are the fits assuming normal hierarchy (NH) and inverted hierarchy (IH). Within the OMSD assumption, measurements for $|\Delta m_{31}^2|$ and $|\Delta m_{32}^2|$ have been combined into one measurement of $|\Delta m^2|$. These values have been taken from [8, 21].

2.4.2 Remaining Unknowns

The measurements of the neutrino oscillation parameters as outlined in section 2.4.1 present a reasonably complete picture of the PMNS mixing matrix framework. However, simply knowing these parameters is not all that there is to it as there are (at least) two more relevant questions whose answers will help complete our picture of neutrino oscillations.

The first of these is the question of the neutrino “mass ordering.” From the vacuum oscillation equations discussed in section 2.2, we can see that it is not possible to obtain

information about the sign of the Δm^2 terms by measuring vacuum oscillations alone. However, due to the presence of $\mp aL$ in the matter oscillation equations discussed in section 2.3, with some knowledge about your source (whether it is emitting neutrinos or anti-neutrinos) then the matter effect can give you access to information about the sign of Δm^2 . For solar neutrinos (which are all created as ν_e) the high density of the Sun’s interior plus the largeness of its radius produces a matter effect significant enough to allow us to determine that $\Delta m_{21}^2 > 0$ [8]. However, the sign of Δm_{31}^2 and Δm_{32}^2 have yet to be determined. We have discovered that the magnitude of these two quantities is much greater than Δm_{21}^2 leaving us with two possible scenarios: $m_1 < m_2 < m_3$ referred to as the “normal” ordering, or $m_3 < m_1 < m_2$ referred to as the “inverted” ordering.

The second question deals with the so-called “octant” of θ_{23} . Our current best measurements tell us that $\theta_{23} \approx 45^\circ$. However due to the ambiguity introduced by the $2\theta_{23}$ in equation 2.9, it is not possible to determine if $\theta_{23} > 45^\circ$ (putting it in the “upper” octant) or $\theta_{23} < 45^\circ$ (putting it in the “lower” octant) without a measurement far beyond our current level of precision that would allow one to fit with the smaller terms in 2.8. Another avenue for determining the θ_{23} octant (discussed in more detail in section 2.6) would be to use a ν_e appearance measurement that could take advantage of the $\sin^2 \theta_{23}$ term in equation 2.14. The value of θ_{23} governs the relative proportions of ν_μ and ν_τ in the three mass states. Knowing the octant of θ_{23} will determine how strongly each of these state couples to ν_μ and ν_τ .

2.5 Measurements of $|\Delta m_{32}^2|$ and $\sin^2 2\theta_{23}$

Since this thesis is dedicated to a measurement of θ_{23} through the ν_μ disappearance channel, it is worth taking a moment to discuss the other experiments that have contributed to this measurement. I begin with a discussion of measurements from atmospheric neutrinos with

the Super-K experiment and include the most recent results from the ν_μ beam oscillation experiments T2K and MINOS+. I will conclude with a quick summary of the ν_τ appearance results from the OPERA experiment.

2.5.1 Results From Super-K

The Super-K collaboration used the same 50-kiloton water Cherenkov detector built originally for the Kamiokande experiment, which is now being used for the T2K experiment [22]. Super-K detects “atmospheric” neutrinos that come from the decay of pions into muons created in the Earth’s upper atmosphere by cosmic rays as described in section 1.5. Each pion decay produces three neutrinos, a ν_μ , a $\bar{\nu}_\mu$, and either a ν_e or $\bar{\nu}_e$ depending on whether a π^+ or π^- was originally produced. For the purposes of their ν_μ disappearance measurement, Super-K treated ν_μ and $\bar{\nu}_\mu$ oscillations as the same, combining all muon-like events into one data sample.

By examining the “fuzziness” of the Cherenkov rings, the Super-K collaboration was able to develop a multivariate method for separating electron-like events from muon-like events. These events were grouped into two energy categories, sub-GeV and multi-GeV based on whether the reconstructed neutrino energy was less than or greater than 1.33 GeV, and counted as a function of the reconstructed lepton direction. The lepton direction is of course closely related to the fraction of the Earth’s diameter that the neutrino traversed prior to entering the detector, and therefore is a proxy for the oscillation baseline. The full analysis separated the muon-like events further into single-ring, multi-ring, and partially contained, to help improve their overall sensitivity. The distribution for the single-ring events is shown in the left plot of figure 2.2.

The method described here of using the lepton direction with respect to θ_{zenith} is referred to as the “zenith angle analysis.” Super-K also performed an analysis separating events by




Figure 2.2: Top: The expected distributions for sub-GeV and multi-GeV single-ring electron- and muon-like events. The boxes indicate a prediction based on a model without oscillations and the solid line indicates the best fit to a two flavor oscillation model. **Bottom:** Measurement contours for 68, 90, and 99% confidence intervals showing the best fit for $|\Delta m^2|$ and $\sin^2 2\theta$ with a two flavor oscillation model for both the θ_{zenith} and L/E analyses. These plots were taken from reference [22].

L/E instead of θ_{zenith} which is outlined in [23]. For both of these methods, a two flavor $\nu_\mu \rightarrow \nu_\tau$ disappearance model was used to fit for $|\Delta m^2|$ and $\sin^2 2\theta$. The best measurement of θ (which is really θ_{23}) comes from the zenith angle analysis which obtained a measurement of $\sin^2 2\theta > 0.92$ at the 90% confidence level using a total exposure of 1489 days [22]. Due to the fact that binning in L/E gives a better depiction of the shape of the disappearance spectrum, the L/E analysis resulted in a better measurement of $|\Delta m^2|$ but a less constrained measurement of $\sin^2 2\theta$. A measurement contour for the two-flavor oscillation fit parameters is shown in figure 2.2 for both methods.

2.5.2 Recent Results from T2K and MINOS+

As of May 2015, additional neutrino oscillation results have been published that are not reflected in the numbers presented in the previous sections. Specifically relevant to the NOvA experiment are the most recent results published by T2K [24] and MINOS+ (an extended run of the original MINOS experiment) [25]. Both experiments are similar to NOvA in terms of the general design and the main physics goals. Each is a beam neutrino oscillation experiment, using a primarily ν_μ beam to probe oscillations in both the ν_μ disappearance and ν_e appearance channels. T2K takes place in Japan using a neutrino beam created at the J-PARC facility, and MINOS+ is based out of Fermilab using the same beam line that is being used for NOvA. Each of them can make a measurement of $|\Delta m_{32}^2|$ and $\sin^2 2\theta_{23}$ through ν_μ disappearance, and $\sin^2 \theta_{13}$ through ν_e appearance. Through the matter effect on the ν_e appearance measurement, they also have some sensitivity to δ , the mass hierarchy, and possibly the octant of θ_{23} .

Both experiments combined ν_μ and ν_e measurements to perform a full three-flavor fit for the oscillation parameters. T2K measures $\sin^2 \theta_{23} = 0.514_{-0.056}^{+0.055}$ (0.511 ± 0.055) and $|\Delta m_{32}^2| = 2.51 \pm 0.10 [\times 10^{-3} \text{eV}^2]$ ($|\Delta m_{31}^2| = 2.48 \pm 0.10 [\times 10^{-3} \text{eV}^2]$), and $\sin^2 \theta_{13} =$

$0.042_{-0.021}^{+0.013}$ ($0.049_{-0.021}^{+0.015}$) assuming the normal (inverted) hierarchy. By combining their results with results from reactor experiments, they are also able to exclude $\delta = [0.15, 0.83] \pi$ ($[-0.08, 1.09] \pi$) at the 90% confidence level and show a weak preference for the normal hierarchy and the upper octant of θ_{23} (although the maximal value is not excluded.) By including measurements of atmospheric neutrinos, MINOS+ presented $\sin^2 \theta_{23} = 0.43_{-0.04}^{+0.16}$ ($0.43_{-0.05}^{+0.19}$) and $|\Delta m_{32}^2| = 2.34 \pm 0.09 [\times 10^{-3} \text{eV}^2]$ ($2.37_{-0.07}^{+0.11} [\times 10^{-3} \text{eV}^2]$) with a slight preference for the inverted hierarchy and the lower octant of θ_{23} . Measurement contours for these values from both experiments are shown in figure 2.3.




Figure 2.3: Left: Sensitivity contours taken from [24] showing the current T2K measurements of $\sin^2 \theta_{23}$ and $\sin^2 \theta_{13}$ vs. Δm_{32}^2 . **Right:** Sensitivity contours taken from [25] showing the most recent MINOS+ measurement of $\sin^2 \theta_{23}$ vs. Δm_{32}^2 as compared to an earlier result from T2K published in 2014.

2.5.3 Results from OPERA

The OPERA experiment is a ν_μ beam oscillation experiment similar to T2K, MINOS, and NOvA. However, OPERA was designed to look for a $\nu_\mu \rightarrow \nu_\tau$ appearance signal using the CERN Neutrinos to Gran Sasso (CNGS) beam. In July of 2015, they reported on five candidate ν_τ events that allow them to claim detection of a $\nu_\mu \rightarrow \nu_\tau$ appearance signal with a significance larger than 5σ [26]. They performed a fit for Δm_{23}^2 resulting in a value of $\Delta m_{23}^2 = 3.3 \times 10^{-3} \text{ eV}^2$ with a 90% confidence interval of $[2.0, 5.0] \times 10^{-3} \text{ eV}^2$. While this measurement did not allow them to fit for θ_{23} , it is an important confirmation of the three-flavor mixing paradigm and is relevant to NOvA whose ν_μ are disappearing into ν_τ .

2.6 Physics Goals of the NOvA Experiment

The NOvA experiment will primarily study neutrino oscillations with two detectors placed 810 km apart in what begins as a coherent beam of muon neutrinos. As these neutrinos travel from the near detector to the far detector, they will mostly oscillate from $\nu_\mu \rightarrow \nu_\tau$ but some will also oscillate from $\nu_\mu \rightarrow \nu_e$. The science that will be done with NOvA can be divided into three basic analysis categories. The first is the ν_μ disappearance analysis which will measure $\sin^2 \theta_{23}$ and $|\Delta m_{32}^2|$. The second is the ν_e appearance analysis which will measure $\sin^2 \theta_{13}$, δ , and can provide insight into the questions of the matter hierarchy and the octant of θ_{23} , depending on what values Nature has chosen for some of these parameters. Lastly, there are a wide variety of other physics goals outside of the scope of the standard beam-neutrino oscillation analyses that are put into a category appropriately labeled “exotic studies.” This category includes studying atmospheric neutrinos, supernova neutrinos, dark matter, magnetic monopoles, sterile neutrinos, etc. These topics, while fascinating, are not discussed further in this document.

As discussed in section 2.3, the beam neutrinos traveling through the Earth’s crust will

affect the ν_e and $\bar{\nu}_e$ appearance probabilities differently. Getting the most out of equations 2.10, 2.14, and 2.15 requires measuring the difference between these two probabilities, which creates the need for taking separate data with a ν_μ beam and a $\bar{\nu}_\mu$ beam. Fortunately, the beam provided to NOvA by Fermilab (the NuMI beam) can run in either a ν or a $\bar{\nu}$ mode. The details of the NuMI beam are discussed further in section 3.1. The current plan for the full NOvA exposure is to run for 3 years in each of these modes, which will maximize the potential physics reach of the ν_e appearance analysis. The absence of δ and the matter effects in equation 2.8 will allow the ν_μ disappearance analysis to treat ν_μ and $\bar{\nu}_\mu$ interactions as the same, combining all 6 years of data into one data set.

2.6.1 The ν_μ Disappearance Measurement

The ν_μ disappearance analysis will proceed by creating an energy spectrum out of data selected as ν_μ CC events. An example of the true energy spectra for contained ν_μ CC events in the far detector is shown in figure 2.4 with the unoscillated spectra in black and an oscillated spectra (assuming the maximal value for θ_{23}) in red. A fit can then be performed to the shape of the measured energy spectrum for $|\Delta m_{32}^2|$ and $\sin^2 2\theta_{23}$ using equation 2.9, or to $|\Delta m_{31}^2|$ and $\sin^2 \theta_{23}$ using equation 2.8. For this second more precise fit, $\sin^2 \theta_{13}$ must be an input parameter which can be taken for example, from the high-precision reactor neutrino measurements.

Shown in the left plot in figure 2.5 are true neutrino energy spectra for contained ν_μ CC events for two different values of $\sin^2 \theta_{23}$. These plots are scaled to the expected exposure for the first 3 years of running with the NuMI beam in ν mode. Note that the relative difference between these two curves is greatest near 1.5 GeV which corresponds roughly to the first oscillation maximum for the 810 km NOvA baseline. Distinguishing the two possible values of $\sin^2 \theta_{23}$ shown in this plot therefore depends greatly on accurate energy reconstruction

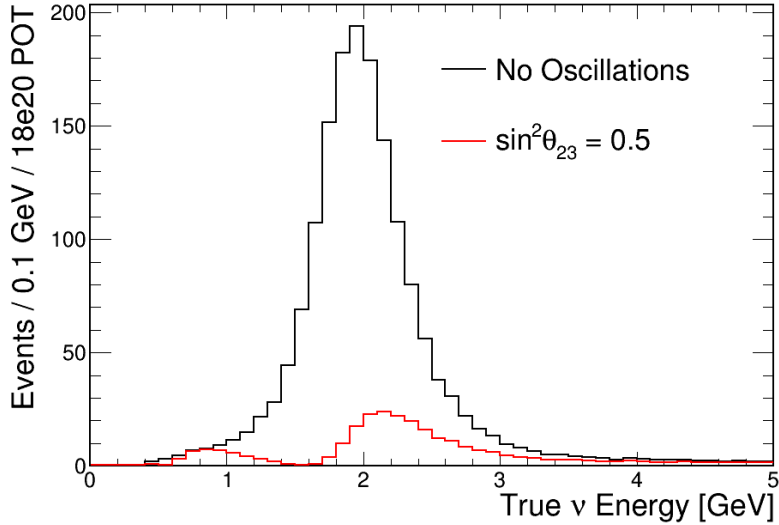


Figure 2.4: Simulated true energy spectra for contained ν_μ CC events in the NOvA far detector. The exposure is assuming 3 years of nominal data taking with a 14 kiloton detector while running the NuMI beam in ν mode.

and an accurate count of the events in this region. With reconstructed energy resolutions of 5 - 10%, and assuming what is expected to be only a small contamination from cosmic and NC background events, the expected sensitivities for the full 6 year exposure are shown in the right plot of figure 2.5 for these two possible values of $\sin^2 \theta_{23}$. Note that while it is technically possible to acquire information about the octant of θ_{23} using equation 2.8, this sensitivity is beyond the NOvA ν_μ disappearance measurement alone. The greatest sensitivity to the octant will come from the ν_e appearance measurement which is described in the next section. A full description of event reconstruction, energy estimation, and background estimations can be found in chapters 6, 7, and 8.

As compared to the contours shown in figure 2.2, NOvA has the potential to surpass the Super-K measurement producing the world's best measurement of $\sin^2 \theta_{23}$ with constraints on $|\Delta m_{32}^2|$ being comparable to values from other experiments. While it is not correct to say that the T2K and MINOS+ results presented in section 2.5.2 are in conflict, it is interesting to note that T2K has a preference for the maximal value of θ_{23} and a slight preference

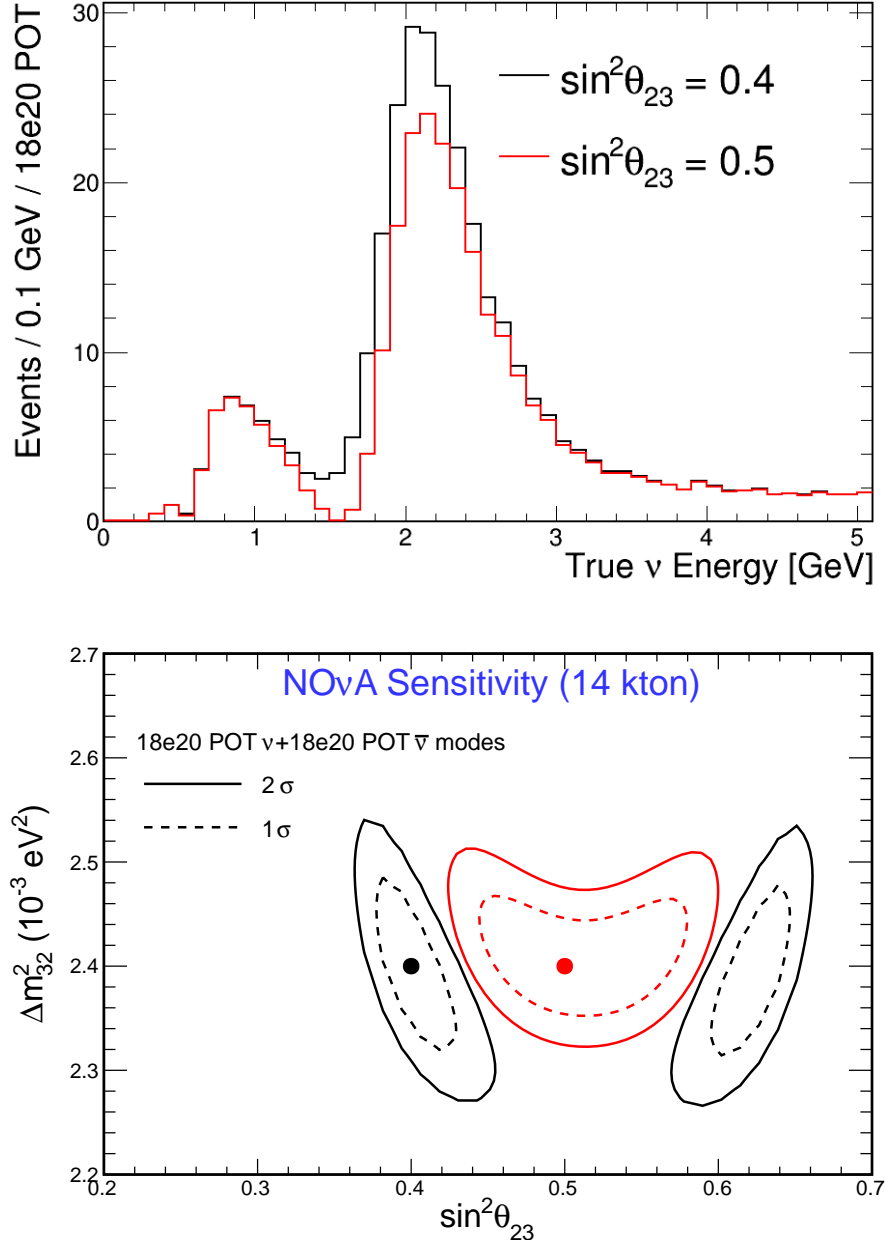


Figure 2.5: Top: An example of the expected true energy spectra for contained ν_μ CC events in the NOvA far detector under two different assumptions for $\sin^2\theta_{23}$. The exposure is assuming 3 years of nominal data taking with a 14 kiloton detector while running the NuMI beam in ν mode. **Bottom:** Example of the expected sensitivity contours for the full NOvA exposure (3 years running the NuMI beam in ν mode plus 3 years running in $\bar{\nu}$ mode) for the two values of $\sin^2\theta_{23}$ shown in the top plot.

for the normal hierarchy and the upper octant, while MINOS+ has a slight preference for the opposite (non-maximal θ_{23} , inverted hierarchy, and lower octant.) The NOvA ν_μ

disappearance measurement has the potential to clarify these measurements.

2.6.2 The ν_e Appearance Measurement

The reactor experimental results of the past few years have measured a small, but non-zero value of θ_{13} . For NOvA, this means potential access to δ and the questions of the mass ordering and the octant of θ_{23} through fitting to the ν_e appearance measurements using equations 2.10, 2.14, and 2.15. This analysis will proceed by separately counting ν_e and $\bar{\nu}_e$ events to obtain both appearance probabilities. These two measurements can be plotted together to form what is known as a “bi-probability” plot.

An example of two bi-probability plots, generated for the full 6 year expected exposure, are shown in figure 2.6. Several things become clear from these plots. First, while the size of the P_{sol} term in equation 2.10 is small, the size of the interference term between P_{atm} and P_{sol} is comparable to the dominant term P_{atm} , which can allow for a large CP-induced asymmetry (depending on the value of δ .) Second, the matter effect separates these ellipses into the blue and red curves for the normal and inverted hierarchies, which can cause potential confusion with any effects from CP asymmetry.

An example of a case that would allow us to make strong statements about the mass hierarchy and δ , would be the normal ordering with $\delta = 3\pi/2$ (the starred points in figure 2.6.) This case would allow us to claim to $> 2\sigma$ that the ordering was normal. If on the other hand, Nature has chosen the inverted ordering and $\delta = 3\pi/2$, then we will not be able to distinguish a best fit of inverted ordering with $\delta = 3\pi/2$ from normal ordering with $\delta = \pi/2$. In this case, input from another experiment with a different baseline such as T2K, would be needed to have hope of distinguishing the ordering. However, we can still claim that we have observed CP violation in this case since a non-zero value of δ would be required to cancel out the matter effect producing equal ν_e and $\bar{\nu}_e$ appearance probabilities.

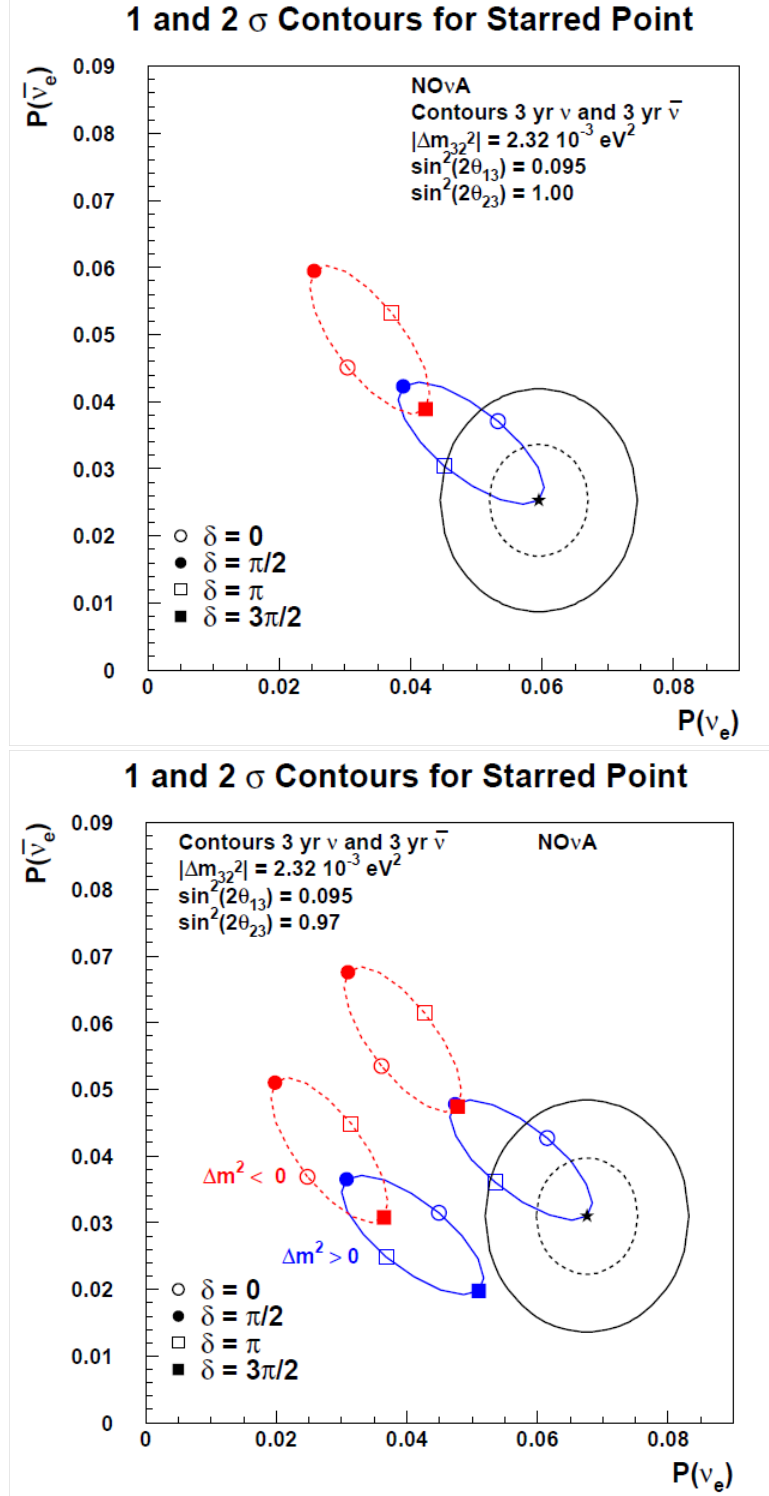


Figure 2.6: Examples of bi-probability plots for different values of $\sin^2 2\theta_{23}$. The top plot assumes θ_{23} is maximal and the bottom plot assumes it is non-maximal, with two sets of ellipses for each of the two possible octants. The star on each plot indicates the “best possible scenario” for making the strongest statements about δ and the matter hierarchy.

If θ_{23} is non-maximal, then the hierarchy ellipses in the left plot in figure 2.6 split into the ellipses shown in the right plot, with the upper two ellipses representing $\theta_{23} > 45^\circ$. As with the case of resolving the ordering, determining the correct octant for θ_{23} depends on the choice of hierarchy and the best fit value for δ . Shown in figure 2.7 are two possible sensitivities to the octant as a function of δ . The top plot corresponds to the case of maximum ambiguity between the hierarchy and δ (measuring the point where the two upper ellipses touch in the right plot in figure 2.6) and the bottom plot corresponds to a best-case scenario (measuring the starred point in the bottom plot in figure 2.6.)

2.6.3 Improving the Measurement of θ_{23}

As outlined above, θ_{23} has already been measured by several experiments. So one might ask why is it interesting to make this measurement again. First, despite the fact that θ_{23} was the first mixing angle to be measured, it is currently the one that we know to the least precision. The results discussed in section 2.5.2 show that some (very minor) tension exists between the current best measurements from T2K and MINOS+ which a more precise measurement could help to resolve. Second, a precision measurement of θ_{23} can help to determine the correct texture for the PMNS mixing matrix. There exists the possibility that θ_{23} is maximal which may be suggestive of a new symmetry within the lepton sector. New symmetries are of great interest to the world of particle physics since they hold potential insight into new physics, and a precision measurement of θ_{23} is necessary to either strengthen or rule out this idea. Lastly, since θ_{23} appears in equations 2.14 and 2.15, a more accurate measurement of this parameter from the NOvA ν_μ disappearance analysis will improve information obtained through the NOvA ν_e appearance analysis. As described in section 2.6.2, this analysis has the potential to shed significant light on many outstanding issues such as resolving the mass ordering and determining the value of the

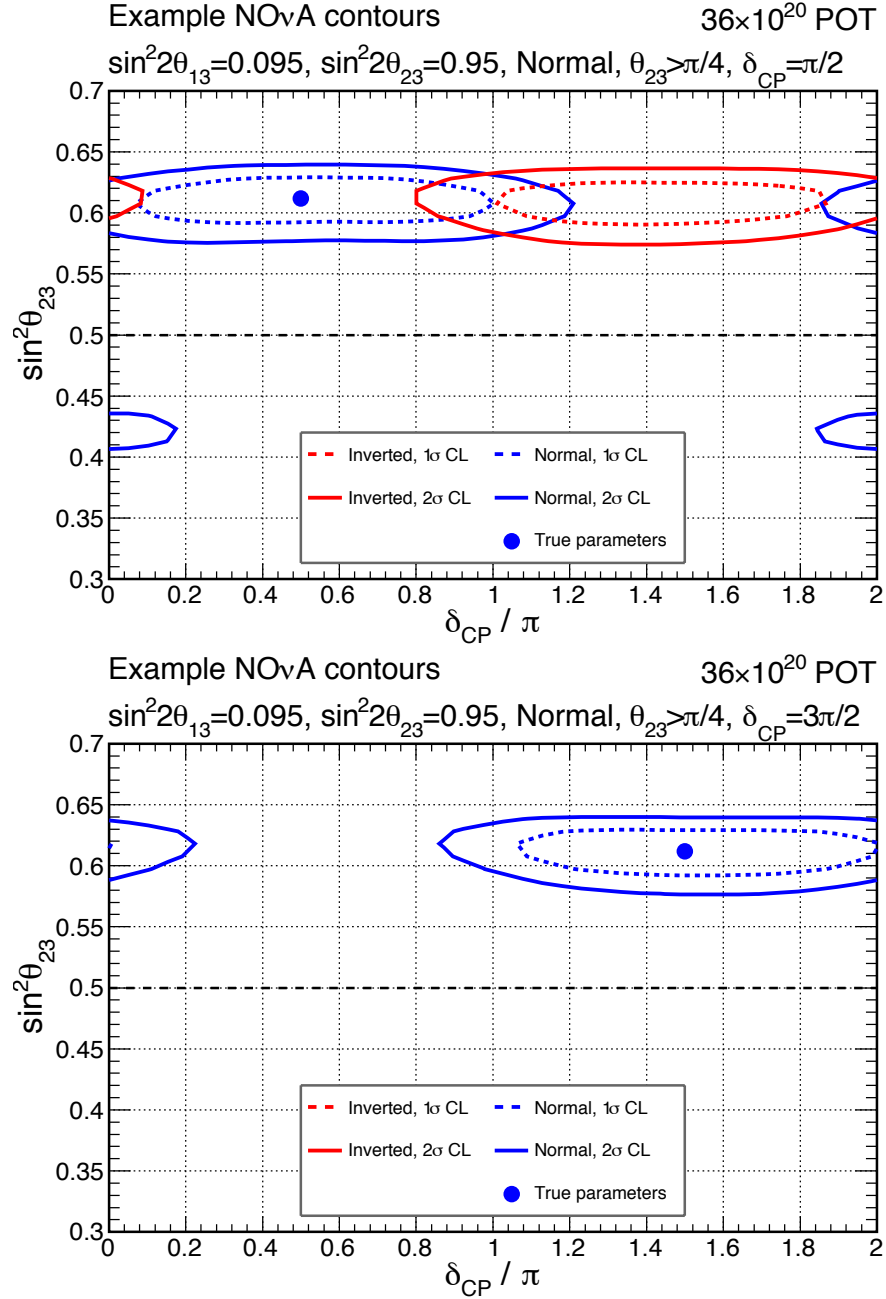


Figure 2.7: Examples sensitivities to the octant of θ_{23} as functions of δ , assuming $\sin^2 2\theta_{23} = 0.95$ and $\theta_{23} > 45^\circ$.

CP-violating phase factor δ .

CHAPTER 3

Design of the NOvA Experiment

The design of the NOvA experiment varies little from what was outlined originally in the technical design report (TDR) [27]. The main details are summarized in this chapter with an emphasis on the items relevant to the later chapters.

3.1 The NuMI Beam

The NOvA experiment uses the NuMI (**N**eutrinos at the **M**ain **I**njector) beam provided by the Fermilab accelerator division as our neutrino source. The Main Injector accelerates protons up to 120 GeV which are grouped together into six batches (see figure 3.1) that span a spill window of 10 μ sec. The protons in each spill collide with a graphite target to create mesons (primarily pions) that produce muon neutrinos as decay products.

As of the writing of this document, the NuMI beam was operating around 300 kW which corresponds to roughly 2.1×10^{13} POT (protons on target) per spill. However, Fermilab is currently engaged in continuing efforts to upgrade the NuMI beam to bring the power up to 700 kW providing 5.0×10^{13} POT per spill, roughly doubling the total NuMI beam power provided for the majority of the MINOS experiment. This will be accomplished by reducing the intervals between spills from 2.2 seconds to 1.33 seconds and simultaneously increasing the beam intensity by 10%. To achieve this increase in total beam power, the Recycler ring (previously used for anti-proton storage for the Tevatron collider program)

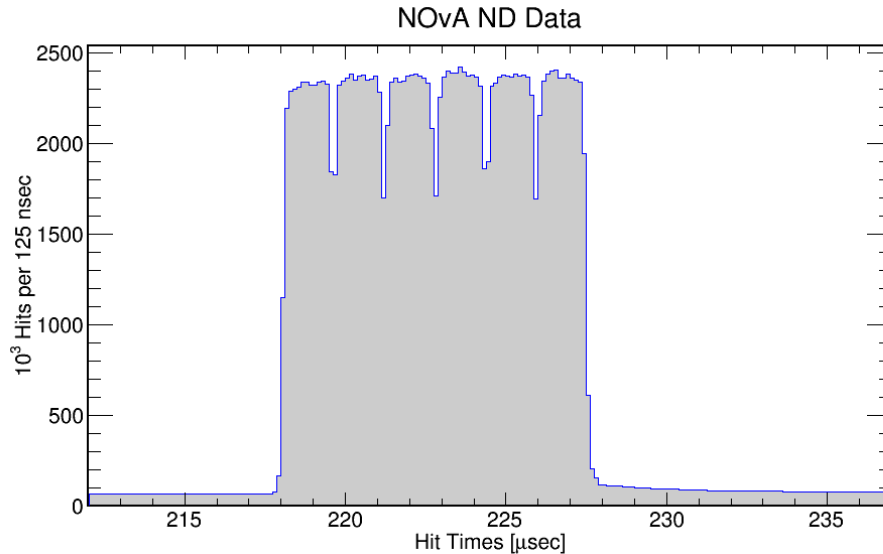


Figure 3.1: The distribution of times for all hits seen with the NOvA near detector with respect to a trigger time issued 218 μsec before the start of the beam spill. This plot includes data from 525,899 NuMI spills taken between Feb. 14, 2015 and Feb. 17, 2015.

has been re-purposed into a pre-injector to the Main Injector. This allows for protons from the Booster to be collected by the Recycler while the Main Injector ramps. These protons are “slip-stacked” into the Main Injector which reduces the cycle time to 1.5 seconds. The final step to bring the cycle time to 1.33 seconds is accomplished by increasing the maximum proton acceleration from 204 GeV/sec to 240 GeV/sec through upgrading the existing power supplies and adding two more RF stations. When complete, the NuMI beam will deliver an estimated 6.0×10^{20} POT per year, assuming 60% beam-up time for 10 operating months.

Once the protons have been bunched together and accelerated up to 120 GeV, they are directed to a 95 cm long graphite target. The target is water cooled and the cooling system has been modified to handle the increase in beam intensity. Proton collisions with the carbon atoms in the target produce a plethora of pions and kaons which are focused into a coherent beam by two parabolic magnetic horns (see figure 3.2) These horns can be run in forward or reverse current mode allowing for the selection of positively or negatively charged mesons leading to either a neutrino or an anti-neutrino beam. The focused meson

beam then travels through a 675 m long evacuated pipe where the mesons decay producing charged leptons and neutrinos. This decay pipe is followed by various hadron and muon detectors for monitoring and roughly 250 m of rock which absorbs all remaining beam products (except of course for the neutrinos) prior to reaching the near detector.




Figure 3.2: A cartoon representation of the NuMI beam.

The NOvA experiment was designed to have both detectors situated off of the main axis of the neutrino beam. While this reduces the overall event rate seen at either detector, it has several main advantages. Given the medium-energy focusing horn configuration (described below) a detector situated 14 mrad off-axis will see a flux of neutrinos produced from pions with a broad spectrum of different energies. However, the decay kinematics of these boosted pions creates neutrinos whose energy is relatively independent of the parent pion energy. For small angles, the flux of neutrinos created by the $\pi \rightarrow \mu + \nu$ in flight decay, seen by a detector with cross-sectional area A and at a distance z from the pion source is given by

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

where θ is the angle between the pion direction and the neutrino direction and $\gamma = E_\pi/m_\pi$.

The neutrino energy as a function of the parent pion energy is given by

$$E_\nu = \frac{0.43E_\pi}{1 + \gamma^2\theta^2}. \quad (3.2)$$

Both of these equations are plotted as functions of E_π for several values of θ in figure 3.3. The choice of being 14 mrad off-axis therefore creates the narrowly peaked neutrino spectra seen in the right plot in figure 3.4.




Figure 3.3: Left: Neutrino flux as a function of the energy of the parent pion for different off-axis angles. **Right:** Neutrino energies as a function of the parent pion energy for different angles with respect to the original pion direction.

The focusing horn positions are configurable and can be chosen to provide a neutrino beam with different characteristics. Three different horn configurations (low-energy, medium-energy, and high-energy) were considered when designing the NOvA experiment, each affecting the total beam flux as well as the shape of the energy spectra. Depicted in figure 3.4 are the expected spectra for the low-energy and medium-energy configurations. The medium-energy tune was chosen for NOvA since at 14 mrad off-axis, it produces the highest neutrino flux at 2 GeV (roughly the oscillation maximum for a detector situated 810 km from the beam source.) The narrowness of this distribution has the advantage that it helps reduce our neutral-current background events which have the potential to fake a signal in the 2 GeV region by “feeding down” from higher energies.




Figure 3.4: Expected charged current ν_μ event spectra for the low-energy (left) and medium-energy (right) horn configurations, calculated prior to oscillations for a detector 810 km from the beam source.

3.2 The NOvA Detectors

3.2.1 Physical Design

As a beam neutrino oscillation experiment, NOvA is designed to function with two separate detectors, a near detector and a far detector (see figure 3.5.) The main purpose of the far detector is to measure the energy spectra for our beam neutrinos, separating muon and electron charged-current interactions from neutral-current interactions. The near detector is designed to measure the unoscillated beam spectra so that it can be extrapolated to the far detector, and to measure the inherent electron neutrino component in the beam. The two detectors are as close to being functionally identical in physical structure and data acquisition architecture as possible to help minimize the impact of systematic uncertainties.

The NOvA detectors are constructed from extruded PVC tubes called “cells” (discussed in section 3.3.1) that span either the full width or the full height of the detector. Both of the detectors are composed of flat planes constructed from multiple modules glued together side-by-side. These planes are arranged in alternating, orthogonal layers oriented normal




Figure 3.5: Both NOvA detectors shown to scale with each other (and with a cartoon person.) The inset shows the alternating, orthogonal plane structure. The far and near detectors are approximately 60 m and 16 m long respectively.

to the neutrino beam (see figure 3.5) which allows for our reconstruction software to track a particle in three dimensions as it passes through multiple planes. In the far detector, 32 planes (24 for the near detector) are glued together into one solid piece called a block. Blocks are stacked together to complete the physical structure of the detectors.

Since NOvA is aiming to simultaneously perform a ν_μ disappearance and a ν_e appearance measurement, it is important for the NOvA design to be able to identify muons and electrons, the signatures of the respective neutrino charged-current interactions. In addition, NOvA must also be able to distinguish electrons from π^0 s which are a common element in neutral-current interactions and will be a component in the backgrounds for the ν_e appearance measurement. To achieve these goals, NOvA was designed to use relatively low Z materials (PVC and mineral oil are both primarily carbon.) The Moliere radius [8] for the NOvA detectors is approximately 11 cm (equivalent to the width of 2 cells) which gives electron showers sufficient width to be distinguished from non-showering particle tracks created by muons, charged pions, and protons. The radiation length for the detectors is

approximately 40 cm. When a π^0 is created in a neutrino interaction, the two photons produced from the very rapid (10^{-17} s) decay of the π^0 will on average travel a reasonable distance from the interaction vertex before they convert (see figure 3.6.) This gap is the hallmark sign of a π^0 in the NOvA detectors and will help to distinguish electrons from π^0 s provided that the neutrino interaction vertex can be identified.

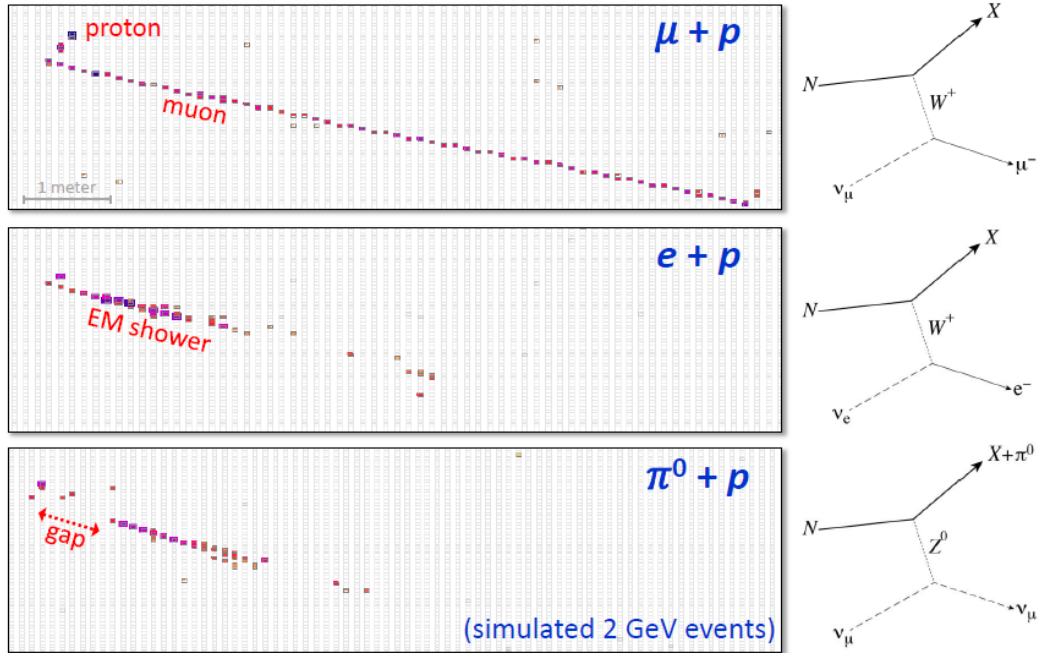


Figure 3.6: Three simulated events showing NOvA’s ability to distinguish muons, electrons, and π^0 s. The distinguishing feature of the electron over the muon is the shower-like nature of its track, and the distinguishing feature of the π^0 over the electron is the gap between the event vertex and the start of the shower.

3.2.2 The Far Detector

The NOvA far detector has been constructed at a site near Ash River, MN. This site was chosen because it is the furthest point away from Fermilab along the NuMI beam line still accessible by road within the United States. At 810 km away from the beam origin, this site is situated at the oscillation maximum for neutrinos produced with the medium-energy horn tune. The detector is built at approximately 10 m below ground level and the neutrino

beam enters traveling upwards at an angle of 3° . The detector itself is composed of a total of 344,064 cells each of which is 15.5 m long. These cells are arranged into 896 planes placed normal to the beam, making the full detector dimensions approximately $15.5 \text{ m} \times 15.5 \text{ m} \times 60.0 \text{ m}$. The total mass of the far detector is 14 kilotons, which is roughly 65% scintillator and 35% plastic.

Since the far detector is built on the surface at an elevation of 1220 ft above sea level, cosmic rays have the potential to be a burdensome source of background events. For the ν_μ disappearance analysis, this background is primarily muons most of which can be removed with an appropriate set of cuts (see section 8.1.3.) For the ν_e appearance analysis, cosmic-ray photons will be the background source as they can mimic electron-like showers in the detector. To help reduce this background, the building that houses the far detector was designed with a concrete enclosure supporting an overburden composed of barite (barium sulfate.) This mineral was chosen for its availability and its photon stopping power due to its high Z nuclei. Together, the 122 cm of concrete and 15 cm of barite provide an equivalent thickness greater than 12 radiation lengths which is sufficient to reduce the expected background from cosmic-ray photons to an acceptable level.

3.2.3 The Near Detector

The NOvA near detector is located on site at Fermilab in a small cavern adjacent to the existing MINOS underground area (see figure 3.7.) It sits 105 m below the surface and 1015 m from the NuMI target hall, and therefore sees a much lower flux of cosmic rays and a much higher flux of neutrino events than at the far detector. The detector is roughly 50 times smaller than the far detector with only 20,192 cells and is approximately $4 \text{ m} \times 4 \text{ m} \times 16 \text{ m}$ bringing its total mass to 300 tons. Much like the far detector, the neutrino beam enters the near detector at an angle of 3° but these neutrinos are headed downwards instead

of upwards.




Figure 3.7: Overhead view of the location of the NOvA near detector in the MINOS underground area.

The smaller near detector is functionally identical to the far detector with two important distinctions. First, the electronics for the near detector have been designed to sample each channel every 125 ns, which is four times faster than the sampling rate at the far detector. The data pileup in the near detector will be greater than in the far detector due to the fact that the near detector sees 5-10 neutrino interactions in a 10 μ sec window whereas the far detector sees 60-70 cosmic rays in a 550 μ sec window spread out over an object that is 50 times larger. Faster sampling helps improve the timing resolution of hits in the near detector which in turn improves our ability to handle pileup (see section 6.3.) Second, since the near detector is much smaller, the downstream end has been designed to include higher density materials to help muons range out more quickly, increasing its effective fiducial volume. This “muon catcher” is composed of eleven layers each of which contains a 10 cm thick plane of steel followed by two orthogonal planes of standard NOvA cells. For economic reasons, the steel for the muon catcher was re-purposed from the muon catcher used with the NOvA prototype detector, NDOS. This means that the near detector muon catcher is the same width as the rest of the detector, but only two thirds of the height (see figure 3.8.)




Figure 3.8: Technical drawing of the near detector showing the smaller size of the muon catcher (right end.) For picture clarity, only some planes in the main body of the detector have been drawn.

3.3 The Smallest Unit of NOvA Readout

3.3.1 The NOvA Cell

The smallest physical unit with which we read out data in the NOvA experiment is the cell as shown in figure 3.9. The basic design is a long plastic tube filled with a liquid scintillator with a plastic fiber to capture the scintillation light. This light is transmitted down the length of the fiber and is converted into an electronic signal by an avalanche photodiode (APD) at the far end of the cell.

The cell itself is made out of extruded PVC, and has interior dimensions of 3.8 cm transverse to the beam, 5.9 cm parallel to the beam, and 15.5 m in length for the far detector and 4.0 m for the near detector. The thickness of the cell walls varies across different types of cells from 2 to 5 mm due to structural and load-bearing considerations. Scintillation light typically bounces off of the cell walls an average of 8 times (determined from simulations) before it is absorbed by the transmitting fibers, so titanium dioxide was added to the PVC to increase the cell wall reflectivity. Cells are not extruded individually but together in groups of 16 as one plastic object (see figure 3.9.) Altogether, the PVC

used to make the cells represents roughly 35% of the mass of the detector.




Figure 3.9: **Left:** Schematic of the PVC cell filled with liquid scintillator. A charged particle induces scintillation light (blue line) that bounces around until it is absorbed by the transmission fiber (green line.) The light is then transmitted through the fiber to the APD. **Right:** Cross section of an extruded PVC group of 16 cells. The dimensions shown are the exterior dimensions of the cell.

The scintillator inside the cells represents most of the remaining 65% of the detector mass and is primarily mineral oil with 4.1% pseudocumene as the scintillant. The peak of the scintillation light is emitted at 375 nm, so a mixture of PPO and bis-MSB is added to shift these wavelengths to 425 nm to match the absorption spectra of the transmission fibers. The attenuation length of the scintillator was measured to be 5 m and from Monte Carlo simulations, the average path length traveled by a photon before being absorbed by the fiber was found to be 40 cm. Some amount of Stadis-425 is also added to the scintillator to increase its conductivity, which helps reduce the build up of static charge that occurs when filling the PVC cells with the non-conductive mineral oil.

The transmitting fiber is 0.7 mm in diameter and is roughly twice the length of a cell. It is composed of a core of polystyrene with R27 dye added as a wave-shifter to push the peak

of the transmitted light to 520 nm where the attenuation length of the fiber is longer. This core is surrounded by two layers of acrylic each of which has a higher index of refraction than the core to facilitate total internal reflection of the transmitted light. Each fiber is looped within the cell with both ends being read out by the same APD pixel. This allows for scintillation light captured by the fiber (which is transmitted in both directions) to be more efficiently detected. Light with a wavelength of 520 nm that travels 15.5 m in a fiber is attenuated by roughly a factor of 10.

3.3.2 The Avalanche Photodiode

The detection of photons transmitted by the fibers is done with an avalanche photodiode (APD). The NOvA APDs are custom built by Hamamatsu to match the physical design of our fiber readout. Each is composed of an array of 32 pixels with each pixel reading out both ends of a single fiber (see figure 3.10.) A single APD reads out 32 cells or exactly one module. The quantum efficiency is 85% for light between 520 and 550 nm, corresponding to the peak wavelengths transmitted by the fibers. To reduce thermal noise, each APD is cooled to -15°C by a thermo-electric device and the excess heat is taken away by a continuously flowing supply of cold water.




Figure 3.10: Left: The custom designed NOvA APD showing the 32 pixel array. **Right:** Close up of the end of the module that interfaces with the APD in which all 32 fibers have been bundled together. In this picture, an individual fiber has been illuminated.

The design of the readout system is such that the trigger threshold and pedestal level can be set for each APD pixel individually. The pedestals are determined by measuring a baseline noise level for each channel by rapidly sampling them when running in a non-triggering mode. This process is known as a “Digital Oscilloscope” or DSO scan. Determining the right triggering threshold to use for each channel is a slightly more complicated decision. A channel’s threshold should be low enough to maximize the detection efficiency at the far end of the 15.5 m long cells, but high enough to prevent degrading our data quality by flooding valid information about real physics with meaningless noise. More importantly, the thresholds should not be so low as to prevent the electronics from being able to keep up with a data rate that is too high. The data acquisition system can handle a maximum data rate equivalent to 6.2 kHz per channel which is roughly 50 times greater than the nominal rate caused by cosmic ray muons in the NOvA far detector. For each channel, the distribution of signal amplitudes from noise is measured and the width of this distribution is used to determine the threshold. Currently, the thresholds are set to 3.9 standard deviations above the pedestal level for each channel resulting in an overall data rate that is comfortable for the data acquisition system [28]. The end result is that the majority of the data we read out is noise, but this allows us to detect low levels of light from real physics at the far end of the NOvA cells. The NOvA reconstruction software described in chapter 6 has proven itself robust enough to be able to reliably remove high levels of noise at the offline stage, leaving typically less than 1 noise hit in spatial and temporal coincidence with a cluster of hits from a neutrino event.

Each APD is connected to a front-end electronics board (FEB) which serves as the intermediary between the APD and the data acquisition system. The APDs are read out continuously by the FEBs removing the need for an external trigger. Internally, the FEBs automatically handle the pedestal subtraction for each channel as well as pulse shaping

for every measured signal. The signal pulses are shaped to give them a characteristic rise time of 460 (140) ns and a fall time of 7000 (4500) ns for the far (near) detector. This is done for three reasons. First, it allows for light detected from individual scintillation photons to be integrated into a single signal. Second, since the FEBs use a multiplexer to sample each channel once every 500 ns for the far detector and every 125 ns for the near detector, the shaping allows for a signal that occurs between two samplings to be seen. Third, when a signal is triggered, we read out a total of four samples for the triggered channel that included the sample that was triggered and the three samples immediately before that. These four samples can then be fit offline using the known pulse shaping parameters to obtain improved timing resolutions on the order of 10 ns (see section 6.1.4 for more information.) This process is known as “multi-point readout.”

3.4 The Data Acquisition System

3.4.1 Overview of the DAQ

The NOvA electronics operate in a continuous readout mode with all APD information being recorded in a buffer farm. The basic purpose of the data acquisition system (DAQ) is to coordinate the readout of all APDs into one concentrated data stream to be transmitted to the memory banks of this buffer farm for temporary storage. Here the data can await a decision about whether it should be kept or rejected. These decisions can be made by online processing modules or by receiving a trigger from an external source such as the accelerator signal from Fermilab indicating a successful NuMI beam spill.

The DAQ system begins with the front end boards (FEBs) that interface with each APD, with one FEB servicing exactly one APD. The FEB takes each shaped, pedestal-subtracted, and time-stamped signal pulse from an APD pixel, digitizes it into 12 bits, then transmits the digitized data to a data concentrator module (DCM). Each DCM communicates with

up to 64 FEBs. The DCMs collect all of the information received from their FEBs within a 50 μsec window and transmits that packet of information (called a “microslice”) to a buffer farm where the data can await a decision about whether or not it will be kept.

3.4.2 The NuMI Triggers

At the NOvA far detector, the ratio of expected signal events to cosmic rays is low enough that a data-driven trigger system is insufficient to ensure that all neutrino events are recorded. Therefore, an external trigger must be issued that corresponds to the occurrence of a NuMI beam spill. When a successful NuMI beam spill is generated at Fermilab, a trigger signal is time stamped and a network packet is sent to both detectors (with the far detector signal corrected for time of flight) so that the correct data can be recorded from the buffer farm. Each NuMI spill trigger consists of a time stamp (T_0) and a spill duration (ΔT .) When one of these triggers is received by the DAQ, the buffer nodes are searched for all microslices from all DCMs that touch the trigger window (defined as the time between T_0 and $T_0 + \Delta T$.) That data is then assembled together into a “data event” and written to disk for permanent storage. For the NuMI beam spills, T_0 occurs roughly 218 μsec before the actual 10 μsec beam spill and ΔT is 500 μsec . This results in the NuMI beam data being roughly centered in time in the data event with a reasonable amount of out-of-time data which can be used for side-band assessment of backgrounds. The size of the out-of-time window on either side of the beam spill also serves as comfortable padding in the event that our timing system is discovered to have drifted by small amounts.

3.4.3 The Timing System

To record accurate time information, NOvA requires a timing system to ensure that the internal clocks of various DAQ hardware components are accurate and remain in sync with each other. Naturally, it is also required that both the near and far detectors (which are

separated by 810 km) be synchronized to an absolute wall clock so that information from the NuMI triggers will be accurately recorded. The absolute wall clock is derived from a GPS signal and has been verified to be correct through the laborious process of moving an atomic clock synced to the MINOS far detector (verified to be in sync with the NuMI beam) from the Soudan mine to the NOvA far detector in Ash River, MN and back [29].

Both the near and far detectors use two identical timing systems with one being the primary and the other serving as a backup. Each system has a master timing unit that is connected directly to a GPS antenna from which it derives absolute time. The master unit drives the clock ticks for the entire system in addition to issuing reset and synchronization commands when necessary. All timing information derived from the master unit is transmitted to the DCMs where it is distributed to each of the FEBs. By issuing timing pings from the master unit throughout the system, timing offsets for each DCM (caused by different lengths of cable) can be reliably determined and applied to the data offline. These DCM offsets have been checked and validated by comparing the times for hits on tracks that cross DCM boundaries to the expected values. For more information on the timing system, see [30].

CHAPTER 4

Performance of the NOvA Detectors

An important part of ensuring that a particle physics experiment is a success is appropriate monitoring of detector performance. Outlined in this chapter are the tools used for this task, beginning with a description of the online monitoring system which provides live feedback on basic performance metrics. Next I described the nearline monitoring system which processes the data within an hour of it being written to disk. The nearline provides a record of various metrics that can be used to monitor detector performance over longer periods of time and for decisions about which data is of high enough quality to be used for analysis. Both the online and nearline monitoring systems have been an integral part of detector commissioning and are monitored by shifters on a regular basis. Following this I discuss the metrics used for data quality decisions and the cuts applied to these metrics to select periods of good detector performance.

The following is a list of terms used liberally throughout this chapter. Some of them are defined elsewhere, but to ease the reader's burden, I have collected their definitions here. For more information about the physical structure of the detector or the DAQ system, see chapter 3, specifically section 3.4.

- **pixel:** The smallest unit of NOvA readout. One pixel reads out one NOvA cell.
- **APD:** “Avalanche Photo-Diode” - One APD is used to collect the data from 32 pixels.
- **FEB:** “Front End Board” - The FEB performs real-time manipulation of the APD

signals such as pulse-shaping and digitization. Each FEB services exactly one APD.

- **DCM:** “Data Concentrator Module” - One DCM collects the data from up to 64 FEBs.
- **diblock:** A physical unit of 64 detector planes. DCMs are always associated with a specific diblock.
- **data event:** All detector activity associated with a specific trigger. A data event is typically 550 μ sec of readout.
- **subrun:** A collection of data events that occurred in succession (often represents tens of minutes of real time data taking.)
- **run:** A collection of subruns that occurred in succession (often represents hours of real time data taking.)

4.1 Online Monitoring

The online monitoring (OnMon) software provides close to real time displays of the detector data. It is designed to leverage the framework used for offline processing by receiving events from the live data stream. The OnMon suite has two main programs: a producer that fills a set of histograms to track data quality, and a viewer that organizes and displays these histograms in a meaningful way. The producer and the viewer are designed to be run together (although they can be run independently) in an environment that gives shifters immediate access to the data as it is being taken.

4.1.1 The OnMon Producer

The producer is the main work-horse of the OnMon framework. It can receive data either by spooling over a file on disk, or directly from the live event dispatcher (part of the DAQ system) through an IO module that accepts data from a socket and casts it into the

format expected by the offline framework. The purpose of the producer is to sift through the raw detector data, organize it into meaningful histograms, and provide access to those histograms by writing them to disk for later use and by serving them to copies of the OnMon viewer running at the same time. When the producer receives the data, it unpacks it into summary classes that are passed to histogram-making classes. The goal of this organization is to do all of this within a single set of loops, rather than force each histogram-making class to re-implement the unpacking, making the code cleaner, more efficient, and more easily extendable.

There are four summary classes: `Raw Event`, `Data Block`, `Microslice`, and `Nanoslice`, which represent the “nested” organizational structure of the raw data. The raw event contains all of the information for the whole event which is broken into multiple data blocks, each data block is composed of multiple microslices, and each microslice contains multiple nanoslices. Each nanoslice summary is therefore associated with a specific microslice summary and so on. The nanoslice summary represents the “smallest” summary object and contains the readout information from a single triggered channel. This includes the time and total charge for that hit, the hardware address in terms of FEB and pixel number, and a few status bits associated with the state of the FEB. While nanoslices are associated with individual channels, microslices are associated with individual DCMs. The microslice summary contains all of the information reported by a DCM during a 50 μ sec interval, which includes the detector, diblock, and DCM, as well as error codes associated with the state of the DCM. Data block summary objects are rather simple, and only contain a single error code pertaining to missing data. Lastly, the raw event summary object contains the highest-level information about the event such as the run, subrun, and event numbers, the trigger type, and an error code set by the DAQ about the completeness of the entire event.

Each of the histogram-making classes must be a subscriber to one of these four levels

of summary information. After each summary level is unpacked by the producer, a call is made to the subscribers for that level, passing them the appropriate summaries. Because of the nested structure of the raw data, a subscriber to one summary level receives the information from all of the levels above it as well. For example, a microslice summary subscriber will receive the microslice summary object in addition to the data block and raw event summary objects. With the appropriate summary objects in hand, each histogram making class can then fill its specific set of histograms.

Since the channel hits represent the majority of the information reported by the detector, most of the histogram making classes are nanoslice summary subscribers. One such example is the class that makes the “hit maps” which are two-dimensional histograms that record channel information displayed in a grid of hardware coordinates, which can be either diblock, DCM, FEB, and pixel, or plane and cell. Two examples of hit maps that record total hit rates are shown in figure 4.4 in the plot gallery (section 4.4) at the end of this chapter. Descriptions of the major relevant metrics recorded in the OnMon histograms can be found in section 4.1.3.

In addition to storing the current state of the detector, the producer automatically retains a series of histograms that represent recent snapshots of the detector state. These are referred to within OnMon as “look-back” histograms. Each histogram that the producer makes has associated with it a schedule used to refresh the look-back copies and a number representing how many look-back copies to keep. The refresh schedules can be for example, every 5 minutes, 30 minutes, 24 hours, or at run or subrun boundaries. When a set of histograms with N look-backs is to be refreshed, the $N - 1$ look-back is copied to the N look-back histogram, the $N - 2$ look-back is copied to the $N - 1$ histogram, etc. until the current version of the histogram is copied to the first look-back, leaving the current histogram blank and ready to be filled with new information. As an example, the hit maps

displayed in figure 4.4 are refreshed every 30 minutes with 6 look-back copies kept in history. Refreshing the look-back copies for each histogram according to their individual schedules, is done by a single histogram management class. This class maintains control over all of the histograms that the producer makes and handles requests made by the histogram making classes such as the booking of a new histogram.

Many histogram parameters, including the look-back schedule, are stored in a text file that is read by the producer at run-time. This run-time configuration is used to handle differences between histograms required for the near and far detectors, so that identical versions of the OnMon software can be easily deployed for each detector. Organizing things in this way also makes changing or adding new histograms relatively easy by keeping all histogram information, which needs to be accessed from multiple places within the code, stored in one location. The information in this text file is parsed into a form that is made available to the histogram management class so that booking new histograms is as easy as possible. The information about each histogram in the text file includes things like the binning for each axis, titles, and a variety of display options used by the viewer (described in section 4.1.2.)

4.1.2 The OnMon Viewer

The OnMon viewer is a GUI that makes heavy use of many of the classes and objects available through the ROOT software platform [31, 32]. It can be operated by itself by reading a set of histograms from a file on disk that the producer created, or by communicating directly with an instance of the OnMon producer running on the same machine. This communication is facilitated through a block of shared memory created and managed by the producer. Multiple instances of the viewer can connect as clients to the same shared memory segment. However in typical running, only one copy of the viewer is attached to

the producer at any given time.

A complete picture of the OnMon viewer window is shown in figure 4.1. The majority of the window is taken up by the canvas displaying the current histogram. In this example, the histogram being displayed shows for all hits reported by a specific DCM, the hit time with respect to the event trigger time $T0$ vs. the hit charge (ADC.) The empty box immediately above the canvas will display information about the coordinates and contents of a histogram bin selected by the user by single clicking on the canvas. To the left of that box are a variety of buttons that allow the user to navigate forwards and backwards through a history of recently viewed histograms. If the viewer is not reading histograms from a file, then continually updated information about the current status of the producer is displayed in the text seen directly above the navigation buttons.

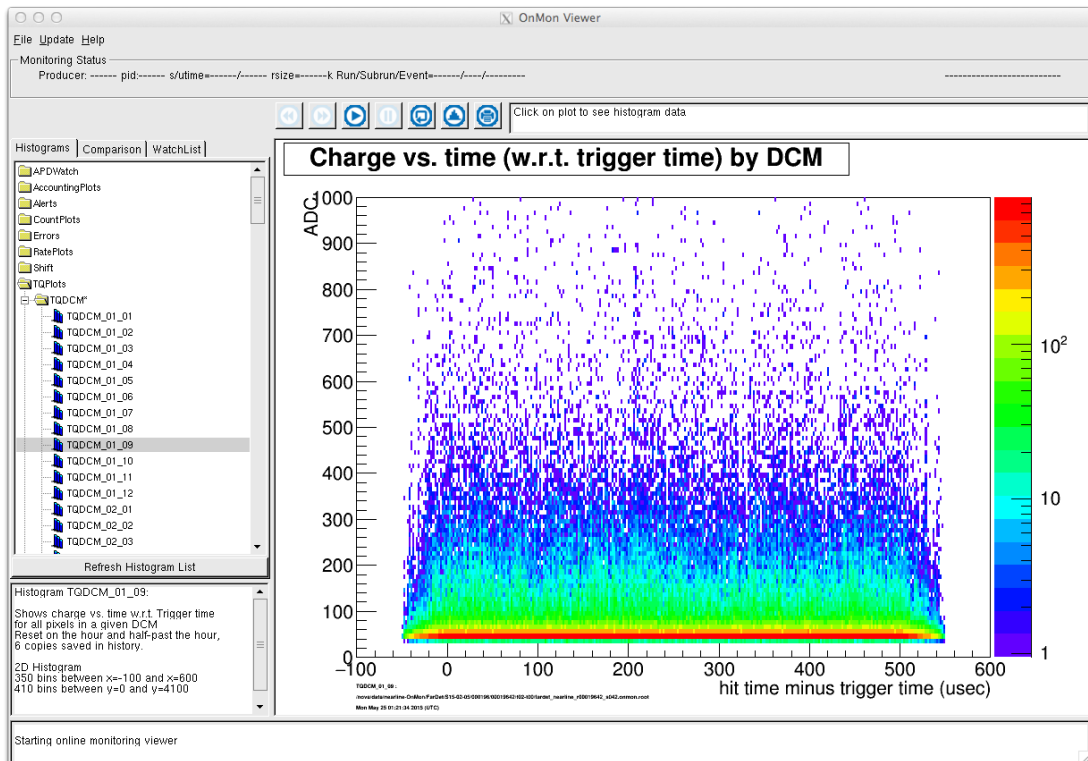


Figure 4.1: The OnMon viewer.

The upper box immediately to the left of the canvas contains three tabs with different

features. The tab being shown (labeled “Histograms”) allows the user to easily browse through all of the available histograms which have been organized into folders. The other two tabs, labeled “Comparisons” and “Watch List”, offer special features that are described below. The box immediately below displays additional information about the current histogram useful to the user. This information includes a descriptive sentence about what the histogram means and how it is filled, the look-back schedule, the number of look-back histograms kept in memory, and information about how the histogram is binned. The last box spanning the width of the viewer window at the very bottom, displays general messages from the viewer to the user, including errors and warnings.

The OnMon viewer has a variety of other features that provide assistance to the user that are worth mentioning. A short list of the most important features includes the following:

1. **Comparison Plots:** These plots are generated on demand by the user for any of the histograms made by the producer. Reference histograms used to generate comparisons can either be from the look-back list or loaded from an example file on disk. Comparison options include differences, ratios, and asymmetries which are computed bin-by-bin for the entire histogram.
2. **Watch List Plots:** A suite of plots can be generated for only a selected set of hardware, allowing the user to monitor the behavior of individual pixels or FEBs, etc. that may be acting suspiciously. The desired hardware is selected from the viewer, and a command is issued to the producer through the shared memory block with instructions to begin filling plots for that hardware. The watch list plots include distributions of times and charge for all hits reported by the selected hardware.
3. **Bin Content Projections:** This option will generate a one-dimensional histogram filled with the bin contents of the current histogram, and can be coupled with options like the comparison plots.

4. **Bin Content Descriptions:** As mentioned above, single-clicking within the main canvas on a histogram bin, will display the bin contents and bin coordinates in the box immediately above the canvas. In the case of a hit map, the bin coordinates are translated into the appropriate hardware coordinates (diblock, DCM, FEB, pixel, plane, and cell.)
5. **Drill-Down Navigation Assistance:** All of the hit maps that OnMon produces are created at two levels: displaying the entire detector and showing a more detailed view for each individual DCM (figure 4.4.) By double-clicking on a specific DCM in the detector-level hit map, the user will automatically be directed to the corresponding DCM-level hit map. For the NOvA far detector, which has 168 DCMs, this option makes navigating hit maps much easier.

Much of the information about the individual histograms contained in the text file described in section 4.1.1 is relevant to the viewer as well. As is the case with the producer, this file is read and parsed at run-time so that the viewer has immediate access to important information specific to each histogram. For the viewer, the relevant information includes any text to be displayed in the description box in the lower left corner of the viewer, information about which organizational folder(s) each histogram belongs in, plot drawing options such as a zoom range or the use of a log scale for a specific axis, and special drawing options such as the additional lines and labels drawn on the plots in figure 4.4. Collecting these options all in one place allows for quick and easy changes to any of the OnMon histograms.

Another advantage of using a single text file to specify all of the histogram information is that by specifying a different set of histograms and histogram configurations in a different text file, the OnMon viewer can display histograms created by other programs. This flexibility allows any program that generates histograms about the detectors to take advantage of the viewer features listed above, especially if those histograms are arranged in the style

of the hit maps. For example, when a DSO scan is performed, the results are cast into the format of the OnMon hit maps, written to a file, and opened with the OnMon viewer. A shifter can perform visual inspections of thresholds and pixel masks generated by this scan before applying them to the detector, using the familiar interface of the OnMon viewer.

4.1.3 OnMon Metrics

OnMon keeps track of a large number of histograms, 76 for the whole detector plus 2 for each diblock and up to 23 for each DCM. This brings the grand total to a maximum of 406 for the near detector and 3968 for the far detector. Only 78 of these require a unique description since the each of the 23 DCM-level plots is a version of the one of the main 76 made for the whole detector. Of course, 76 descriptions is still too much information for this document. Described below are only the most useful or interesting metrics, specifically the ones used to make decisions about data quality. For space conservation, all of the plots described in this section appear in the plot gallery (section 4.4) at the end of this chapter.

1. **Hit Counts:** Hits are counted at three hardware levels: DCM, FEB, and pixel and are displayed in the hit map style. Since a total hit count does not give a complete picture of the information being reported by by a specific piece of hardware, these hit counts are broken into three additional categories by the charge (ADC) reported for each hit: low ($\text{ADC} < 175$), “MIP” ($175 \leq \text{ADC} \leq 3200$), and high ($3200 < \text{ADC}$).
2. **Live Time:** A count of the total live time summed over all processed events.
3. **Hit Rates:** For each hit count plot described above, there is a corresponding hit rate plot computed by dividing the hit counts by the total live time. While the hit count histograms are filled continuously, the hit rates are recomputed after every 100 processed events. Two examples of hit rate plots are shown in figure 4.4.
4. **FEB Rates vs. Time:** This plot displays the total hit rates for each FEB as a function of time (see figure 4.5.) Each horizontal stripe is a single FEB, indicating

fluctuations in performance and possible drop-outs.

5. **Reporting Hardware vs. Time:** For each data event, OnMon counts the number of hardware elements (DCM, FEB, and pixel) that report information, and plots those numbers vs. time. These plots can be sensitive to hardware drop-outs and can show long-term fluctuations in the detector noise or data rates.
6. **Hit Times Within the Trigger Window:** This plot displays the time for all hits with respect to the start of the trigger window (T_0 .) Two versions of this plot are made, one showing the full trigger window, and one zoomed in on the NuMI beam spill window. For the near detector, large-scale timing problems can be evident by spikes of activity outside of the normal NuMI beam spill window. An example of the zoomed version of this plot can be seen in figure 4.12. The version showing the full trigger window can be seen in figure 4.14. This particular example was taken during a time when the beam spill was outside of the expected window (indicated by the vertical lines.)
7. **Microslice Counts vs. Time:** This plot displays the total number of microslices in each data event as a function of time. For a typical 550 μ sec data event, this number should be $11 \times$ the number of DCMs being read out. A departure from this number is an indication that some DCMs are having trouble reporting all of their information. An example from near detector data (where the total number of microslices should be 154) is shown in figure 4.5.
8. **Alerts and Errors:** The status bits (alerts) and error codes (errors) are reported by different hardware elements as discussed in section 4.1.1. Errors represent information reported intermittently whereas alerts, which represent the status of a particular hardware element, are reported continuously. OnMon keeps track of 5 different alerts and 10 errors. Both alerts and errors are recorded in two different ways: by which hardware element reported the alert/error in the hit map style, and by a total count

in each category vs. time. An example of each of these can be seen in figure 4.6. In this example, the plot of all alerts vs. time shows that an FEB somewhere in the detector has the alert bit “FEB buffer empty” set high. The corresponding hit map shows the offending FEB in red.

4.2 Nearline Monitoring

While the online monitoring system processes the live data stream, the nearline monitoring system processes the data within a few minutes of it being written to disk. The main purpose of the nearline is to create a permanent record of detector performance, including some basic reconstruction metrics, for each subrun. This permanent record serves two primary purposes. First, it is used for long-term (weeks to months) detector performance monitoring. Displaying the long-term detector behavior, including some simple reconstruction metrics, offers a different view of the data not apparent through the live version of OnMon. Second, this permanent record is used for general data quality decisions, including determining which subruns represent good data and which channels should be masked prior to reconstruction. The entire nearline system functions through a series of fully automated scripts that handle everything from data processing and transfer through the generation of performance plots that are regularly published to a website.

4.2.1 Nearline Processing

The nearline metrics are generated by processing the raw data after it has been written to disk through the OnMon producer and a simple set of reconstruction algorithms. As discussed in section 4.1.1, the OnMon producer leverages the offline framework to process data and can be run independently of the viewer using files on disk as input. Even though a version of the OnMon producer is often processing the live data stream, rerunning it over

the data on disk is necessary for a couple of reasons. First, the live version depends on being properly operated by the shifter and is often started after data taking has already begun. Second, only a fraction of the total data is seen by the live version due to the fact that the rate at which OnMon processes data and the rate at which the dispatcher serves data to OnMon are not the same. Nearline reprocessing of the raw data through the OnMon producer provides a controlled way of generating reliable OnMon metrics that includes all data taken.

The reconstruction used for nearline processing includes a clustering algorithm and a simple tracker. The clustering algorithm (described in detail in section 6.3) is designed to group hits together based on their spatial and temporal proximity, isolating clusters that originated from individual physics events and separating hits judged to be noise. These clusters are referred to as “slices” and are typically individual neutrino interactions or cosmic rays. All hits labeled as noise are put into a single cluster called the “noise slice.” The tracker was written specifically for quick processing of cosmic ray data and is appropriately referred to as “cosmic tracker.” It works by fitting a straight line to the slice hits in each view to create two, two-dimensional tracks which are then merged into a single three-dimensional track. If the quality of the 2D tracks is sufficiently low, which can be caused by a variety of different problems, then the 2D tracks can remain unmerged. This tracking method works well for cosmic rays given their highly linear topology. Both algorithms are relatively simple and very quick, and are therefore appropriate for generating nearline reconstruction metrics.

Nearline processing takes place across multiple computers that are dedicated to this task. On these machines, an automated set of scripts checks for new data on disk every few minutes, and spawns jobs to process any new files. These scripts have checks in place to ensure that files are not processed more than once and that a maximum number of running

jobs is not exceeded. The output from these nearline processing jobs includes an OnMon histogram file and a separate file containing histograms filled with reconstruction metrics based off of slices and cosmic tracks. An additional set of automated scripts handles the regular transfer of these files to a permanent storage location and clean up of old files on the local disks.

The raw data files processed by the nearline contain data events separated by trigger type. Since different triggers can represent data with different topologies (number of hits per slice, average charge deposited per hit, etc.) a decision must be made to process triggers that present a clear picture of the detector behavior so that accurate data quality metrics can be obtained. A good choice of highly consistent data for both of the detectors are cosmic rays. Since the far detector sits on the surface, any 550 μsec readout of data will contain a snapshot of cosmic ray activity. Therefore, the far detector nearline processing combines both the calibration and NuMI beam triggers. Cosmic rays are rare enough in the near detector that they must be captured through one of the data-driven trigger streams (called the “DDT-Activity” trigger.) The NuMI beam spill triggers also represent important yet qualitatively different information about the near detector performance. Processing the NuMI triggers can give insight into things like the detector response as a function of beam power and the synchronization of the timing system with respect to the NuMI beam spills. The beam events themselves occur in high enough quantities per subrun that they can be used in place of the DDT-Activity triggers to determine many of the OnMon metrics such as the hit rates. Currently, we process both the NuMI and DDT-Activity trigger files independently but we use only the information from the NuMI trigger files to assess near detector data quality.

Once the nearline output files have been transferred to the permanent storage location, they can be used for a variety of tasks. The three primary uses are to generate performance

plots that are posted to a monitoring website, to generate a list of good subruns to be used for analysis, and to determine bad channel masks. The monitoring website is discussed in section 4.2.2 and determining the list of good subruns is discussed in section 4.3. The bad channel masks are determined from the hit rate histograms like the ones shown in figure 4.4 in the OnMon files. The information in these histograms is used to determine which channels are too hot or too cold and should be masked prior to reconstruction. Through these three operations, the files created by the nearline processing framework have provided the foundation for all data quality decisions and have been absolutely essential for commissioning both of the NOvA detectors.

4.2.2 Nearline Metrics

Once the nearline output files have been transferred into permanent storage, they can be used to assemble plots that monitor the long-term performance of the detectors. On a dedicated machine, another series of automated scripts looks for new nearline files and spools over them to generate a set of performance metric plots that are regularly posted to a website. These scripts make plots for three time intervals: the past 24 hours (which are refreshed every 10 minutes), the past week (refreshed every hour), and the past 30 days (refreshed once per day.) The total time delay between the raw data originally being written to disk and information about that data appearing in the nearline 24-hour plots, is on the order of one hour.

There are two nearline websites available, the main nearline page and what is referred to as the nearline “frontpage.” The main nearline page displays a total of 56 plots broken into 8 categories for each detector and for each of the three time intervals. The near detector contains one additional category with 4 plots related to the NuMI beam time peak. This page includes a short description of each plot, links to related websites, and java-enabled

buttons that allow the user to easily select and view only a specific set of plots. The nearline frontpage includes just the 4 “most important” plots, collected in one place for a simple overview of detector health. Each of these plots is accompanied by an example of what the plot should look like (i.e. through a reference plot) with a link to instructions on what actions to take if this is not the case.

Described below is a sample of some of the most informative nearline plots. All of the corresponding pictures have been assembled in the plot gallery (section 4.4) at the end of this chapter.

1. **FEB Hit Rate Spectra vs. Time:** This plot shows the FEB hit rate spectrum computed from the first histogram shown in figure 4.4, for each subrun as a function of the subrun start time. An example of typical behavior during good running is shown in figure 4.7. A “bad” example, taken during early commissioning when the detector electronics were not yet being cooled, is shown in figure 4.14.
2. **Average Number of Noise Hits per Subrun vs. Time:** This plot shows the average number of noise slice hits for each subrun as a function of the subrun start time. Large scale oscillations (days to weeks) in this plot have been linked to environmental factors such as the humidity level in the detector hall [33]. An example is shown in figure 4.7.
3. **Number of Active FEBs per Subrun:** This plot shows for each subrun, a count of how many FEBs were active (had a hit rate > 0) for that subrun plotted as a function of the subrun start time. Similar plots are kept for the number of active DCMs and active pixels. The example in figure 4.8 shows the typical behavior for the far detector: over a period of a few hours, a small number of FEBs will stop reporting information but can be brought back through issuing a refresh signal from the timing system. Determining which FEBs are dropping out is accomplished with the FEB

Dropout Counts plot.

4. **FEB Dropout Counts:** With information from multiple successive subruns, the nearline keeps track of how many times each FEB switches from a reporting state to a non-reporting state (classified as an “FEB dropout.”) For each FEB, the number of dropouts is recorded in the OnMon hit map style. These FEBs can then be added to a list to be monitored and possibly replaced. An example is shown in figure 4.8.
5. **Percentage of Empty Spills:** For each subrun, the number of events that contain zero hits is counted and plotted as a percentage of the total number of events in that subrun as a function of the subrun start time. This plot normally hovers at or very near zero. An example of a period when the DAQ was experiencing trouble is shown in figure 4.14.
6. **Number of Slices per Spill:** This plot shows the 1D distribution of the number of slices per “spill” (i.e. - trigger) for all spills within the 24-hour, week, or 30-day time interval. An example is shown in figure 4.9.
7. **Average Number of Slices per Spill vs. Time:** This plot shows the mean of the Number of Slices per Spill histogram for each subrun plotted as a function of the subrun start time. An example is shown in figure 4.9. Many types of poor detector behavior (DCMs out of time sync, partially complete or empty events, unusually high noise rates, etc.) can have a large impact on the total number of slices produced.
8. **Trigger and Spill Rates vs. Time:** This plot shows the rate of recorded calibration and NuMI triggers for each subrun as a function of the subrun start time. Superimposed is the delivered beam spill rate determined from a database maintained by the accelerator division. This plot helps ensure that beam trigger signals at both detectors are not being missed. An example is shown in figure 4.10.
9. **Track Fractions vs. Time:** This plot shows the percentage of 2D, 3D, and contained 3D tracks produced by the cosmic tracker for each subrun as a function of the subrun

start time. An example of typical far detector behavior in which $> 99\%$ of tracks are 3D is shown in figure 4.10 and an example from an early commissioning period when this metric was much less stable is shown in figure 4.14. Many types of poor detector behavior (DCMs out of time sync, partially complete or empty events, unusually high noise rates, etc.) can be seen in this plot through an increase the percentage of 2D tracks.

10. **Track Start Positions:** Shown in this plot are XZ and YZ projections for the starting positions of all 3D tracks produced by the cosmic tracker. An example from far detector data, which shows the dominance of top and side entering cosmic rays, can be seen in figure 4.11.
11. **Timing Peak Distribution:** This plot shows the distribution of the hit times for all hits with respect to the trigger time (T_0) taken directly from the OnMon files (see the description of the “Hit Times Within the Trigger Window” plot in section 4.1.3). An example of the version zoomed in on the NuMI beam spill window is shown in figure 4.12. This example was taken from a time when the NuMI beam was running in a 6 + 2 batches, partial slip-stacking mode.
12. **Timing Peak Distribution vs. Time:** Shown in this plot is the timing peak distribution for each subrun as a function of the subrun start time. This plot can show not only variations in the intensity of the beam over time, but any drifts or jumps in the location of the beam window as well. An example is shown in figure 4.12.
13. **PE Distribution for Slices and Tracks:** This plot shows the PE distributions for all hits broken into 3 categories: all hits in the noise slice, all hits in non-noise slices, and all hits on tracks made by the cosmic tracker. “PE” is the number of photoelectrons seen in a given hit (a simple conversion from units of ADC.) An example

from far detector data is shown in figure 4.13. Problems with channel thresholds, APD gains, and increased noise rates can be seen in this plot.

4.3 Data Quality Selection

Because of the way that the nearline processing is done (described in section 4.2.1), the output files are representations of the average detector behavior over a subrun. The major data quality decisions, which are based on the nearline output files, are therefore made on a subrun by subrun basis. A subrun for the far (near) detector is approximately five minutes (one hour) of real-time data taking. Intermittent failures, such as temporary interruptions in the DAQ, can happen during an otherwise stable subrun. Additional data quality cuts are therefore made at the data event level to remove these problems.

4.3.1 Subrun Level Cuts

Many of the OnMon and nearline metrics described in this chapter expose the interesting subtleties of detector behavior and have been extremely helpful in diagnosing problems, and commissioning both of the NOvA detectors. However a simple and reliable set of data quality cuts can be created from just a small subset of the nearline information. The major data quality cuts were optimized for each detector through a series of studies [34, 35]. The cuts for each detector are described below. More information about the subrun-level data quality cuts can be found in [36].

Far Detector Cuts:

1. **Diblock Cut:** This cut is designed to determine how much of the active hardware can be labeled as “good” and works through a bottom-up series of hardware assessments starting with the pixels. A pixel is considered good if the total hit rate is between $10^{0.5}$ Hz and $10^{3.5}$ Hz. An FEB is considered good if the number of bad pixels is < 7 .

A DCM is considered good if the number of bad FEBs is < 9 . A diblock is considered good if all DCMs are good. Lastly, the subrun is considered good if there are at least four good consecutive diblocks.

2. **Median Detector MIP Hit Rate:** For all reporting pixels in the detector, the median hit rate in the MIP range (see the description of the “Hit Counts” plot in section 4.1.3) is required to be between 13 Hz and 23 Hz.
3. **Live Time:** This cut was designed to ensure that a subrun was long enough to obtain reliable information about pixel hit rates, and to remove short subruns which are in general, indicative of bad data. Subruns are required to have a total live time > 1.0 sec, which results in approximately 20 MIP hits per pixel and is roughly 1800 events (if each event is 550 μ sec long.)
4. **Reconstruction:** There are two reconstruction cuts used. The first ensures that the number of slices produced (normalized to the number of active channels) is neither too low nor too high, and the second removes data in which the cosmic tracker produced too many bad tracks. The slice cut requires the number of slices per event per 10^4 active channels to be between 1.2 and 3.2, and the track cut requires the fraction of 2D tracks to be $< 15\%$. Subruns that fail these cuts typically represent extremely noisy data, or data taken with all detector DCMs out of time sync.
5. **File Time Stamp:** This category of cuts removes the very small number of subruns that have experienced failures that can occur within the DAQ when assembling nano and microslices together into data events and writing the events to disk. The cut requires the subrun start time to be before the end time, the time stamp for the subrun to be after the year 2012, and the subrun to have at least one non-empty data event. In figures 4.2 and 4.3, this category is labeled as “other.”

Near Detector Cuts:

1. **Diblock Cut:** The criteria for a “good diblock” is the same as for the far detector with two exceptions. First, not all near detector DCMs have 64 FEBs, so a near detector DCM is considered good if $> 80\%$ of the FEBs are good. Second, for the subrun to be considered good, all 4 diblocks (that is, the whole detector) must be good.
2. **Median Detector MIP Hit Rate:** This is the same as for the far detector, with the accepted hit rate range being 12 Hz to 20 Hz. Since hit rates in the near detector are proportional to the beam intensity, these rates are normalized to a “nominal” NuMI spill of 2.5×10^{13} POT.
3. **Number of Triggers:** The near detector cuts use a count of the number of NuMI triggers seen within the subrun as a proxy for live time.
4. **Reconstruction:** The near detector reconstruction cut only requires the mean number of slices per event to be between 3.5 and 5.5. This number is also normalized to a “nominal” NuMI spill since the number of slices is also proportional to the beam intensity.
5. **Fraction of Empty Spills:** Empty spills (a data event with no slices) can be recorded for the near detector when global DAQ failures occur, or when the NuMI beam is not running. Near detector good subruns are required to have a fraction of empty spills $< 3\%$.
6. **Timing Peak:** This cut determines whether or not there were any shifts in the near detector timing system by examining the times at which beam events occurred. To do this, a fit is performed to a plot of the subrun slice times (the average of all hit times within the slice) to determine the start and end of the NuMI timing peak. The start time is required to be between 217 μsec and 219 μsec , and the end time is required to be between 227 μsec and 229 μsec .
7. **File Time Stamp:** This category is the same as for the far detector.

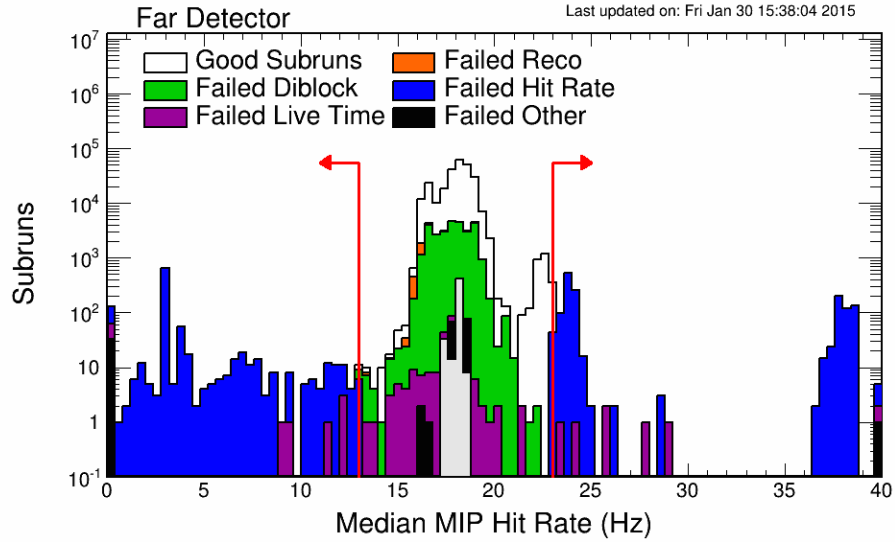


Figure 4.2: Example distribution of far detector subruns in terms of the median MIP hit rate variable. “Other” refers to the file time stamp cut.




Figure 4.3: Example distribution of far detector subruns in terms of the 2D track fractions variable. “Other” refers to the file time stamp cut.

Two examples of the subrun distributions from far detector data with respect to the cuts listed above are shown in figures 4.2 and 4.3. Each plot is color coded by which cut caused each subrun to fail, showing that often a bad subrun fails multiple cuts. Plots of this type are also posted regularly to the main nearline website. For the far detector,

an additional group of subruns identified by shifters as “untrustworthy” due to alternate detector running conditions, are manually labeled as bad. For the far detector, a total of approximately 250,000 subruns were recorded, 72.4% of which were labeled as good. For the near detector, a total of approximately 3,400 subruns were recorded, 86.7% of which were labeled as good. A summary of the number of far and near detector subruns that failed each of these cuts is listed in table 4.1.

cut category	far det. (%)	near det. (%)
file time stamp	0.1	0.1
live time	< 0.1	-
number of triggers	-	6.0
timing peak	-	3.9
empty spills	-	7.0
hit rates	0.7	5.1
good diblocks	23.5	3.2
2D track fraction	0.2	-
number of slices	0.7	4.0
manual removal	3.0	-
total bad subruns	27.6	13.3

Table 4.1: Percentage of subruns for both detectors that failed each of the subrun-level data quality cuts. The percentages do not add up to the totals due to subruns failing multiple cuts.

4.3.2 Event Level Cuts

The first of the data event level cuts is based on the “DCM edge metric” [37]. This metric was designed to catch intermittent failures in which multiple DCMs are out of time sync with the other DCMs in the detector, resulting in broken tracks that could mimic contained events. This occasionally happened at the beginning of a run and was often resolved after a period of unsynced running. The metric is a count of the fraction of the tracks within a data event that successfully cross the boundary between two DCMs. For a typical far detector event, this number is relatively stable near 0.6. A cut of 0.2 was found to be sufficient to

remove out of sync data events for the detector [37]. This cut is only applied to the far detector.

The near detector uses a different metric to remove data events in which information from DCMs was missing. It was found from a sample of near detector subruns that 0.03% of NuMI trigger events had one or more DCMs that reported no hits [38]. To remove these events, the number of hits in each DCM is counted on an event-by-event basis. Events with no hits reported by one or more DCMs are cut. Of course, having a properly functioning DCM that reports no hits within a 550 μsec NuMI beam event can happen by chance if no particles pass through that DCM and no noise hits are reported by any of the APDs connected to that DCM. However, DCMs typically report 10-20 noise hits every 550 μsec , so a properly functioning DCM reporting no hits will be rare enough that cutting these events was judged to be an acceptable loss. In addition to this, a specific group of near detector cells were found to be susceptible to light leaks [38]. This was discovered during a period of running with the lights turned on in the near detector cavern. To remove these events, a cut is applied that requires the fraction of hits outside of the NuMI beam spill window in the affected cells out of the total number of hits to be less than 45%. A physical “light-blocking” mechanism has since then been installed to prevent future light leaks.

A variety of metrics regarding the quality of the NuMI beam are provided to all Fermilab experiments by the Accelerator Division on a spill by spill basis. Several of these metrics were chosen for the event level cuts to ensure that the beam was of high enough intensity and on target [39,40]. These cuts require the number of POT to be $> 2 \times 10^{12}$, the measured current in the focusing horns to be between -202 and -198 kA, the X and Y position of the beam to not stray more than 2.0 mm from the center of the target, and the beam width to be between 0.57 and 1.58 mm.

4.4 Gallery of OnMon and Nearline Plots

This section includes the plots fully described in sections 4.1.3 and 4.2.2.

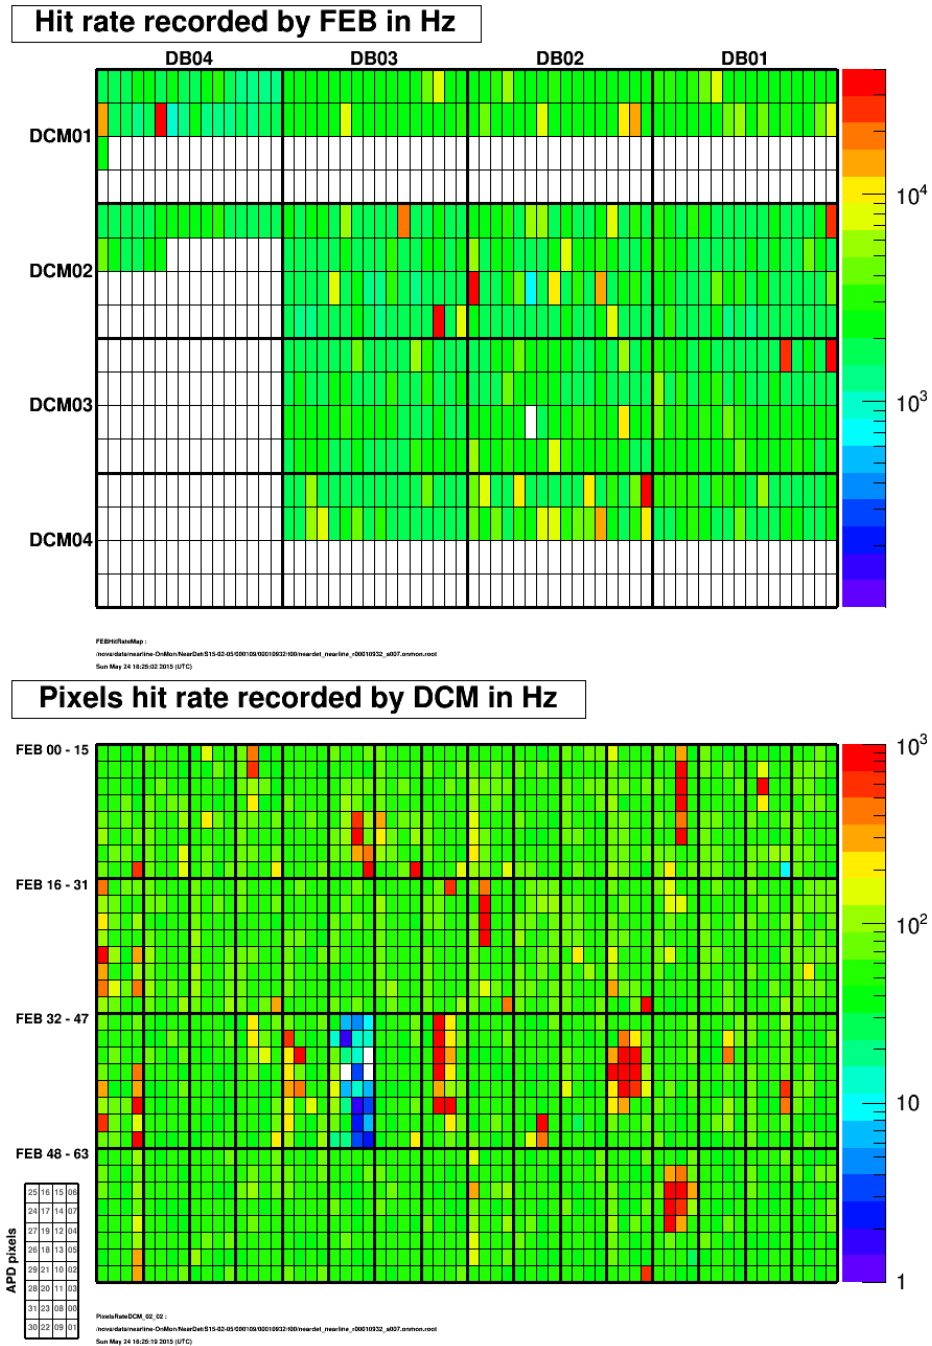



Figure 4.4: OnMon Plots - Hit map for the NOvA near detector displaying the average pixel hit rate for each FEB (top.) Hit map representing DCM 2 in diblock (DB) 2 showing the hit rates for each individual pixel (bottom.) The large boxes represent a single DCM and the small boxes are an individual FEB. White spaces are either non-reporting or uninstrumented hardware.

FEB Hit Rates by Diblock vs UTC Hour



Total number of micro slices per event vs UTC Hour

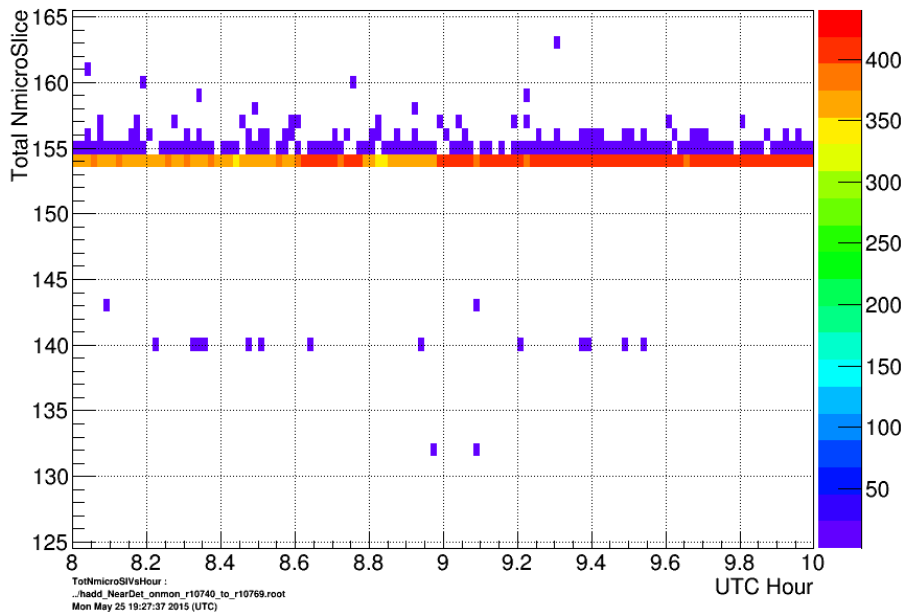


Figure 4.5: OnMon Plots - FEB rates vs. time (top) and total number of microslices per data event vs. time (bottom.)

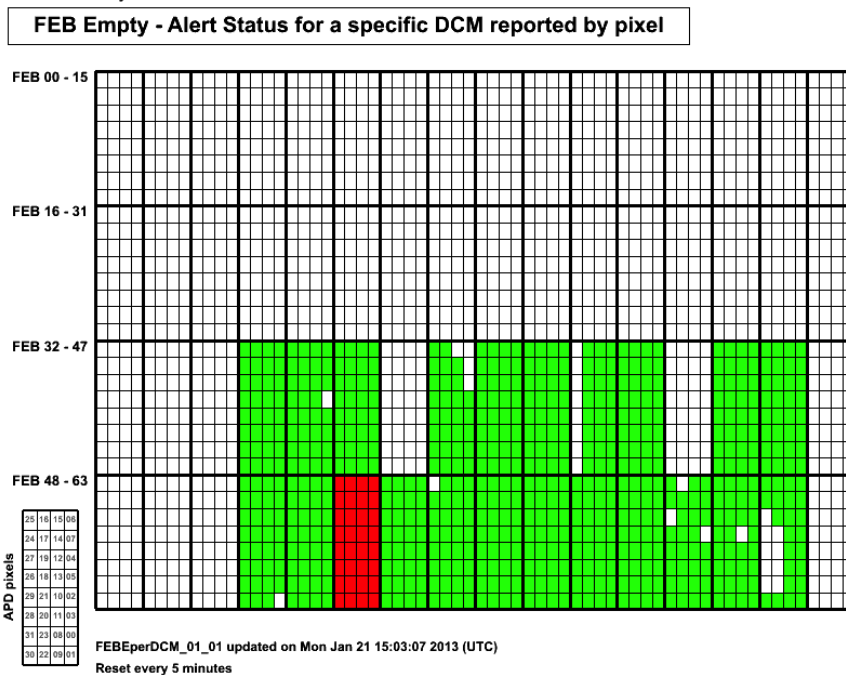
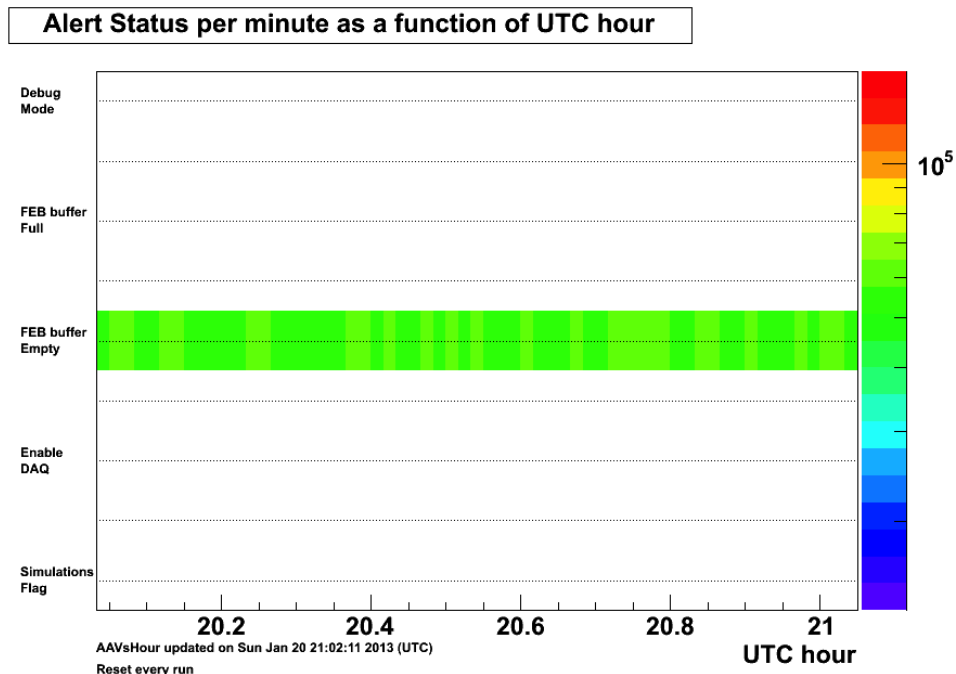


Figure 4.6: OnMon Plots - All alerts reported vs. time (top) and a DCM-level hit map for the “FEB buffer empty” alert (bottom.)


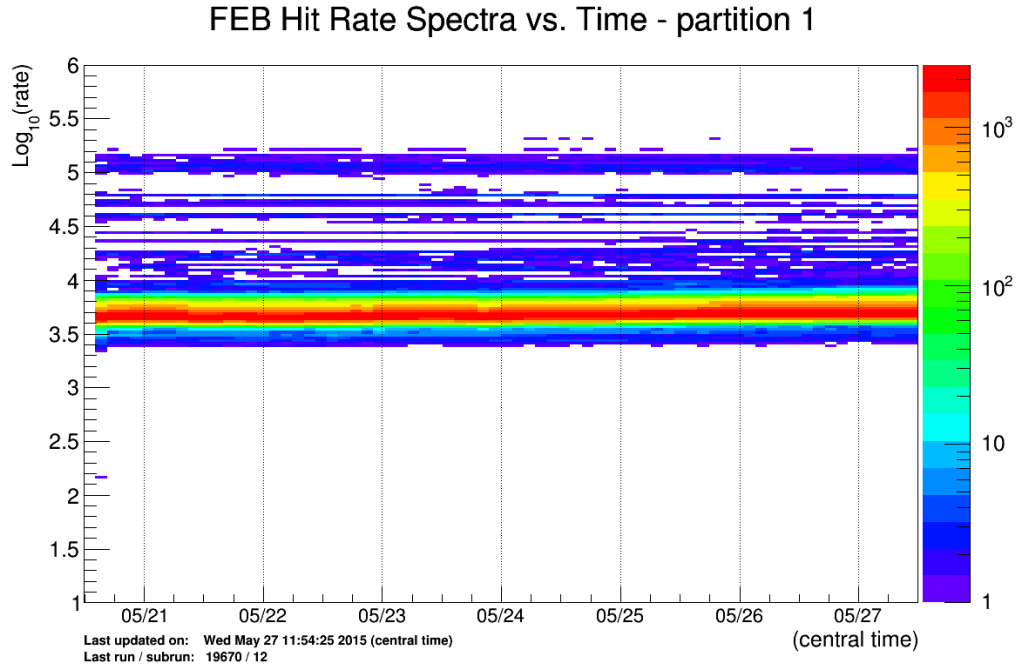


Figure 4.7: Nearline Plots - FEB hit rate spectra vs. time (top) and average number of noise hits per subrun vs. time (bottom.)

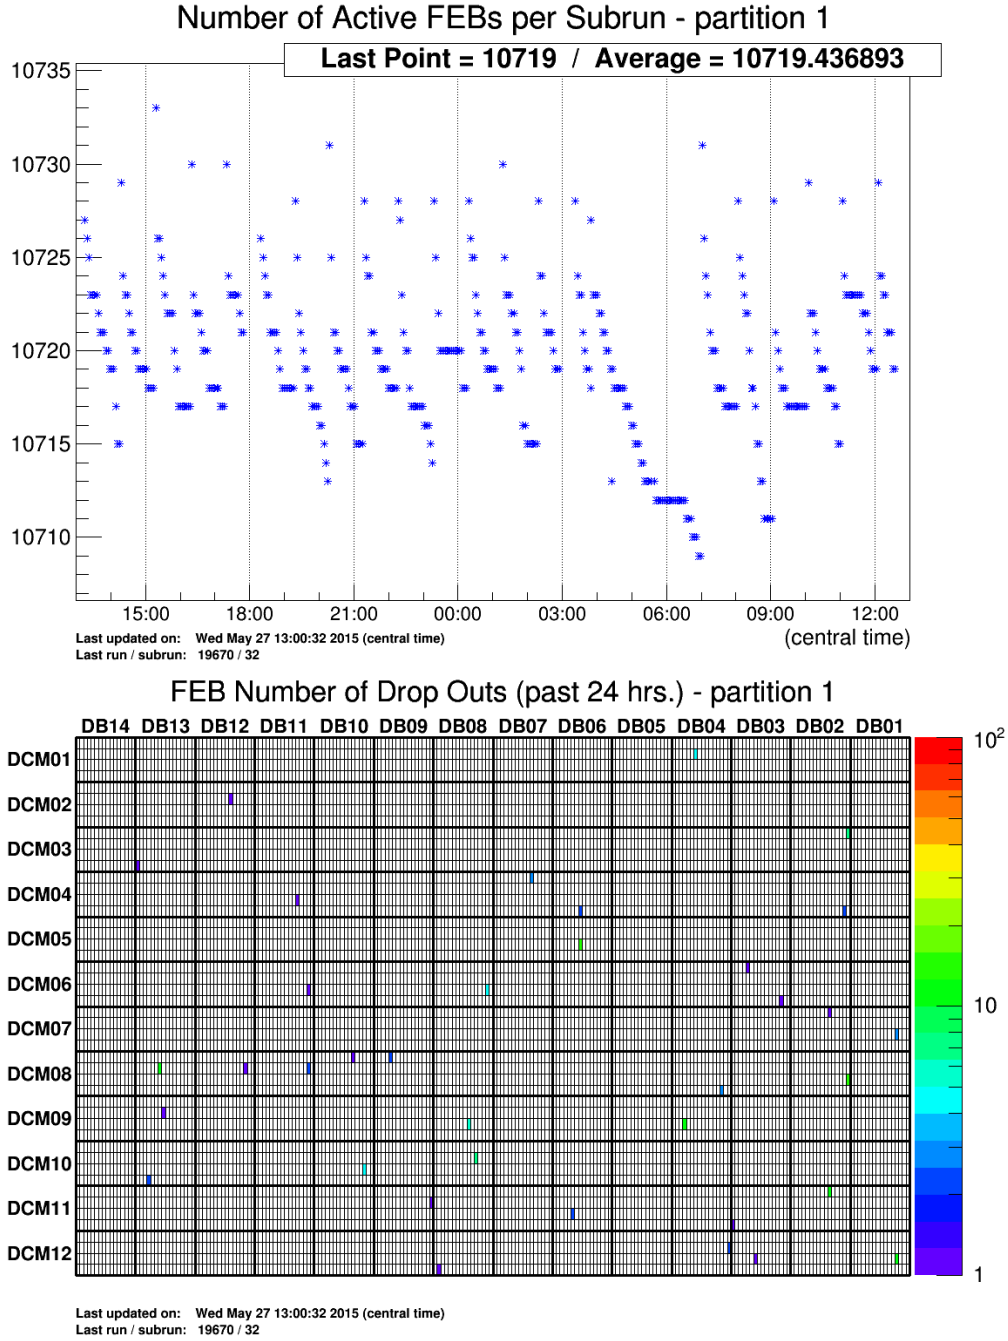
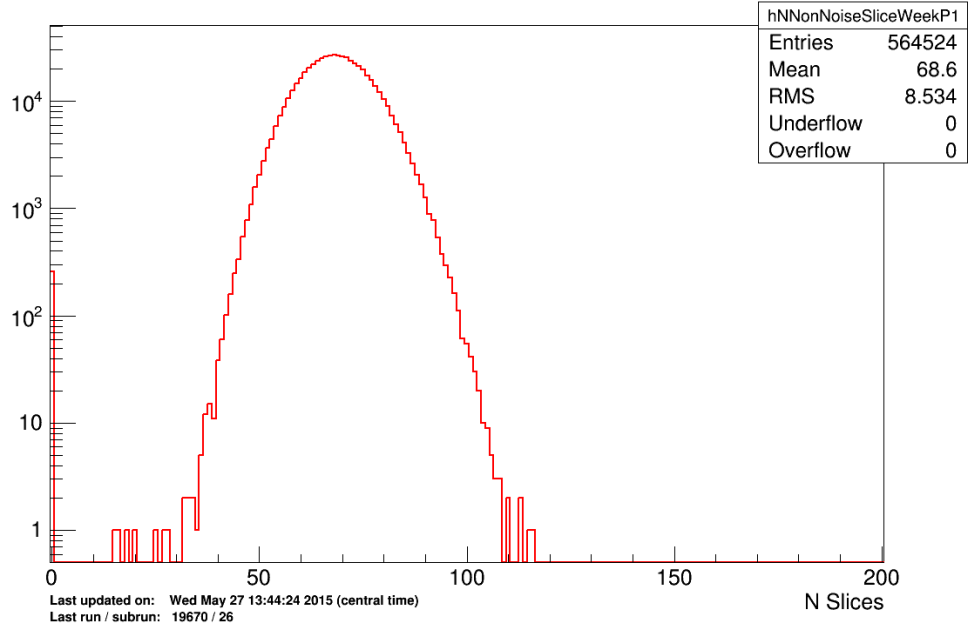


Figure 4.8: Nearline Plots - Number of active FEBs per subrun vs. time (top) and FEB dropout counts (bottom.)

Non-Noise Slices (past 7 days) - partition 1



Number of Slices per Subrun - partition 1

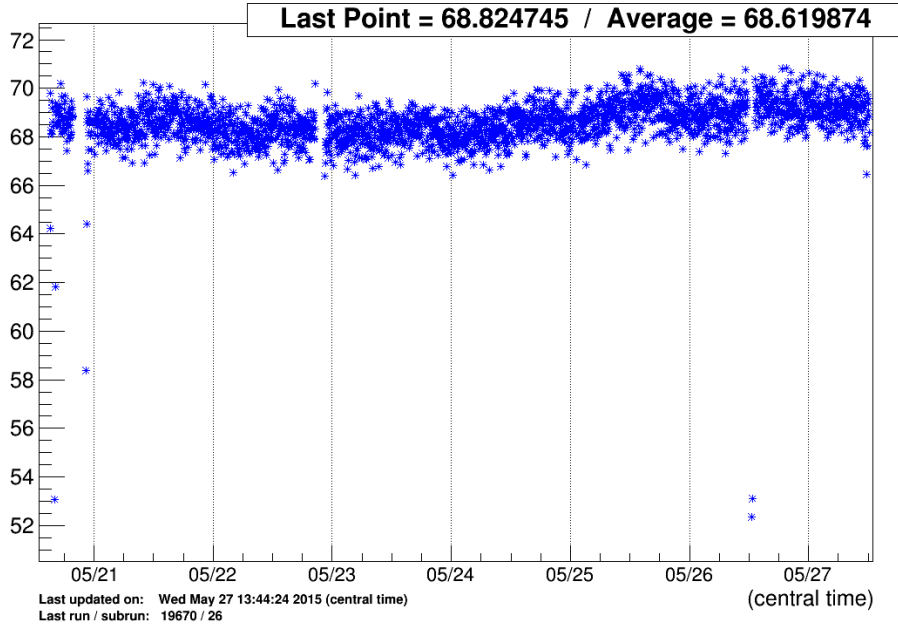
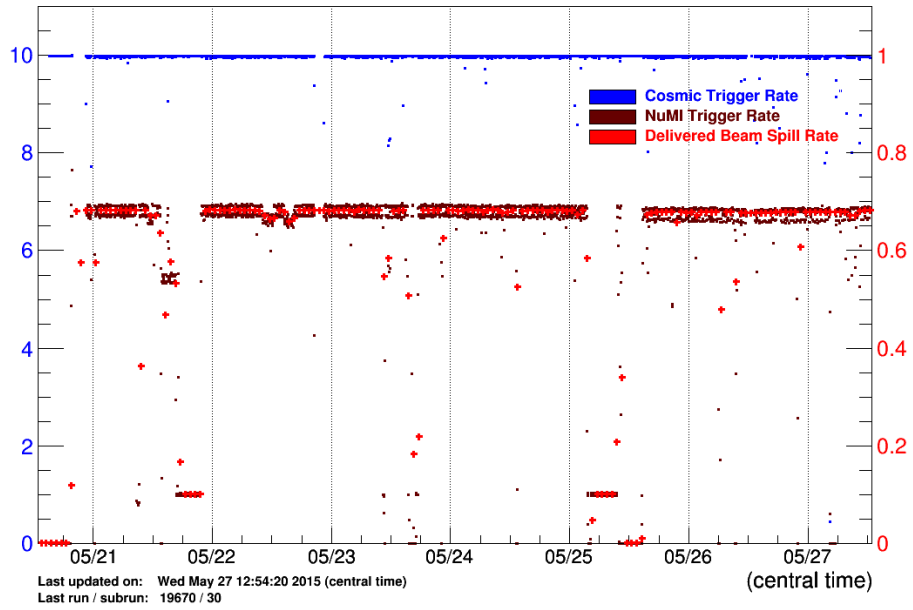


Figure 4.9: Nearline Plots - Distribution of the number of slices per spill (top) and average number of slices per spill vs. time (bottom.)

Average Trigger and Spill Rates (Hz) - partition 1



Track Fractions - partition 1

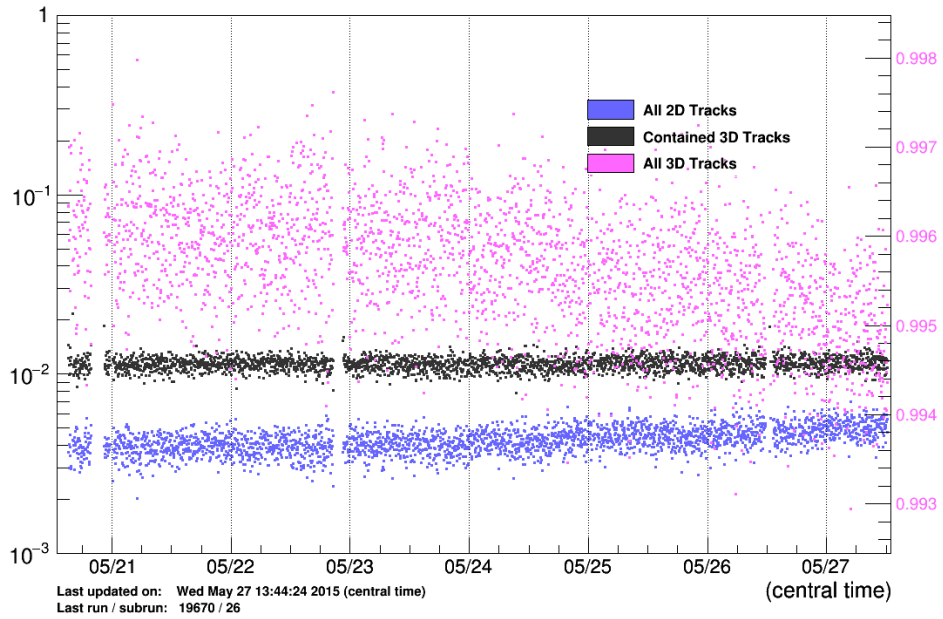


Figure 4.10: Nearline Plots - Calibration and NuMI trigger rates vs. time (top) and track fractions vs. time (bottom.)





Figure 4.11: Nearline Plot - Track start positions.





Figure 4.12: Nearline Plots - NuMI timing peak (top) and NuMI timing peak vs. time (bottom.)

PE Distribution (past 24 hrs.) - partition 1




Figure 4.13: Nearline Plot - PE distribution for slices and tracks.




Figure 4.14: Nearline examples of bad detector performance. They are (clockwise from the top left) FEB hit rate spectra vs. time, track fractions vs. time (3D tracks in red, 2D in blue, contained 3D in green), percent empty spills, and timing peak distribution (from OnMon.)

CHAPTER 5

Simulating NOvA Data

The full NOvA simulation suite includes many steps, which are run sequentially saving the output from each step to reduce the complexity and total computational time. The first step is to simulate the NuMI beam. This is followed by modeling the generation of neutrino interactions and cosmic rays within the detectors, and the propagation of any particles through the detector geometry. The last simulation step is to model the detector response to these particles ending with files that match the output format of the raw data. This chapter briefly describes these steps and includes examples of both real and simulated data.

5.1 Simulating the NuMI Beam

The simulation of the NuMI beam begins by modeling the hadrons produced when 120 GeV protons collide with the graphite target, and tracking them as they pass through the magnetic focusing horns and subsequently decay into neutrinos. All of this is handled by FLUKA [41] and Geant4 [42] through the FLUGG [43] interface. The output from the beam simulation is a flux file filled with neutrino “rays” where each ray is a neutrino of a given flavor and energy, traveling in a specific direction. Each of these rays retains information about its parentage such that downstream reweighting of the hadron production models can occur if desired.

5.2 Simulating Detector Events

Neutrino interactions within the NOvA detectors are simulated using GENIE [44], with some additional code that interfaces with GENIE to serve the specific NOvA simulation needs. Taking the flux files as input, GENIE combines information about interaction cross sections with a user supplied geometry to determine if an individual neutrino ray interacts within the detector. If it does, then GENIE will determine the relevant information about the interaction including the vertex location and the interaction kinematics. The default ν_μ CC interaction cross section used by GENIE is shown in figure 5.1. GENIE also handles propagating the primary particles created by the neutrino interaction through the nucleus, modeling inter-nuclear scattering and absorption. The output from GENIE is a list of the particles that escaped the nucleus and their kinematic variables. For cosmic rays, the generation of initial particle 4-vectors is done using the CRY [45, 46] simulation package. This package generates the flux list of particles created in cosmic ray showers well above the detector that can be handed off to the next simulation step.

Neutrino interactions are sufficiently rare in the far detector (approximately one for every 10,000 spills) to be simulated one at a time. The beam flux at the near detector however, is high enough to produce multiple neutrino interactions within the detector per spill as well as a large amount of activity from interactions that occur in the surrounding rock (“rock events.”) Accurately simulating the near detector data events therefore requires including a large volume of rock in the detector geometry model. Full simulations of NuMI beam spills (assuming a 700 kW beam) for the near detector with a sufficient amount of rock typically produce 700-800 neutrino interactions, of which only about 20 result in any energy deposited in the detector [47, 48]. Of these 20 interactions, roughly 3-4 have a vertex within the detector with the rest being rock events.

To save computational time and to dramatically reduce file sizes, we developed an alter-




Figure 5.1: The GENIE default cross section for ν_μ CC interactions on an isoscalar target, taken from [44]. The shaded band indicates an estimated uncertainty on the free nucleon cross section.

nate strategy for generating the near detector interactions. First, a large sample of single rock events was assembled by generating neutrino interactions one at a time, only keeping those that originated in the surrounding rock and deposited energy in the detector. Next, full beam spills were simulated with a geometry that included the detector but not the surrounding rock. Each of these beam spill simulations was then mixed with a Poisson-distributed random number of rock single events to emulate the expected near detector data. The mean for this Poisson distribution was determined by simulating single rock events, counting only the ones that touched the detector, and normalizing by the ratio of the expected number of protons-on-target (POT) per spill to the total accumulated number of POT for all of the rock singles generated. A cross-check of the mean computed by this method was done by simulating roughly 10,000 spills with the full rock plus detector geometry and counting only the relevant rock events per spill [48].

The next stage of simulation takes all particles created by either GENIE or CRY and

propagate them through the detector geometry model, simulating particle trajectories and energy deposition through processes such as multiple scattering, decays, and secondary particle generation. Geant4 [42] is used again for this step with the output being a list of “Fiber in Liquid Scintillator Hits” (FLSHits) which represent true energy deposited in the detector.

5.3 Simulating the NOvA Readout Response

The response of the NOvA detectors is modeled with two software modules specific to the NOvA experiment [49,50]. The first converts FLSHits into scintillation light and determines how many of those photons are collected by the fibers in the cells. The fiber collection efficiency is a function of the position along the cell and was determined from a ray tracing simulations. The collected photons are split, with half traveling in each direction along the fiber, and attenuated according to measured results from fiber quality control tests, in order to compute the number of photons that are expected to arrive at the APD. The true number of photons captured by the APD is calculated through Poisson sampling, using the expected number as the Poisson mean. The final simulated APD signal combines the signals from the photons arriving at the APD with a model of the known APD noise response.

The second module simulates the readout response of the NOvA FEBs to the APD signals. As described in section 3.3.2 the FEBs perform in real time the pulse-shaping and pedestal subtraction, resulting in a digital output signal. This module simulates these processes while introducing electronic noise. For cells with FLSHits, Gaussian-Markov processes are used to add noise to the shaped pulse from both current and voltage sources. Noise in cells without FLSHits is modeled by drawing from a distribution of unclustered hits taken from real data. Shown in figure 5.2 is a comparison of the ADC distribution for cosmic-ray hits between simulations and far detector data.




Figure 5.2: A comparison of the ADC distribution for hits on cosmic tracks that cross more than 20 planes between simulations made with CRY and far detector data. Cosmic tracks are produced by the tracking algorithm described in section 4.2.1.

5.4 Fine Tuning the Simulations With Data

There are often unanticipated detector response effects not modeled by the simulation. Through a comparison to the near detector data, two such effects have been discovered that have been incorporated into the simulation. Those effects are the quenching of scintillation light modeled by Birk’s Law, and a reduction in the overall APD output caused by large energy deposits in a single channel referred to as the “APD sag” effect.

5.4.1 Tuning the Birks-Chou Suppression Parameters

In a scintillating material, the observed light yield is proportional to the energy deposition rate dE/dx . However at high energies, the light yield begins to quench due to not enough charge being present in the surrounding medium. This is well modeled by an empirically derived relationship known as Birks-Chou attenuation [51] in which the light yield is taken to be

$$LY \propto \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2}, \quad (5.1)$$

where k_B and k_C are material dependent parameters. The original NOvA simulation model used only the Birks parameter k_B with a number that did not appear to match measured values of dE/dx for protons in data [52]. A study was done that showed that including the second order Chou parameter k_C allowed the simulation to be tuned so that it matched the near detector data. This tuning produced values of $k_B = 0.04$ cm/MeV and $k_C = -0.0005$ cm²/MeV² which matched the data well [53,54].

5.4.2 Modeling the APD Sag Effect

Through testing the APD response with a light-injection system, it was discovered that a large amplitude pulse on one APD channel caused the baseline for all other channels on the APD to drop, resulting in a lower pulse height for any signals seen in those channels [55,56]. Figure 5.3 shows an example of this ‘‘APD sag’’ behavior, which is likely a simple effect caused by the limitations of the electronics. Bench measurements of the APD response in




Figure 5.3: Results showing the readout for all 32 channels on a single APD when light was injected into just channel 11. Channel response was measured in units of ADC as a function of time in units of 125 ns. The plot on the right zooms in on the y -axis showing the sag effect on the other 31 channels. This plot was taken from [56].

this situation showed that the sag effect was consistently 1.86% across different channels and for different amounts of injected light [56]. Modeling this behavior in the simulation was achieved by incorporating a function of the shape shown in figure 5.3. This model performed well when tested against the data by looking for cell hits with large energy deposits and

examining the response of the other cells on the same APD [57, 58].

5.5 Event Displays

The end result of all of the stages of the simulation described above are files containing simulated events that have been cast into the same format as the raw data, which allows the simulations and the data to be easily treated in the same way for the purposes of reconstruction and analysis. Figures 5.4 through 5.7 show real and simulated data drawn with the NOvA event display tool for both detectors. The event display is a basic ROOT [31, 32] GUI that draws the information about the event on four main “canvases.” The top two canvases display the hits in the XZ and YZ planes which are reading out the vertical and horizontal cells respectively. The bottom left canvas shows a 1D histogram of the hit times with respect to the trigger time T_0 , and the bottom right canvas shows a 1D histogram of the ADC spectrum for all hits.

Figure 5.4 shows a 550 μsec readout of far detector cosmic ray data and figure 5.5 shows a simulation generated with CRY for the same amount of time. In each event display, the hits are colored and scaled by their ADC values. Small, dark blue hits are primarily noise hits and large, red hits are primarily caused by high energy physics. Long cosmic ray tracks are clearly visible in each picture, with large energy deposits caused by stopping particles or by tracks with long path lengths through a cell. The 1D hit time histogram shows spikes of activity from each cosmic ray on top of a relatively constant background of noise hits.

Figure 5.6 shows NuMI beam data taken with the near detector and figure 5.7 show a NuMI beam simulation generated with GENIE assuming a 350 kW beam to match the average beam power at the time the data was taken (the beam enters from the left.) Each of these event displays shows the hits scaled by their ADC values but colored by time and is zoomed in on the 10 μsec NuMI beam spill window. Hits that occurred outside of this

window are drawn in gray. The large gray hits in figure 5.6 are caused by particularly hot pixels or FEBs that will be masked off prior to reconstruction. Both the data and the simulation clearly show neutrino interactions inside the detector volume with activity from rock events, mostly in the form of through-going muon tracks. The apparent sparsity of tracks at the far right side is due to the larger gaps between the near detector planes in the muon catcher.




Figure 5.4: An event display showing cosmic ray data taken with the far detector. Hits are colored and scaled by charge.




Figure 5.5: An event display showing simulated cosmic ray data with CRY for the far detector. Hits are colored and scaled by charge.




Figure 5.6: An event display showing NuMI beam data taken with the near detector. Hits are scaled by charge and colored by time.

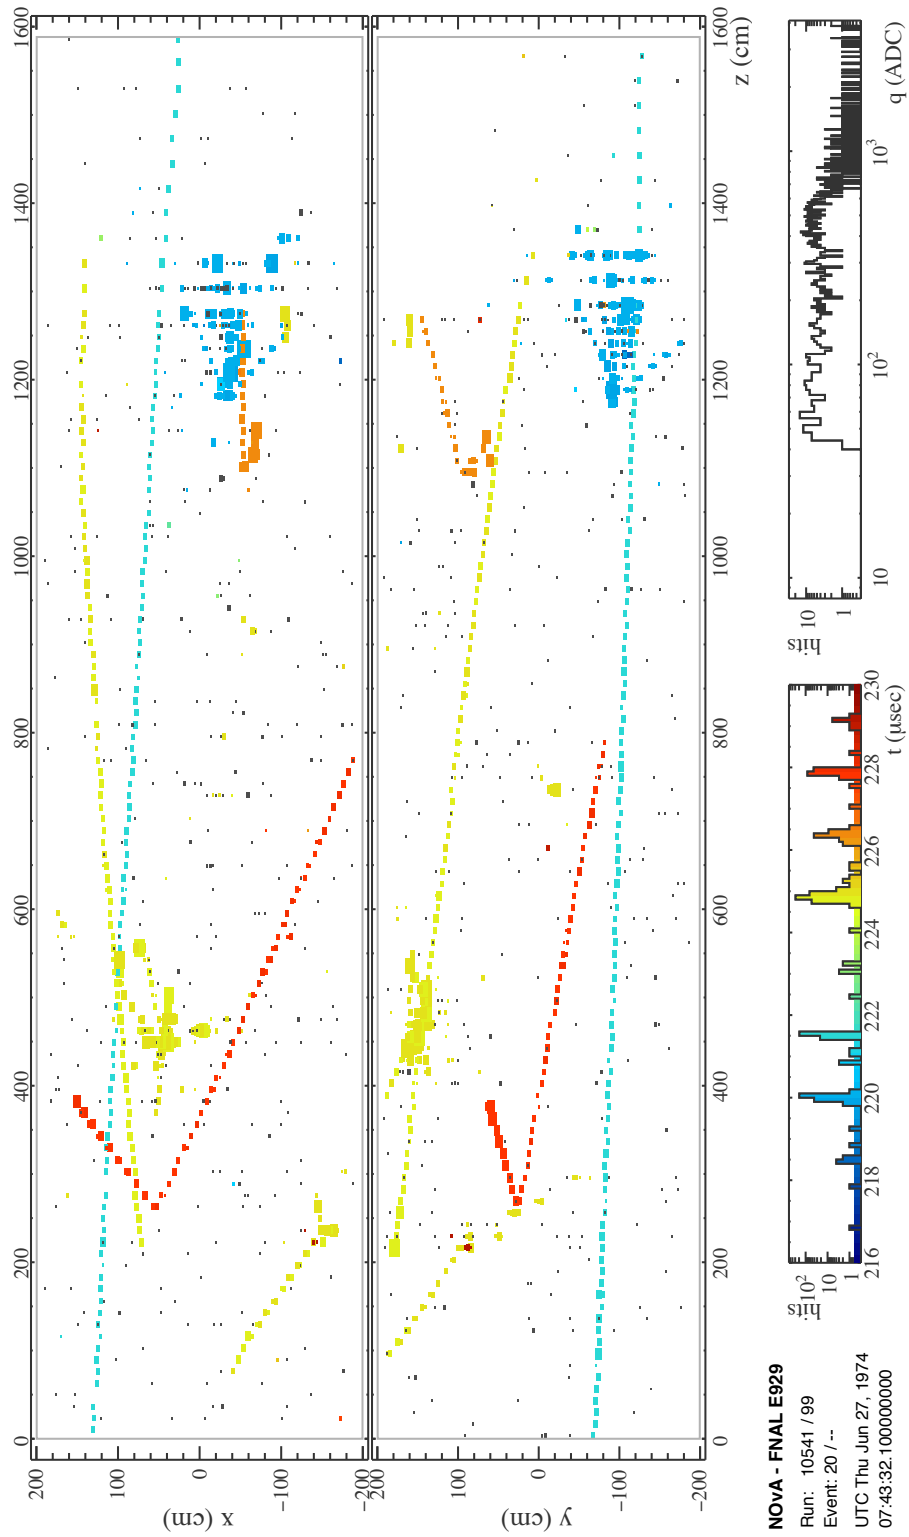


Figure 5.7: An event display showing simulated NuMI beam data for the near detector. Hits are scaled by charge and colored by time.

CHAPTER 6

Calibration and Event Reconstruction

This chapter describes in detail all of the steps necessary to reconstruct a NOvA neutrino event. It describes the calibration processes for the individual hits, how those hits are clustered, how major event features are identified, and how particle tracks are reconstructed from subgroups of hits within each cluster. Particle identification and event energy estimation are discussed in chapter 7.

6.1 Calibration

There are really two main parts to the calibration of the cell hits in the NOvA detectors; the energy calibration and the timing calibration. The purpose of the energy calibration is (obviously) to put all cells in the detector on an equal footing such that light measured by any one of them can be properly converted into an amount of energy deposited in the detector. This occurs in three parts. First, the light attenuation properties for each individual fiber in the cells must be measured which accounts for cell to cell variations in fibers and scintillator quality. Second, an absolute energy scale must be derived to convert attenuation-corrected signals into GeV. Third, the response of each channel must be monitored for changes over time due to things like seasonal temperature variations and scintillator degradation. The purpose of the timing calibration is to compute offsets between each DCM's clock within the detector and to compute the timing resolution of the cell hits.

More information on the attenuation and absolute energy calibrations, the drift calibration, and the timing calibration can be found in [59,60], [61], and [30] respectively.

The units of cell hit readout described in this section are as follows:

- **ADC:** The peak value readout by the electronics in the raw data.
- **PE:** A simple linear conversion of the ADC value into a unit intended to approximate the number of photo-electrons seen in the readout.
- **PECorr:** The number of photo-electrons seen corrected for attenuation.
- **GeV:** PECorr converted into absolute energy units.

6.1.1 Attenuation Calibration

The attenuation calibration is done separately for each cell in both of the NOvA detectors using cosmic-ray muons. At a rate of approximately 100,000 Hz, these muons are available in great excess at the far detector and the DAQ need only take any random 550 μsec snapshot to capture of order 50 of them. For the near detector, the rate is only 5 Hz but cosmic-ray muons can still be captured using an appropriate data-driven trigger.

A sample of through-going muon tracks are selected for the attenuation calibration. Muons are chosen because the rate at which they lose energy (dE/dx) is well known and relatively flat for energies between 1 and 100 GeV. Picking through-going tracks helps avoid the end of the track where the dE/dx curve rises rapidly. This selects a sample of muons that are depositing a relatively consistent amount of energy in the detector. As a quality cut, only the cell hits along these tracks for which the muon path length through the cell is considered reliable are used. These cells are chosen based on whether or not the two adjacent cells within the same plane were also hit. After taking a substantial amount of data, a data set for each cell is accumulated of ADC/cm as a function of where the hit occurred along the length of the cell as determined through three-dimensional track reconstruction. An example of this kind of data set is shown in figure 6.1.




Figure 6.1: Left: Attenuation fit for a specific cell with cosmic-ray data taken using the near detector. The ends of the cell are indicated by the vertical dashed lines. Data outside of these regions comes from various track reconstruction failures and is not used in the fit. The attenuation fit is shown in blue. **Right:** Example plot of PECorr/cm vs. distance from the track end, generated from far detector cosmic-ray data used to compute the absolute energy scale.

The mathematical form of the fit shown in the left plot in figure 6.1 is given by

$$y = C + A \left(e^{x/X} + e^{-(3L/2+x)/X} \right), \quad (6.1)$$

where X is the attenuation length, L is the length of the cell, and x is the distance measured from the center of the cell. The $3L/2$ in the second exponential term accounts for the fact that any light absorbed by the fibers will propagate in both directions to the readout. Benchmark measurements of the fibers showed that they have both a long and a short attenuation length. Using the constant term C was found to account for the longer attenuation length sufficiently well in the data.

It was also seen in the data that the attenuation curves, like the one shown in figure 6.1, have a “rolloff” feature near either end of the cells which can be explained by the geometric properties of reflected light in different locations within the cell. A good fit to this feature

was found empirically with the following function

$$y' = \begin{cases} 1 - \alpha_R(x - x_R) & x > +x_R \\ 1 & otherwise \\ 1 - \alpha_L(x - x_L) & x < -x_L \end{cases} \quad (6.2)$$

with the full attenuation fit being the product of equations 6.1 and 6.2. Lastly, any remaining small deviations from this fit model are accounted for by applying a locally weighted scatterplot smoothing (LOWESS) method described in [60]. This method uses 20 points spread out equally across the cell to fit the residuals. The seven free parameters in equations 6.1 and 6.2 along with the 20 LOWESS points for each cell are then stored in a database. This allows any cell hit measured in PE to be converted into PECorr provided an estimate of the hit location within the cell during offline processing.

6.1.2 Absolute Energy Calibration

Once the attenuation calibration is complete, then all light detected by any cell can be treated equally. The next step then is to determine the absolute energy scale that will allow measurements of PECorr to be converted into GeV. This process is not done for the individual cells, but for the detector as a whole using a sample of cosmic-ray muons that stop within the detector. These muons are selected by requiring that the track end point (taken as the lowest point) be within the fiducial volume of the detector.

Cells along this track are selected using the same good path length requirement described above. This selection allows for the construction of a plot of PECorr/cm vs. distance from the end of the track, with this distance being measured from the track end point to the center of the cell. These values are compared to a set of expected values calculated by integrating up the Bethe-Bloch curve walking back along the track from the end so as to properly account for energy loss in the different detector materials. A fit scales the measured

PECorr/cm values to these expected values in GeV/cm. To ensure consistent light levels, only the cells between 1 and 2 meters from the track end are used in generating this fit [62]. An example of this distribution generated with far detector data is shown in figure 6.1.

6.1.3 Drift Calibration

The drift calibration provides a correction for any effects that may alter the light levels seen in the detector over the duration of the experiment. Such effects could include things like degradation of the scintillator and the aging of other detector materials, or changes in the response of the electronics. Corrections are computed at the level of individual FEBs using cell hits from through-going cosmic ray muon tracks selected by the same criteria used for selecting the cell hits used in the attenuation calibration.

The cell hits selected for the drift calibration are organized by FEB. For all of the cell hits associated with each FEB, the mean PECorr/cm is computed, truncating the upper 10% of the values to avoid effects of long tails on the mean. This mean is then tracked on a run-by-run basis, where a “run” represents roughly one hour of data for the far detector and roughly three days of data for the near detector. A piece-wise linear fit is applied to this data to correct the response of each FEB over time. An example plot from far detector data and the fit correction is shown in figure 6.2.

6.1.4 Timing Calibration

The timing calibration has two main purposes. The first is to compute the timing offsets between the DCMs. This is done by examining the hit times from cosmic ray tracks that cross through multiple DCMs and is double checked by comparing the computed numbers to pings issued to the DCMs by the timing distribution units. In short, the timing offsets are accurately computed and the process is described in detail in [30]. Of more relevance to the reconstruction algorithms described below is the determination of a function that




Figure 6.2: An example of the drift calibration as applied to far detector data for Diblock 3, DCM 6, FEB 24. The mean channel response and error on the mean is shown in black with the fit for the drift correction shown in red.

predicts a cell hit’s timing resolution based on the measured PE for that hit.

As described in section 3.3.2, the FEB shapes each signal pulse from the APDs with a 460 (140) ns rise time and a 7000 (4500) ns fall time for the far (near) detector readout. The FEBs are multiplexed to readout each channel once every 500 ns in the case of the far detector and every 125 ns in the case of the near detector. Simply taking the time of the first triggered sample as the the cell hit time (known as “single-point readout”) the expected timing resolution for each hit will be $500/\sqrt{12} \approx 144$ ns for the far detector and $125/\sqrt{12} \approx 36$ ns for the near detector. However by reading out four consecutive samples (known as “multi-point readout”) an offline fit for the light arrival time can be applied to the shaped pulse to get a much better estimate of the true hit time.

The timing resolution for the hits is determined empirically from the data using through-going cosmic ray muon tracks in the case of the far detector, and through-going “rock” muon tracks in the case of the near detector. Once the offline timing shape fit has been applied, the times for the hits on these tracks have additional corrections applied for the time of flight of the particle along the track and the time for the light to propagate through the fiber

to the APD readout. After these corrections, the time difference between a given hit and all other hits (within 25 PE) in that DCM is computed and plotted in a two-dimensional histogram verses the charge Q measured in PE for the hit. For each Q bin, the RMS is calculated and taken as the timing resolution for that Q . These RMS values are then plotted as a function of Q and fit with a function of the form

$$\sigma_t = \frac{p_0}{p_1 + Q^{p_2}} + p_3, \quad (6.3)$$

where the four p_i values are fit parameters. The results for these timing fits are shown for both detectors in figure 6.3. Note that for single-point readout, the timing resolution flattens out to the expected value for high Q hits. For multi-point readout, the timing resolution for higher Q hits is around 10 ns for the far detector and around 5 ns for the near detector.




Figure 6.3: Timing resolution plots determined from far detector (top two plots) and near detector (bottom two plots) data. The two plots on the left show the single-point readout mode and the two plots on the right show the multi-point readout mode.

6.2 NOvA Reconstruction Challenges

As described in the previous chapters, the fundamental unit of NOvA data is the cell hit. Since each cell spans the full width of the detector, each cell hit represents information from 2 spatial coordinates; X,Z or Y,Z. To achieve full three-dimensional event reconstruction, information from each view must be combined, which can only be done if hits that originated from the same particle are appropriately grouped together across both views. The first step in NOvA event reconstruction is to sort all of the hits from a 550 μ sec readout window into clusters, identify which hits within those clusters are likely to have come from the same particles, and apply an algorithm to create the individual particle tracks. Once proper three-dimensional tracks have been reconstructed, this information can be used to identify particle types and estimate the total event energy.

The reconstruction philosophy described in this chapter centers around first identifying a global event vertex for each cluster of hits, so that tracks (specifically short tracks) emanating from this point are easier to find. Identifying both long tracks (such as those created by a primary lepton) and short tracks (such as those created by a recoiling proton) is important to achieve the best possible reconstructed event energy resolution. I have outlined below a summary of the major event reconstruction steps, with each of these steps occupying a section of this chapter. Naturally, a heavy emphasis has been placed on describing the steps that I contributed to (indicated in **bold**.)

1. **Grouping cell hits into clusters that are spatially and temporally separated.**
2. **Identifying simple but major, linear event features.**
3. Choosing a global vertex based on these linear event features.
4. Making sub-clusters of hits that represent a single track or shower.
5. **Creating 3D tracks from these sub-clusters.**

6.3 The Slicer4D Clustering Algorithm

For the NOvA experiment, a “hit” is the information from a single channel that was triggered by being above threshold. As outlined in section 6.1, each hit contains the information from four consecutive samplings of that channel so that an offline timing fit can be applied to achieve an improved timing resolution. Each hit contains five additional pieces of information: the plane and cell numbers for that channel, the total ADC and improved hit time (determined from the offline timing fit) and the estimated timing resolution for each hit.

In the NOvA detector, a data event is a collection of hits that occurred within a specified time interval (typically 550 $\mu\text{sec.}$) The hits in these data events can be divided into two categories, signal (originating from high energy particles traveling through the detector) and noise. Within a given data event, the signal hits can come from multiple independent physics events. For example, in the NOvA near detector, we anticipate hits from multiple neutrino interactions and in the NOvA far detector, we anticipate hits from neutrino interactions and hits from cosmic rays together in the same data event. This initial process of separating hits into groups is referred to within the NOvA experiment as “slicing.”

The goal of any slicing module is to accurately separate the signal hits from the noise and to further separate the signal hits into clusters of hits that originated from the same source (ideally, one cluster for each neutrino interaction or cosmic ray.) The true clusters can often overlap in both space and time, therefore an intelligent algorithm that can make the best use of both spatial and temporal information is required. I designed the “Slicer4D” clustering algorithm to do this by using an expanding, density-based clustering algorithm initially outlined in [63].

6.3.1 The Expanding Density-Based Clustering Method

The clustering algorithm outlined in [63] is designed for scenarios in which clusters are separable from each other by regions in which the density of points in some parameter space drops below some critical density. This algorithm works well even when the clusters have irregular shapes and when the density of background noise hits is relatively constant within the parameter space. A local density is computed for each hit by counting the number of neighbors that are within a certain “distance” of that hit. A strength of this algorithm is that this “distance” can be calculated with any kind of score function that is appropriate for the specific clustering application. The score function I used for clustering events in NOvA uses information about the spatial and temporal separation between two hits to determine if they are neighbors and is described in the next subsection.

The clusters that this algorithm makes have two types of points: core points and border points. Core points have at least the minimum number of neighbors within the critical distance. Border points have less than the minimum number of neighbors, but are included in the cluster if and only if they are the neighbor of a core point. The algorithm makes clusters by expanding the cluster around the core points. It begins by looping over all points. When it finds a core point, it begins a cluster and adds that point and its neighbors to the cluster. It then expands out to all of the neighbors of that core point and asks if they are also core points. If so, then it adds their neighbors to the cluster and continues expanding. The clustering is done when every expanding branch terminates in a border point. The algorithm then returns to the original list of points and seeks out then next unassigned core point. At the end, any point not assigned to a cluster is labeled as noise.

6.3.2 The Neighbor Score Function

The neighbor score function is what determines whether or not two hits are to be considered neighbors. Most of the particles passing through the NOvA detector are all traveling at or very close to the speed of light. Therefore, hits that originated with the same individual physics event (one neutrino or one cosmic ray) are those that are separated by light-like intervals in 4D space. To take advantage of this, the neighbor score (NS) for a pair of hits includes a causality term, two terms to penalize hits that are far separated in space, and a term to penalize hits that have low PE. It is defined as follows:

$$NS = \left(\frac{|\Delta T| - |\Delta \vec{r}|/c}{T_{res}} \right)^2 + \left(\frac{\Delta Z}{D_{pen}} \right)^2 + \left(\frac{\Delta XY}{D_{pen}} \right)^2 + \left(\frac{PE_{pen}}{PE} \right)^5 \quad (6.4)$$

where T_{res} is the timing resolution of the two hits added in quadrature, D_{pen} is the distance penalty, PE_{pen} is the PE penalty, and PE is the number of photo electrons for both hits added in quadrature. These parameters are described below in 6.3.4.

For hits that occur in the same view, the neighbor score is calculated as shown above with $\Delta \vec{r}$ being a two-dimensional quantity. For hits that are in opposite views, ΔXY is zero and D_{pen} in the ΔZ term is replaced with a smaller term, OppViewPlPen (opposite view plane penalty) with $\Delta \vec{r}$ reducing to 1D. The charge spectrum of the noise in the NOvA detector falls off roughly like $Q^{-2.5}$ which is why the PE penalty term is raised to the power of 5. The timing resolution is calculated on a hit-by-hit basis based on the PE for each hit according to a fit function outlined in [30]. Two hits are considered to be neighbors if their neighbor score is less than a parameter ϵ . This parameter is adjustable, represents the overall sensitivity of the algorithm, and fulfills the role of “distance” in the causality neighborhood space.

6.3.3 Algorithm Speed Optimization

For the far detector, the number of hits (both physics hits and channel noise) that occur within a 550 μsec readout window is expected to be of order 10,000. As one might expect, computing the neighbor score function for every pair of hits is incredibly time consuming and unnecessary. In fact, most hits are way too far apart in time to have any chance of being related by a light-like space-time interval. To optimize the code, I skip this computation for hits that are too far apart in time.

First, I sort the hits by time. Then while looping over the pairs of hits, I compute their time difference ΔT and if that difference is greater than 1000 ns, I skip the remaining computations for that hit pair. Since the hits are sorted by time, the remaining hits in the list must also be more than 1000 ns apart so they are skipped as well. I chose the time interval of 1000 ns based on the assumption that all particles in a given event are traveling at or close to the speed of light. A particle traveling at the speed of light would traverse the long axis of the far detector (roughly 60 m) in 200 ns, so five times this interval is an appropriately safe choice. I tested these changes over some simulated cosmic rays for the far detector and found that the algorithm was roughly 50 times faster while maintaining identical performance [64].

6.3.4 Tuning the Parameters

All of the adjustable parameters within the slicing algorithm are listed in table 6.1. The main parameters that had the biggest impact on the performance were the two distance penalty terms D_{pen} and OppViewPIPen , and the neighbor score cutoff value ϵ .

I set three of the parameters listed in 6.1 prior to performance tuning based on prior knowledge of the detector performance, PE_{pen} , MinPts , and MinHitsPerView . Through my initial explorations of the Slicer4D algorithm, I found that changes to the PE_{pen} term had

variable	description
MinPts	The minimum number of points within a distance of ϵ to be considered a core cluster point.
ϵ	The neighbor score cutoff value for two hits to be considered neighbors.
D_{pen}	The distance (in cm) for which the distance penalty terms will be 1.
OppViewPlPen	The distance (in planes) for which the Z-distance penalty term will be 1 for hits in opposite views.
PE_{pen}	The number of photo electrons for which the PE penalty term will be equal to 1.
MinHitsPerView	The minimum number of hits in a given view in order to make a slice.

Table 6.1: Input parameters used by Slicer4D.

essentially no impact on the algorithm performance. With the far detector cells being 15 m long, significant attenuation will occur for physics hits produced at the far end of one of these cells resulting in a signal amplitude on par with the level of electronic channel noise. To eliminate the possibility of penalizing real physics hits at the far end of the NOvA cells with low photo-electron counts, the PE_{pen} term was set to a value of 0, effectively turning off that piece of the neighbor score function. Since the Slicer4D algorithm is a density based algorithm, one needs to decide the critical density that defines the slice borders. In terms of the Slicer4D parameters, this is a combination of the variables ϵ and MinPts. To determine the optimal critical density, one has the freedom to set one of those two variables and tune the other. After experimenting with a few values and closely examining the data, I set MinPts to 4 leaving ϵ to be tuned. Lastly, since MinHitsPerView represents a lower limit on the size of the slices that we considered to be of interest, it was set to 3. I made this choice based on the observation that the vast majority of slices made out of randomly coincident noise hits had less than 3 hits in each view. All of the other parameters listed in table 6.1 were first tuned to the Monte Carlo, then to the data as described below.

To test any set of these parameters, I used two primary metrics. The first is “complete-

ness” and represents the extent to which a slice includes all of the hits that belong to a single physics event (neutrino interaction or cosmic ray.) Completeness is defined as

$$\text{completeness} = \frac{E_{evt,sl}}{E_{evt,tot}}. \quad (6.5)$$

$E_{evt,sl}$ is defined as the visible energy from the event within the slice where “visible energy” is energy deposited only in the scintillator. $E_{evt,tot}$ is defined as the total visible energy deposited by the physics event in the detector. The other metric I used is “purity” which represents the extent to which a slice is contaminated by hits from other physics events. Purity is defined as

$$\text{purity} = \frac{E_{evt,sl}}{E_{tot,sl}}, \quad (6.6)$$

where $E_{tot,sl}$ is defined as the total visible energy deposited by all physics events within that slice.

My first attempt at tuning the algorithm was done with a Monte Carlo method, using simulated neutrino interactions for the near detector and simulated cosmic rays for the far detector. For this method, I defined a primary figure of merit as the number of slices for a fixed data set with completeness > 0.9 and purity > 0.9 . I varied the configurable parameters not initially set over ranges that caused the algorithm to perform poorly, then to perform well, then to perform poorly again so as to ensure that the optimal values were covered. I did a basic grid search within this parameter space to pick out the parameters that maximized the figure of merit. Examples of the distributions of the figure of merit for different values of ϵ and D_{pen} are shown in figure 6.4. This method was effective at picking out a reasonable set of numbers, however a data-driven method was developed due to the obvious fact that the data and the Monte Carlo are never quite the same. More information on the initial Monte Carlo tuning method can be found in [65].




Figure 6.4: Distributions of the FOM from the Monte Carlo method for tuning Slicer4D vs. two of the Slicer4D configurable parameters, ϵ (top) and D_{pen} (bottom). The plot colors correspond to the number of parameter set combinations with a specific value of ϵ or D_{pen} that resulted in that particular FOM. Both plots were made with near detector simulations that included rock events and were used in the initial grid-search method for tuning the Slicer4D parameters.

The fundamental premise behind the data-driven tuning method was that the “right” values for the parameters are functions of the NOvA event topologies and various aspects of the DAQ system, and we should therefore be able to determine them from the data. The exception being the neighbor score cutoff value ϵ , but this can be chosen through extensive hand-scanning once the other parameters have been set and double checked for acceptable performance with the completeness and purity metrics.

With the parameter MinPts set to 4, the distance to the third closest neighbor plays a critical role for each hit in determining whether or not that hit belongs in a slice or will be labeled as noise. To determine D_{pen} and OppViewPIPen, I ran Slicer4D with very liberal parameters to ensure that all hits that belong together were in fact within the same slice. Then for all hits in a given slice, I plotted the two-dimensional distance between each hit and the third closest neighbor for hits within the same view and the difference in plane numbers for each hit and the third closest neighbor for hits in opposite views. My assumption was that these plots will have a Gaussian distribution for hits that belong together sitting on top of a power law distribution created by the random coincidence of hits that do not belong together. Kinks in these plots should give an indication of the correct values of D_{pen} and OppViewPIPen that separate these two populations. Examples generated from near detector data are shown in figure 6.5.




Figure 6.5: Distance to the third closest neighbor for hits within the same view used for determining the Slicer4D parameter D_{pen} (left) and for determining OppViewPIPen (right). Lines are superimposed on each plot illustrating the “kink” used to choose the values for D_{pen} and OppViewPIPen.

Through this data-driven method, I determined the near and far detector values for D_{pen} to be 75 cm and 100 cm respectively. Given that the strong interaction length in scintillator is between 90 and 100 cm, these numbers are within a sensible range. The final values for OppViewPIPen were 8 and 4 for the near and far detectors. The full list of final

tuned values can be found in table 6.2.

The last number to be tuned is the neighbor score cutoff value ϵ . The original paper [63] upon which I based the Slicer4D algorithm provides a procedure for determining ϵ from a set of sample data as follows:

1. Compute the neighbor score for the third closest neighbor (called the “3-dist”) for all points in your sample and store these numbers in a list.
2. Sort the list of 3-dist values in descending order and plot those numbers as a function of position in the list (hit index).

Since the noise hits occupy regions of lower density, then the sorted list should be almost entirely noise at the top and almost entirely signal at the bottom. Plotting this list should then give some indication of an appropriate value for ϵ .

An example of a 3-dist plot and the data I used to generate it are shown in figure 6.6. It is evident from this 3-dist plot that two characteristic scales emerge, indicated by kinks in the data. The first kink occurs around 1 on the vertical axis, indicating the beginning of the boundary between signal and noise, and is taken as a minimal value of ϵ . The second kink occurring around 10^3 on the vertical axis represents the typical spacing of noise hits. The drawback with this plot is that it gives you an indication of how far apart signal hits are from other signal hits that belong together in the same cluster, without giving you any indication of how far apart separate clusters are from each other. The goal of the slicing algorithm is to minimize event pileup by putting as many hits together from the same physics events as it can while not being so greedy as to combine separate clusters into one. So we want to set ϵ to be larger than 1 but not so large that it merges groups of hits that do not belong together.

To determine the best number for ϵ , I processed a set of data events with a range of different values. I scanned each of these events and recorded any “failures” of the algorithm. “Failures” were clusters that I judged by eye to have certain undesirable characteristics such




Figure 6.6: Event display (top) for the data event used to generate the 3-dist plot shown on the bottom.

as inappropriately split or merged clusters or clusters that excluded (included) hits that clearly did (not) belong. From the compiled list of these failures, I calculated an overall efficiency as the ratio of the number of slices without mistakes to the total number of slices for each of the different values of ϵ . I took the best number for ϵ to be the central value in the region of highest efficiency. More information on this method can be found in [66].

parameter	ND value	FD value
MinPts	4	4
ϵ	5.0	2.0
D_{pen}	75.0	100.0
OppViewPIPen	8	4
PE_{pen}	0	0
MinHitsPerView	3	3

Table 6.2: Parameter values used by Slicer4D for both the NOvA near and far detectors.

6.3.5 Results and Performance

The results of slicing a single far detector data event (defined as the full readout from a 550 μ sec window of time) are shown in figure 6.7. This particular data event was the result of a NuMI beam spill trigger and was taken at a time when the far detector construction was not yet complete (thus the empty space on the right-hand side.) In the first picture in this figure, the hits are colored by time and multiple cosmic tracks on a uniform (in time and space) background of noise hits are evident. The second picture shows the results of applying the slicing algorithm. In this picture, hits that have been grouped together into the same slice are shown in the same color and hits labeled as noise are drawn in light gray.

The completeness and purity numbers defined in equations 6.5 and 6.6 for near and far detector simulations are shown in table 6.3. These simulation samples include near and far detector neutrino interactions, as well as simulated cosmic rays for the far detector. I generated the numbers in this table with the parameters listed in table 6.2 using roughly 100,000 data events for each sample. Since slices typically contain hits from multiple physics events, the best matched physics event for a given slice is defined as the event with the highest purity. The numbers listed in table 6.3 are for the best matched physics event for each slice.

It is possible to have a random group of coincident noise hits that end up in the same




Figure 6.7: Event display from a far detector data event prior to (first picture) and after (second picture) slicing has been applied.

slice together with no hits from a legitimate physics event. Naturally, the probability for this to occur drops dramatically as the size of one of these noise-only slices increases. For the completeness and purity metrics however, this leads to an excess of very small slices that have completeness and purity values of zero. To remove these slices from consideration, a simple cut requiring a slice to have at least 10 hits in each view is also applied. This is a reasonably safe cut for any analysis. For comparison, a typical 1 GeV beam-induced muon

is expected to leave more than 50 hits in each view. Applying this cut improves the numbers in table 6.3 and is a better representation of the slices that are used for physics analysis.

simulation type	number of slices	mean comp. (%)	mean purity (%)
All Slices:			
Near Det. GENIE	3.93×10^5	89.1	98.7
Far Det. GENIE	2.86×10^5	31.8	34.1
Far Det. CRY	5.80×10^6	95.4	96.5
Slices with at least 10 hits per view:			
Near Det. GENIE	2.66×10^5	94.4	98.5
Far Det. GENIE	8.53×10^4	97.6	100.0
Far Det. CRY	4.92×10^6	99.3	99.3

Table 6.3: Completeness and purity numbers (expressed as percents) for near and far detector simulated neutrino interactions (GENIE) and far detector simulated cosmic rays (CRY.) Shown are all slices and all slices passing the requirement of at least 10 hits per view.




Figure 6.8: Completeness and purity distributions for near detector GENIE events with at least 10 hits in each view.

Shown in figures 6.8, 6.9, and 6.10 are the completeness and purity distributions for the near detector GENIE, far detector GENIE, and far detector CRY events respectively. Each of these plots shows only the slices with at least 10 hits in each view. The relatively sharp cut off for purity values less than 0.5 is caused by the choice described above to pick the best-matched physics event for each slice by purity. Since the far detector GENIE events




Figure 6.9: Completeness and purity distributions for far detector GENIE events with at least 10 hits in each view.




Figure 6.10: Completeness and purity distributions for far detector CRY events with at least 10 hits in each view.

are generated in a mode that places one and only one neutrino interaction in each data event, we expect the purity numbers for the far detector GENIE sample to be 1.0.

6.4 The Multi-Hough Transform

Once clusters of hits that belong to the same physics event have been formed, the next step is to identify the event vertex. The Hough transform is a relatively standard algorithm used for pattern recognition problems. It is designed to identify major lines or features in a two-dimensional image composed of points or pixels. Since the NOvA data is composed of a two-dimensional image for each view, this makes the Hough transform an ideal algorithm as a first step to finding the global event features and the vertex. The method described

below is based off of a modified, two-point Hough transform originally outlined in [67].

6.4.1 Filling the Hough Space Parameter Map

The Hough transform works by parameterizing lines in a polar coordinate space where ρ is the perpendicular distance from the origin to the line and θ is the angle between ρ and the x -axis. The relationship between x , y , ρ , and θ is given by

$$\rho = x \cos \theta + y \sin \theta. \quad (6.7)$$

This modified Hough transform is applied to each individual detector view (XZ and YZ) by taking pairs of hits in our event, calculating the line that passes through this pair in terms of ρ and θ , and filling a parameter space map (Hough map) with a Gaussian smeared weight. The Gaussian weight is calculated according to

$$\text{weight} = e^{-\frac{(\rho-\rho_0)^2}{2\sigma_\rho^2}} e^{-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}} \quad (6.8)$$

$$\sigma_\rho = \frac{3}{\sqrt{12}}[\text{cm}] \quad (6.9)$$

$$\sigma_\theta = \frac{3[\text{cm}]}{d\sqrt{6}}, \quad (6.10)$$

where d is the distance between the two hits in the detector in cm. In this manner, major event features will cause a build up of votes in certain regions of the Hough map resulting in peaks at the location of each of the lines that characterize the event.

I placed some restriction on which pairs of hits are allowed to vote together. First, pairs of hits that are farther apart than $\sqrt{15000}$ (cm) are prevented from voting together. This distance cutoff is a tune-able parameter and it serves two purposes; it cuts down on the number of pair combinations, and it prevents too heavy of an emphasis from being placed on hits that are far apart since the uncertainty in θ is smaller for hits that are farther away

from each other. Second, since hits in the same horizontal row can have exactly the same x or y coordinate, these hits are also not allowed to vote together unless they are farther apart than $0.25\sqrt{15000}$ (cm). I found this restriction necessary to reduce an otherwise strong tendency to produce lines that were completely horizontal and not representative of the data.

Once the Hough space map is created, I apply some smoothing by averaging hits in the map using a Gaussian smoothing weight. The size of the window over which the smoothing is applied is an adjustable parameter and is currently set to ± 3 bins on either side of the bin to be smoothed in both ρ and θ , with the bins in the Hough space map being 1 cm wide in ρ and 1° wide in θ . I found that this smoothing increased the accuracy of the lines, but it reduced the overall speed of the algorithm. Since the majority of the Hough space map is empty, I only apply the smoothing calculation to the non-zero bins which increases the speed without compromising the performance.

6.4.2 Calculating the Threshold

I require the peaks in the Hough space map to be above a threshold value in order to be labeled as a valid line. Since each event can be composed of a different number of hits, I calculate this threshold value separately for each Hough map. To calculate the threshold, I loop over all bins in the Hough map and calculate both the average and standard deviation, σ , of the bin heights. I take the threshold to be the average bin height plus $N\sigma$ where N is an adjustable parameter. I found through some “hand tuning” that N could be lowered to zero (allowing for small legitimate peaks in the Hough map to be found) while maintaining good results.

6.4.3 Refining the Peak Value

After the Hough map is created, filled with votes, and smoothed, I use the center of the tallest bin in the Hough map as the ρ and θ values for the line of interest. But, since the center of the bin is not necessarily the true ρ and θ value that best represents the line, I apply some refinement of the peak value by averaging over a 7×7 square of bins (centered on the peak bin) where the average is weighted by the bin value and de-weighted by the distance (in terms of number of bins) to the central bin. So that the central bin carries more weight than the four immediately adjacent bins, I set the de-weight distance for the central bin (which would normally be zero) to $\sqrt{2}/2$.

6.4.4 The Multi-Hough Loop

If all peaks above the threshold value are taken as legitimate lines, then the modified Hough transform described above has a strong tendency to produce too many lines. After trying many different values for the Hough map threshold, I discovered that there was no clear value that consistently kept the good lines and rejected the lines that did not belong. So I developed a procedure for removing the hits associated with an already identified Hough line and regenerating the Hough map to look for smaller lines. With the hits associated with the first line removed, the dominant features in the new Hough map are less representative of noise and much more representative of shorter, legitimate, physics tracks.

After a line is identified from the tallest peak in the Hough map, I loop over the list of all hits to determine which ones are associated with that line. Any hit within 6 cm of the line will be reassigned a weight of zero (each hit in the hit list has a weight assigned to it, with the default values all being set to one.) This 6 cm cutoff distance is another configurable parameter, which I chose because it represents roughly the size of one cell. After the hits are removed, I reset the weights of the most upstream and most downstream

hits (as determined by the Z coordinate and allowing for ties between at most two hits) to one, so as to not remove hits potentially associated with another line. Lastly, I “scrub” the hit list by removing hits that have no neighbors within 30 cm. I use this new list to generate the next iteration of the Hough map.

The same procedure outlined above is applied to the new Hough map, the tallest peak is taken to be a legitimate line, that peak value is refined, and the hits associated with that line are removed. In addition to these steps, I test the new line against all the other lines found so far to determine if it is identical to any of the others. I consider two lines to be identical if their ρ values are within 15 cm and their θ values are within 0.02 radians. If a line is found to be identical to another by this definition, then I remove the hits associated with it but do not add it to the list of Hough lines before proceeding to the next iteration. Finally, I repeat this “Multi-Hough” loop. To exit the loop, one of two criterion must be met. Either no more peaks are found above the threshold value, or a maximum number of lines is reached. The maximum number of lines is a configurable parameter with a default value of 10.

6.4.5 Results and Performance

Figure 6.11 depicts the results of the Multi-Hough transform from a simulated, ν_μ CC event in the NOvA near detector on the surface (NDOS). For simplicity, only one view (the YZ view) is shown as the algorithm steps through its iterative process. Shown next to each event is the corresponding Hough map. Note that by removing the hits on each good line, the noise is drastically reduced in the Hough map for the following iteration.

Given that the main goal of this algorithm is to provide input to a vertex finding algorithm, one of the primary metrics that I used to asses performance is the perpendicular distance from the Hough lines to the true Monte Carlo interaction vertex. Since the Hough




Figure 6.11: An example of the Multi-Hough algorithm. Shown are the results after four iterations through the peak finding loop. The figure on the right is the corresponding Hough space map for each event on the left.

transform works in two dimensions, I calculate this separately for each view in the detector.

In general, this can be thought of as an indication of how correct the intercept of the line

is. The validation data sample I used was composed of simulated events in our far detector

and was divided into 8688 ν_μ CC, 8802 ν_e CC, and 9343 NC events. Shown in figure 6.12 are the plots for the perpendicular distance to the first and second Hough lines found in the event for the XZ view. The results for the Hough lines in the YZ view are essentially the same.

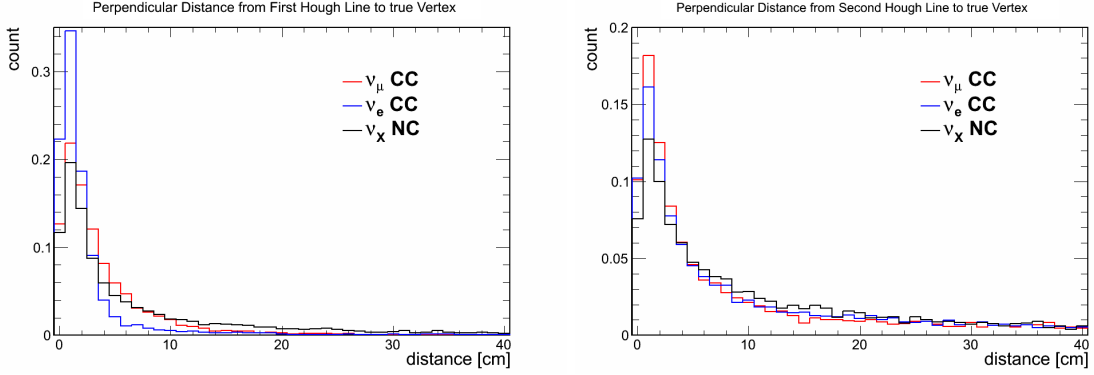


Figure 6.12: Perpendicular distance from the first (left) and second (right) Hough lines to the true vertex in the X view (area normalized to 1).

data type	number of entries	mean [cm]	RMS [cm]	overflow
ν_μ CC 1 st line	8688	4.1	5.1	113
ν_e CC 1 st line	8802	2.7	5.0	145
ν_X NC 1 st line	9343	6.7	8.5	761
ν_μ CC 2 nd line	7844	8.2	10.0	1419
ν_e CC 2 nd line	8366	8.8	10.1	1493
ν_X NC 2 nd line	7100	9.9	10.2	1482
ν_μ CC 3 rd line	5758	8.4	10.3	1002
ν_e CC 3 rd line	6533	8.8	10.1	955
ν_X NC 3 rd line	4052	9.8	10.1	802

Table 6.4: Perpendicular distance from the Hough lines to the true vertex in the X-view.

The second validation metric was a matching of Hough lines to simulated particle trajectories using the dot product between the two. The criteria I used for a particle trajectory to be matched to a Hough line were as follows:

1. The particle could not have been a neutrino, neutron, or π_0 .
2. The Hough line must have passed within 10 cm of the origin of the particle.
3. The particle track must have been longer than 20 cm.
4. The particle track can not have been matched to any other Hough line.

I considered a Hough line and particle trajectory with the greatest normalized dot product to be a match if the above criteria was met. Shown in figure 6.13 are the results of the dot product matching for the first and second Hough lines in the XZ view (again, the YZ results are similar.) This metric can be thought of as how correct the slope of the Hough line is.




Figure 6.13: Dot Product with the first (left) and second (right) Hough lines and the best matched MC particle trajectory in the X view (area normalized to 1.)

6.5 The Elastic Arms Method

The Elastic Arms algorithm is designed to use the output from the Multi-Hough algorithm as a seed for finding the global event vertex. It is based on a method of the same name listed in [68] which is sometimes referred to in the literature as the “method of deformable templates.” This method is described briefly here. For more information on the method of Elastic Arms, see [69], [68], [70], and [71].

data type	number of entries	mean	RMS	underflow
ν_μ CC 1 st line	8688	0.991	0.058	158
ν_e CC 1 st line	8802	0.996	0.039	92
ν_X NC 1 st line	9343	0.980	0.077	538
ν_μ CC 2 nd line	7844	0.948	0.137	401
ν_e CC 2 nd line	8366	0.952	0.139	1741
ν_X NC 2 nd line	7100	0.954	0.125	1002
ν_μ CC 3 rd line	5758	0.933	0.156	564
ν_e CC 3 rd line	6533	0.920	0.176	1912
ν_X NC 3 rd line	4052	0.935	0.150	823

Table 6.5: Dot product between the Hough lines and the best matched MC particle trajectory in the X-view. The numbers in the underflow bin represent unmatched lines.

6.5.1 The Algorithm

The basic template for a NOvA event is a vertex with one or more particle tracks emanating outwards from that vertex. In the Elastic Arms method, each particle track is approximated by an “arm” (a vector pointing away from the vertex) whose direction can be adjusted to match the event topology for an identified vertex seed. For the application of this method to NOvA data, the number of arms is taken to be the largest number of Hough lines found for the event in either the XZ or YZ views. To determine the location of the vertex, a list of vertex candidates must be generated and evaluated.

From this list of vertex candidates, the best vertex is chosen as the one that minimizes an energy cost function given by

$$E = \sum_{i=1}^N \sum_{a=1}^M V_{ia} M_{ia} + \lambda \sum_{i=1}^N (\sum_{a=1}^M V_{ia} - 1)^2 + \frac{2}{\lambda_v} \sum_{a=1}^M D_a. \quad (6.11)$$

Here M_{ia} is a distance measure from detector hit i to arm a , V_{ia} is the strength of the association between hit i and arm a , and D_a is the distance from the vertex to the first hit

on arm a . The parameters λ and λ_v control the strength of the second and third terms. The last term in this equation is unique to the NOvA application of the Elastic Arms method and penalizes arms whose first hits are far from the vertex. This term was added to prevent the best vertex from being the one that was always the farthest away from all of the hits, which in the extreme will always minimize the energy cost function. The hit/arm association term V_{ia} is given by

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta\lambda} + \sum_{b=1}^M e^{-\beta M_{ib}}}, \quad (6.12)$$

where $e^{-\beta\lambda}$ represents the likelihood that the hit is not associated with any arm and β can be interpreted as a range over which hits are allowed to be associated with arm a .

The list of vertex candidates is generated from information about the Hough lines and the spatial distribution of the hits in the event. From the Hough lines, vertex candidates are formed from the intersection points of the major lines in each view. The hits are sorted by their Z coordinates and additional vertex candidates are formed from selected hits at fixed intervals (2%, 5%, ... 50%) in this list. Note that this introduces a bias to favor vertices at lower values of Z but that this is exactly what one expects for beam neutrino events in the NOvA detectors.

The next step is to set the directions of the arms for each vertex candidate. To determine the arm directions, a list of possible vectors is generated from the directions of the Hough lines (matched between views by their peak heights in the Hough map) plus a minimum bias sample of vectors formed from the vertices of a dodecahedron. The arms are then set one-by-one by choosing the direction from this list that minimizes equation 6.11 before moving on to the next arm. Care is taken to ensure that the same or very similar arms are not reused for each vertex.

With a list of vertex candidates, each with a set of carefully chosen arms, equation

6.11 can now be evaluated to determine which vertex candidate will be deemed the “best one.” Once this vertex has been chosen, a process of simulated annealing is applied to allow the vertex to settle into an optimal location. The annealing is accomplished by varying the parameter β in equation 6.12 from low values (representing high temperatures) to high values (representing low temperatures.) This process allows the vertex to smoothly seek out the global minimum of equation 6.11 while avoiding potential local minima within that function.

6.5.2 Results and Performance

Shown in figure 6.14 is an event display from a far detector data event with a selected neutrino candidate. The Hough lines are drawn in red and the selected Elastic Arms vertex is drawn as a blue “X”. The Hough lines clearly pick out the major event features and align with the major particle tracks, and the Elastic Arms vertex sits at or very close to what appears to be the global event vertex.




Figure 6.14: Reconstructed far detector neutrino candidate event showing the Hough lines and the Elastic Arms vertex.

Together, the Multi-Hough and Elastic Arms algorithms achieve average event vertex

resolutions of 11.6, 10.9, and 28.8 cm for ν_μ CC, ν_e CC, and NC events respectively. For both the ν_μ CC and ν_e CC events, 68% of the vertices are within 10 cm of the true vertex (38 cm for the NC events.) This puts the vertex for the CC events within approximately 2 cell widths of the true vertex the majority of the time. The distribution of vertex resolutions broken into these three categories is shown in figure 6.15.




Figure 6.15: 3D Event vertex resolution with the Elastic Arms algorithm.

6.6 Particle Clustering with Fuzzy-K

Once the global event vertex has been identified by the Elastic Arms method, the next step is to assign a cluster membership to each cell hit within the event, with each cluster representing the hits from a single particle track or shower. This is accomplished with a possibilistic fuzzy-k means algorithm described in this section. The term “possibilistic” means that the sum of each hit’s membership across all clusters is not forced to be one, which allows for isolated hits to be treated as noise. Full details on this method can be found in [72], [73], and [74]. The process of creating reconstructed particle tracks from these clusters is described in section 6.7.

6.6.1 The Algorithm

The basic idea behind this method is that when sitting at the global event vertex for either view, the cell hits within the event should appear as peaks of deposited energy in a one-dimensional angular space around that vertex. The Fuzzy-K algorithm determines how many cluster centers (peaks) are present and assigns a cluster membership to the hits in the event. The first step in this process is to compute an angle and an angular uncertainty for each hit. The angles range from $-\pi$ to π with 0 corresponding to the Z-axis of the detector, and the angular uncertainty is given by

$$\sigma = \frac{1.745}{d} + 0.000173d + 0.0204, \quad (6.13)$$

where d is the distance from the vertex to the center of the cell hit in cm. This uncertainty was determined empirically from a simulated sample of 1 and 2 GeV muons by comparing the computed angle for cell hits with the initial track direction [72]. It was designed to account for the multiple scattering that will occur for particles as they travel farther from the global vertex.

Finding the cluster centers and determining the membership for each hit is done with an iterative process that begins by assuming there is only one cluster centered on the region of highest density in the one-dimensional hit angular space. A degree of membership for each hit j in each cluster center i is computed according to

$$U_{ij} = e^{m\sqrt{cd_{ij}/\beta}}, \quad (6.14)$$

where d_{ij} is the distance to cluster centers given by

$$d_{ij} = \left(\frac{\theta_j - \theta_i}{\sigma_j} \right)^2. \quad (6.15)$$

Here σ_j is given for each hit by equation 6.13, m represents the “fuzziness” factor which allows hits to retain partial membership in multiple clusters, c is the number of cluster centers, and β can be interpreted as a normalization term that represents how spread out the hits are expected to be around a “normal” cluster center. The cluster centers are then updated according to

$$\theta'_i = \theta_i + \frac{\sum_{j=1}^n \frac{U_{ij}^m}{\sigma_j^2} (\theta_j - \theta_i)}{\sum_{j=1}^n \frac{U_{ij}^m}{\sigma_j^2}} \quad (6.16)$$

and the process is repeated until $\Delta\theta$ is less than a specified tolerance (currently set to 10^{-7} radians) for all cluster centers.

As a final step, a check is performed to make sure that multiple cluster center seeds have not converged onto the same value. Then, a new one-dimensional distribution of cell hits in angular space is computed from all hits with less than 1% membership in any cluster. This distribution is used to identify the next most dense peak which is then added to the list of cluster center seeds. The next iteration of finding cluster centers is started with this new cluster center added to the list. This iterative process is repeated until either all cells belong to at least one cluster or until a maximum number of clusters is reached.

Since this clustering process is done separately for each view, the last step is to match clusters between the views. To do this, a temporary track is formed out of every pair of two-dimensional clusters from each view. For each of these possible pairings, the cumulative cell hit energy distribution (normalized to the total energy for each cluster) is computed for each view as a function of distance along the track. Comparing this distribution between the views allows for the application of a Kupier metric which takes the sum of the largest absolute positive and negative vertical distances between the two distributions as a metric for each of these possible cluster combinations. The view matching proceeds by pairing together the clusters from each view best matched by this metric, and then continuing in this fashion until all clusters that can be matched have been matched.

6.6.2 Results and Performance

Shown in figure 6.16 is an event display from the same selected far detector data event seen in figure 6.14. The Elastic Arms vertex has been drawn as a blue “X” and the Fuzzy-K clusters are outlined in red, blue, and green.




Figure 6.16: Reconstructed far detector neutrino candidate event showing the matched, three-dimensional clusters produced by the Fuzzy-K method.

The metric used to assess the performance is completeness as defined in equation 6.5 for hits produced by the primary lepton in charged current interactions. For ν_e CC events, the average completeness is 88%, 95% for QE events and 86% for non-QE events. For ν_μ CC events, these numbers are 93%, 98%, and 92% respectively. The distribution of completeness values as a function of true, visible, lepton energy is shown in figure 6.17.

6.7 Particle Tracking with Break Point Fitter

The Break Point Fitter (BPF) tracking algorithm relies on the formalism outlined in [75]. Using the Fuzzy-K 3D prongs and Elastic Arms vertex as input, it constructs particle trajectories using a model of Coulomb multiple scattering and energy loss for each 3D




Figure 6.17: Distribution of completeness values as a function of visible energy for the clusters formed by the primary lepton in a CC event produced by the Fuzzy-K method.

prong. It is considered an “optimal” tracker which brings all available information to bear on the reconstruction of the track parameters, using cell hit locations and measurement errors, scattering information, and correlations between upstream and downstream hits. The primary goals of the BPF tracker are to accurately track particles through all 4π steradians, to incorporate energy loss through different detector materials into the multiple scattering model, and to allow for tracks to be fit under different particle hypotheses. I will discuss how these goals are achieved in the sections below.

6.7.1 The Optimal Track Fit Expression

The BPF algorithm works by assuming that particles are traveling primarily along the z axis of an xyz coordinate system. Within this coordinate system, measurements of the track location are made in orthogonal planes at fixed locations in z to yield a set of values x_i and y_i , each with an associated uncertainty σ_i . As the track propagates, it is allowed to scatter at M locations among the n measurement planes. The locations of these scattering planes are somewhat arbitrary; I will discuss how we optimized the locations of the scattering

planes for the application of BPF to NOvA in section 6.7.2.




Figure 6.18: Schematic representation of the track model used by BPF. The measurements x_i and their uncertainties σ_i are shown as dots with error bars. The track is allowed to scatter through the angles α_J at the scattering planes indicated by the dashed vertical lines.

A schematic representation of the BPF track model is shown in figure 6.18. Using this model, the location of the track trajectory along the i^{th} measurement plane ξ_i can be expressed as

$$\xi_i = a + bz_i + \sum_{J=1}^M \alpha_J (z_i - Z_J) \Theta(z_i - Z_J), \quad (6.17)$$

where a and b are the intercept and slope of the initial track direction, α_J is the scattering angle at the J^{th} scattering plane (assumed to be small), and $\Theta(z_i - Z_J)$ is the Heaviside function which is 1 for arguments > 0 and 0 otherwise. The use of the Heaviside function ensures that only the upstream scattering angles (for which $Z_J < z_i$) affect the trajectory at the i^{th} measurement plane.

The complete track trajectory expressed in equation 6.17 is specified by a , b , and the set of α_J values. To optimize these parameters, we define a χ^2 goodness of fit as

$$\chi^2 = \sum_{i=1}^n \frac{(x_i - \xi_i)^2}{\sigma_i^2} + \sum_{J=1}^M \frac{(\beta_J - \alpha_J)^2}{\sigma_J^2}, \quad (6.18)$$

where σ_i is the error on the i^{th} measurement, σ_J is the RMS scattering angle computed for scattering plane J (given below by equation 6.19) and β_J is the expected scattering angle at the J^{th} scattering plane. For an unmagnetized detector (such as the NOvA detectors) particles will scatter left as often as they scatter right, so β_J will be zero at all scattering planes. Minimizing this χ^2 by taking $\partial\chi^2/\partial q = 0$ with $q = a, b, \{\alpha_J\}$ results in a set of linear equations which can be easily solved through a matrix inversion to yield the optimal track parameters. The details of the construction of this matrix are outlined in [75].

6.7.2 Adapting the Optimal Track Fit Expression for NOvA

To achieve the first major goal of BPF, providing accurate tracking in all 4π steradians, the track fit is performed in a “track-based” reference frame in which the primary direction of the track is taken to be the z axis. Determining the track-based frame $x'y'z'$ begins by fitting a straight line to all cell hits in the Fuzzy-K 3D prong using the Theil-Sen method [76]. This method uses the median slope between pairs of points, making it more robust against outliers than a least-squares fit. The new z' axis is taken to be $z' = \text{unit}(m_x, m_y, 1)$ where m_x and m_y are the slopes of the fitted lines from each view. The next axis is picked by computing $z' \times x$ and $z' \times y$. If $z' \times x$ has the greater magnitude, then the new y' axis is taken to be $y' = \text{unit}(z' \times x)$ and vice versa for x' if $z' \times y$ is bigger. With two axes set, the third axis is simply the cross product of the first two. The origin of this new coordinate system is set using the coordinates of the vertex produced by the Elastic Arms algorithm.

The next step constructs the lists of measurements and their errors from the cell hits, and rotate those lists into the track-based frame. Since an individual cell hit only contains 2D positional information, (x, z) or (y, z) , each hit is first converted into a 3D point using a straight-line interpolation of the 6 closest hits in z from the other view to complete the third coordinate. The cell hit uncertainties are taken to be the cell width/ $\sqrt{12}$ for σ_x and

σ_y , and plane width/ $\sqrt{12}$ for σ_z . The uncertainties for the completed coordinate (x in the y -view, y in the x -view) are estimated to be the sum in quadrature of the cell width/ $\sqrt{12}$ and the interpolated slope m times σ_z . Based on the distribution of χ^2 per degree of freedom values given by equation 6.18 from fitted tracks, it was found necessary to inflate these uncertainties by 25% to shift the peak of the χ^2 distribution to 1. This was deemed reasonable since the cross-section for a cell is roughly a square with a hypotenuse that is approximately 25% longer than the width in z , and assuming a value for σ of width/ $\sqrt{12}$ ignores the fact that particles can pass through the corners of the cells. This set of 3D hits is rotated into the track-based frame and the uncertainties are propagated into σ'_x , σ'_y , and σ'_z for each hit. The hits are sorted by z' so that every x', z' and y', z' pair can be used with equations 6.17 and 6.18 to determine the track fit parameters.

Once we have rotated into the track-based frame, the next step is to determine where to place the scattering planes along the z' axis. A simple solution would be to choose to put one scattering plane in between each measurement plane. However given the 41 cm radiation length of the NOvA detectors, the typical Coulomb scattering from one cell to the next is much smaller than the cell size itself. Including too many scattering planes is therefore computationally unnecessary. The number of scattering planes should be chosen to be as small as possible without impacting the accuracy of the tracker results.

The decisions about where to put the scattering planes are made by the BPF algorithm using a Coulomb multiple-scattering model. It begins by approximating the track trajectory with a straight line drawn from the first hit to the last hit along the z' axis. By walking backwards along this trajectory, the algorithm tabulates the distance traveled in units of radiation lengths X_0 and the total energy loss dE/dx , accounting for the different detector materials traversed. The end result is a list, as a function of z' , of accumulated radiation lengths and particle energy estimates, made by assuming that the particle starts with zero

kinetic energy at the track end.

The algorithms that compute the path-length dE/dx integration through different detector materials are handled by a special BPF class called `Path` that uses a set of dE/dx tables taken from [77]. Given a track trajectory, this class will perform the dE/dx integration as described above, to compute the energy, momentum, and the relativistic variables $\beta\gamma$ for any point along the trajectory. Computing these variables of course relies on knowing the mass of the particle that made the track. Because of this, and because this knowledge of particle mass appears again in equation 6.19 discussed below, the BPF algorithm requires an assumption to be made about the particle type in order to perform the fit for the track parameters. In the current implementation of BPF, each Fuzzy-K 3D prong is fit three times under the assumption that the track is a muon, a pion, and a proton. A decision can be made later as to which of these three track fit assumptions should be used. Using the `Path` class achieves the second goal of the BPF tracker; incorporating energy loss through different materials into the scattering model, and fitting the tracks under different particle hypotheses satisfies the third BPF goal.

Once the list of radiation lengths and energy estimates for the z' coordinates along the approximate track trajectory have been accumulated, the multiple scattering model is used to place the scattering planes. A particle with charge ± 1 that has traveled a distance s , measured in units of the radiation length X_0 , is expected to deviate from its original trajectory by an angle ψ given by

$$\psi = \frac{13.6 \text{ MeV}}{\sqrt{3}\beta p} \sqrt{s}(1 + 0.038 \ln s), \quad (6.19)$$

where β is the particle velocity in units of c , and p is the particle momentum. ψ in this equation is the angle from a 3D track projected onto a 2D plane, to be in line with the expressions derived in section 6.7.1. This equation is derived in [8] using a Gaussian approx-

imation to Coulomb scattering. The scattering surfaces are placed by walking backwards a distance d along the approximated trajectory from either the track end or the most recently placed scattering plane, using steps taken between two measurement planes, until one of two criteria are met:

1. the predicted transverse scattering distance, $d \cdot \psi$, exceeds 3 mm
2. d becomes greater than 100 cm or greater than $2X_0$

at which point the algorithm backs up one step, inserts a scattering plane, and sets σ_J for that scattering plane to the value of ψ at that location. The first criteria ensures that a scattering plane is placed before the particle is expected to scatter by an amount roughly equivalent to the thickness of the cell walls (picked as the smallest resolvable distance relevant to the tracking algorithm.) This criteria also results in more scattering planes being placed in the regions where the particle is expected scatter the most, for example near the end of the track. The second criteria enforces a maximum allowed distance between adjacent scattering planes, so that a track can't propagate too far without being given the opportunity to scatter. In addition to these criteria, the BPF algorithm dictates that at least one scattering plane is placed along the track and that no more than one scattering plane is allowed between two measurement planes. Figure 6.19 shows an example of the scattering plane placement along a 2 GeV/ c muon track.

With the scattering planes and σ_J values set, the BPF algorithm can invert the matrix generated from equations 6.17 and 6.18 to fit for the track trajectory. Since these equations are derived for a track contained within a 2D plane, a separate matrix is generated for each detector view resulting in two 2D track fits which are simply projections of the 3D trajectory into the $x'z'$ and $y'z'$ planes. After the track fit is performed, the BPF algorithm loops over all of the cell hits in the Fuzzy-K prong and computes a χ^2 for each hit based on the first sum in equation 6.18. Any hit with a value greater than 25 for this χ^2 is considered to be an outlier that could possibly throw off the track fit. These outliers are removed from

the list of cell hits and the entire track fitting procedure is repeated using this scrubbed list. The final track fit is rotated back into the detector frame.

As discussed above, the BPF algorithm performs the track fit three times under the assumption that the track was produced by a muon, a pion, and a proton. For each of these particle assumptions, the track produced includes the trajectory points from equation 6.17, the energy-momentum 4-vector for the particle computed by the `Path` class using the full trajectory, the total χ^2 per equation 6.18, and the number of degrees of freedom (NDOF) for the total χ^2 . Since equation 6.18 is a sum over the n hits and the M scattering planes, and the fit trajectory given by equation 6.17 includes $M + 4$ free parameters (the M scattering angles plus a value for a and b in each view) the NDOF will be $n + M - (M + 4) = n - 4$, where n is the number of hits used for the final track fit after the “outliers” had been removed. All of this information is stored with each track so that it can be used by any analysis methods performed later.

6.7.3 Results and Performance




Figure 6.19: This plot shows the location of the scattering planes for the muon fit assumption (tall blue lines) and the proton fit assumption (short purple lines) for a simulated 2 GeV/ c muon with a track length of approximately 850 cm. The scattering planes locations are measured along the track length s .

The plot in figure 6.19 is an example of where the BPF algorithm chose to put the scattering planes under the muon fit and proton fit assumptions, for a simulated 2 GeV/ c muon. First, it is clear that the scattering planes get closer together as the particle reaches the end of its trajectory where it is expected to scatter more. Second, at the beginning of the track, the proton fit assumption scattering planes are farther apart than the muon fit assumption scattering planes. This occurs because the proton would have to have more momentum than the muon in order to produce a track of the same length, and would therefore be expected to scatter less (under the Coulomb multiple scattering model) at the beginning of the track. Third, the last four scattering planes for both the muon and proton fit assumptions occur at exactly the same locations due to the restriction that there can not be more than one scattering plane in between each measurement plane.

Shown in figure 6.20 is an event display from the same selected far detector data event seen in figure 6.16. For simplicity, only the tracks made under the muon fit assumption are drawn. The zoom in on the end of the longer track (assumed to be a real muon) shows the tracker's ability to follow the path outlined by the cell hits, and includes a more dramatic scattering at the track end where the particle had the least momentum.

In an attempt to isolate the performance of BPF from the other reconstruction modules, I opted to use samples of simulated single particles over using simulated neutrino events, which can often be difficult to reconstruct and can introduce poor performance caused by upstream failures in the other modules. Since the analysis I discuss in chapter 7 focuses on identifying the muon track in ν_μ CC interactions, I will focus here on the performance of BPF with respect to reconstructed muons, specifically examining the reconstructed momentum and track directions.

In generating samples of single muons to use for BPF performance assessment, I discovered that in about 1% of events, the Elastic Arms algorithm showed a tendency to displace

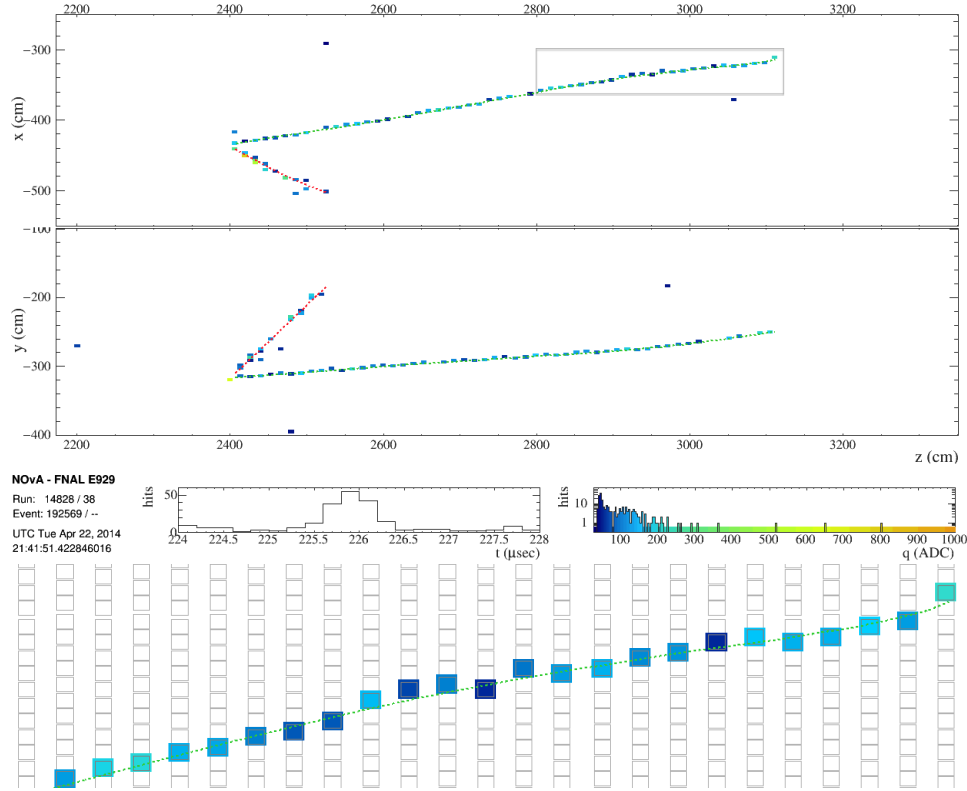


Figure 6.20: Top: Reconstructed far detector neutrino candidate event showing the tracks produced by the BPF algorithm. Only the tracks produced by the muon fit assumption are shown. **Bottom:** Zoom in on the track trajectory shown in the gray box in the top figure. Cells borders are outlined in light gray and cell hits are color coded by the charge deposited according to the scale shown on the bottom right of the top figure.

the event vertex forward along the muon track. This effect is noticeably more pronounced in single track events where a second track is not present to help identify the vertex. An example can be seen in figure 6.21. While the longer blue prong in this figure contains most of the muon track, the short green prong will have a poorly reconstructed value for the muon momentum, which led to a 6% broadening of the momentum fractional bias plots [78]. Since this failure occurs within the Elastic Arms algorithm and not within BPF, I removed this small sample of events by requiring the BPF reconstructed momentum to be at least half of the true momentum. All performance plots that follow were made using this “half-p” cut.

To assess BPF’s ability to reconstruct the momentum of a muon, I generated four samples of 1000 single muons in the center of the far detector with momenta of 0.5, 1.0, 2.0, and




Figure 6.21: An example of the event vertex being pulled forward along the track by the Elastic Arms algorithm, justifying the “half-p” cut. The event shown is a simulated single muon with a momentum of 2 GeV/c traveling from left to right. The Elastic Arms vertex is drawn as a blue “X” with 2 back-to-back Fuzzy-K 3D prongs. The blue prong contains most of the muon and the short green prong points back towards the true muon origin.

4.0 GeV/c, reconstructing each sample with the BPF muon fit assumption. The 0.5, 1.0, and 2.0 GeV/c samples were generated with a uniform distribution of initial track angles between 0 and 60 degrees with respect to the detector z -axis. For containment reasons, the 4 GeV/c sample was restricted to initial angles between 0 and 15 degrees. The distributions of momentum fractional bias values, defined as $(p_{\text{reco}} - p_{\text{true}})/p_{\text{true}}$, for these four samples are shown in figure 6.22 with the mean and RMS values for each distribution listed in table 6.6. The RMS values in this table indicate that the BPF momentum resolution remains relatively constant at 5% across a range of muon momenta, with the resolution improving for larger true momentum values. The mean values in this table suggest that BPF has a consistent tendency to overestimate the true momentum by about 2%. This effect is most likely due to differences between the dE/dx tables in [77] used by the BPF Path class and those used by GEANT4 to simulate the muon samples. These differences will be corrected in future analyses.




Figure 6.22: Momentum fractional bias plots for four simulated single muon samples. The mean and RMS values for these distributions are listed in table 6.6.

sample	mean	RMS
0.5 GeV/ c μ	0.028	0.061
1.0 GeV/ c μ	0.019	0.053
2.0 GeV/ c μ	0.020	0.050
4.0 GeV/ c μ	0.020	0.048
1.0 GeV/ c π	-0.463	0.226
1.0 GeV/ c p	-0.040	0.078

Table 6.6: The mean and RMS values for the fractional bias distributions shown in figures 6.22 and 6.23.

As a quick assessment of BPF’s ability to reconstruct pion and proton momenta, I also generated a sample of 1000 single pions and 1000 single protons each with a momentum of 1.0 GeV/ c with the same 60 degree spread in initial angles. The pion sample was reconstructed with the BPF pion fit assumption and the proton sample with the BPF proton fit assumption. I restricted myself to 1.0 GeV/ c particles because 0.5 GeV/ c protons leave too few hits to produce a reliable track, and 2.0 GeV/ c pions and protons begin to have strong interactions which are not included in the BPF multiple scattering model. The

distributions of momentum fractional bias values are shown in figure 6.23 with the mean and RMS values for each distribution listed in table 6.6. BPF seems to do reasonably well with the protons, resulting in a momentum resolution of about 8%, but very poorly with the pions. After examining a few event displays from the single pion sample, it was clear that this sample suffered from a variety of things that would cause the reconstruction to fail, including hard scattering of the pion causing improper placement of the vertex, absorption of the pion into a nucleus resulting in the ejection of one or more protons, and mid-flight decay of the pion into a muon. Examples of these types of failures can be seen in figure 6.24. Since the reconstruction discussed in this chapter is aimed at events with neutrino-like topologies which often have vertices well defined by multiple particle tracks, failures seen in reconstructing isolated single pions might not impact the pions that we expect to see in real neutrino data. The analysis I discuss in chapter 7 does not use either the BPF pion or proton reconstruction information, so it is unlikely to be impacted by these kinds of failures. However, any future analysis that does wish to use this information will require looking further into the BPF algorithms ability to reconstruct tracks for particles other than muons.

The other major BPF tracking performance metric I examined was the reconstructed initial direction of the track. For this assessment, I generated five samples of 1000 single muons with momenta of 0.5, 1.0, 2.0, 4.0, and 8.0 GeV/ c . All muons in these samples were generated at exactly 5 degrees with respect to the detector z -axis. The distributions of reconstructed track angles for these samples are shown in figure 6.25 with the mean and RMS values for each distribution listed in table 6.7. The central values for these distributions are all around 5 degrees with resolutions less than a few degrees. To test some larger angles, I generated three more samples of 1000 single muons with 2 GeV/ c momenta each, at 30, 60, and 80 degrees with respect to the detector z -axis. The distributions of reconstructed




Figure 6.23: Momentum fractional bias plots for the 1.0 GeV/c simulated single pion and single proton samples.

track angles for these samples are shown in figure 6.26 with the mean and RMS values for each distribution listed in table 6.7. The central values for these distributions all line up with the true angles with resolutions of a few degrees. The 80 degree sample has an additional cluster of values centered around 90 degrees. Some of the muons in this nearly vertical sample will scatter so as to remain within the same detector plane, so a small group of events with a reconstructed track angle of 90 degrees is expected.

sample	mean (degrees)	RMS (degrees)
0.5 GeV/c μ , 5°	5.83	2.27
1.0 GeV/c μ , 5°	5.47	1.37
2.0 GeV/c μ , 5°	5.19	0.75
4.0 GeV/c μ , 5°	5.15	0.43
8.0 GeV/c μ , 5°	5.12	0.25
2.0 GeV/c μ , 30°	30.0	0.74
2.0 GeV/c μ , 60°	60.3	4.34
2.0 GeV/c μ , 80°	80.5	3.88

Table 6.7: The mean and RMS values for the reconstructed track angle distributions shown in figures 6.25 and 6.26.




Figure 6.24: Examples of the types of pion reconstruction failures that led to poor estimates of pion momenta. In the top figure, the pion (pink track) has a strong interaction knocking a proton (purple track) out of a nucleus which deposits the majority of the energy in the Fuzzy-K 3D prong (blue outline). In the bottom figure, the pion has had several hard scatters causing the Elastic Arms vertex (blue “X”) to be placed in the wrong spot.




Figure 6.25: Distribution of reconstructed track angles for five samples of single muons all generated at 5 degree angles to the detector z -axis. The mean and RMS values for these distributions are listed in table 6.7.




Figure 6.26: Distribution of reconstructed track angles for three samples of single 2 GeV/ c muons, generated at 30, 60, and 80 degree angles to the detector z -axis. The mean and RMS values for these distributions are listed in table 6.7.

CHAPTER 7

The Break Point Fitter Analysis Suite

In this chapter, I outline the details of the analysis steps performed on the reconstruction products described in chapter 6, that are unique to my analysis. I give a more general description of the ν_μ CC disappearance steps (event selection, near-to-far energy spectra extrapolation, etc.) in chapter 8. The methods described here include using the results of the Break Point Fitter tracking algorithm (see section 6.7) to identify muon tracks and to estimate energy and energy resolution on an event by event basis for ν_μ CC interactions. At the end of this chapter, I show some examples of performance with respect to simulated data.

7.1 Summary of Motivations and Methods

The primary motivation behind my development of the BPF ν_μ CC disappearance analysis was to improve our event energy estimations and our final sensitivity contours for $\nu_\mu \rightarrow \nu_\mu$ oscillations. Improved event energy estimations can be accomplished by taking advantage of Fuzzy-K's ability to pick out hadron tracks, rather than simply accounting for the hadronic energy in each event calorimetrically. The BPF track model can lead to event energy improvements using BPF's built-in momentum estimate for muon tracks which automatically incorporates energy loss through different detector materials. As shown in figure 2.5 the difference between measuring $\sin^2 \theta_{23} = 0.4$ and $\sin^2 \theta_{23} = 0.5$ depends greatly on accurate

energy estimations of the events specifically near 1.5 GeV. Our final sensitivity contours can be improved if, in addition to improving the energy estimations for all events, we can also provide an estimate on an event by event basis of the energy resolution. Events near 1.5 GeV that are estimated to have good energy resolution can be treated as “more important” for the final fit than events with poor energy resolution, leading to better contours. In principle, this could make a bigger difference for the “maximal” oscillation case ($\sin^2 \theta_{23} = 0.5$) where a few improperly binned events at the oscillation maximum would pull the best fit for the oscillation parameters away from the “maximal” solution.

The first step of this analysis identifies which track is the most “muon-like.” This is important for ν_μ CC energy estimation since the muon carries a significant fraction of the total event energy and muon energies can be well estimated using information about the track length and dE/dx . To find the muon track, I developed a “muon ID” algorithm based on variables computed from track parameters. This method, the input variables used, and the results are discussed in section 7.2. With the most muon-like track identified, the next step is to estimate the energy and energy resolution for each event. For this step, I developed a multi-variate regression method which is described in section 7.3 along with some example results.

7.2 Muon Identification

Accurately identifying the muon track is important for ν_μ CC event energy estimation, and will allow the BPF energy estimator to use the momentum computed from the BPF tracks with the help of the BPF path-length dE/dx integrator, `Path` (described in section 6.7.2.) However, determining which track is the most muon-like with greater than 95% accuracy can be challenging. Any individual characteristic of a true muon track (length, dE/dx , etc.) often shares a common phase space with tracks made by non-muons. Therefore, I chose a

method using a k-nearest-neighbor (kNN) algorithm [79], to combine the selective power of several variables. For this method, I picked four quantities based on their discriminating power. These variables are track length, the ratio of hits on the track to hits in the Fuzzy-K prong, the total χ^2 from the BPF track fit (equation 6.18), and a dE/dx log-likelihood value. Track length, the hit ratio variable, and the total χ^2 all come straight from the reconstruction results. However, the dE/dx log-likelihood variable is not part of the reconstruction results and is calculated afterwards. I will discuss the method I developed to compute this variable in the next section.

7.2.1 The Muon Track dE/dx Log-Likelihood Calculation

Calculating a dE/dx log-likelihood value for each track is based on the idea that particles will lose energy along the length of their track according to some expected value that will be different for different particle types. Measuring the energy loss rate (dE/dx) along a given track and comparing it to the value expected for a muon can therefore help to determine if the track is “muon-like” or not. For the range of muon energies expected in NOvA ν_μ CC events, this energy loss rate (known as the “stopping power”) is dominated by electromagnetic interactions leading to energy transfer from the penetrating particle to the surrounding material through atomic excitation and ionization. This is modeled well by the Bethe formula given by

$$\left\langle -\frac{dE}{dx} \right\rangle = \frac{Kz^2Z}{2A\beta^2} \left[\ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - 2\beta^2 - \delta(\beta\gamma) \right], \quad (7.1)$$

where W_{max} is the maximum energy transfer in one collision given by

$$W_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2}, \quad (7.2)$$

where z and M are the charge and mass of the incident particle, Z , A , and I are the atomic number, atomic mass, and mean excitation energy of the absorbing material, $K = 0.307075$ MeV mol⁻¹ cm², m_e is the electron mass, and $\beta\gamma = p/Mc$. $\delta(\beta\gamma)$ is a density effect correction due to polarization in the absorbing material [8]. Equation 7.1 is accurate to within a few percent for muons with approximately $0.01 \leq p \leq 100$ [GeV/ c] traveling through intermediate- Z materials [8]. A derivation and more detailed description of the Bethe equation can be found in reference [8].




Figure 7.1: Stopping power for μ^+ in copper. Equation 7.1 is accurate to within a few percent in the region labeled “Bethe.” This plot was taken from reference [8].

A plot showing an example of the stopping power of muons in copper is shown in figure 7.1. Equation 7.1 is a reasonable approximation in the region labeled “Bethe.” Since this equation represents the mean stopping power, it is pulled towards the larger values that occur in single, high energy collisions. An expression for the “most-probable” value of dE/dx would therefore be more desirable. The most-probable accumulated energy loss for a particle traveling a distance x through a material is derived in [8] using a Landau-

Vavilov distribution. However, this expression is written in terms of x (which is typically very different for different cell hits) making its implementation as an estimation of the instantaneous value of dE/dx difficult.

A log-likelihood approach can be taken by generating a model for the probability of getting a specific dE/dx value for a particle with a given $\beta\gamma$. I accomplished this by plotting dE/dx values vs. $\ln(\beta\gamma)$ measured from true muon tracks reconstructed under the BPF muon fit assumption. Using $\beta\gamma$ instead of the particle momentum allowed me to make one set of plots that work for particles of different masses. I used the log of $\beta\gamma$ which naturally produces bins of uniform fractional size, picking bin widths of 5% to reflect the 5% BPF muon momentum resolution discussed in section 6.7.3. The values for dE/dx and $\ln(\beta\gamma)$ are calculated for each cell intersected by the BPF track trajectory. The value for dx is calculated by determining the track entrance and exit points within the cell. The value for dE is taken from the measured energy deposited in the cell, which has been corrected for attenuation using the middle of the track segment within the cell as the estimated distance to the APD readout (a coordinate referred to as “ W ”). The distance along the track from this middle location to the end of the track is then passed to the BPF path-length dE/dx integrator `Path`, which returns the estimated value of $\beta\gamma$ for that cell under the muon fit assumption. These values are plotted in a 2D histogram (referred to hereafter as the “dEdxLL plot”) and each vertical strip one bin wide is normalized to an area of 1.0 so that each bin represents for each cell hit, the probability of getting a specific value of dE/dx for a particle with a given value of $\beta\gamma$.

Before generating the final versions of the dEdxLL plot, I considered the possibility that my dE/dx calculations could be skewed by two things. First, the W coordinate could affect dE since tracks passing through cells farther from the readout will be more susceptible to attenuation and noise. Second, steep tracks (with long path lengths through a cell) will

likely have less reliable estimates of dx and will be more affected by the fact that the method for calculating dE/dx described above computes only one value of dE/dx for each cell since dE/dx will be changing throughout the cell. To test these effects, I generated dEdxLL plots using four samples of 3 GeV/ c muons simulated in the far detector. The first two samples tested the effects of different W coordinates and were simulated in the upper-West corner of the detector (i.e. close to the APD readout) and the lower-East corner (i.e. far from the APD readout.) The dEdxLL plots for these two samples along with an asymmetry plot comparing the two can be seen in figure 7.2. The last two samples tested the effects on tracks with different values of dx by comparing a sample of “shallow” muon tracks, generated between $\pm 15^\circ$ with respect to the detector z -axis, to a sample of “steep” muon tracks, generated with angles between 70° and 90° . The results of comparing the last two samples can be seen in figure 7.3.

The plots in figures 7.2 and 7.3 are very similar but the asymmetry comparisons show some non-trivial differences due to the effects listed above. To get a better representation of the distributions of dE/dx values for tracks in different locations of the detector and at different angles, I applied a simple solution of generating four different dEdxLL plots for two ranges each of W and dx values. The computed values of dE/dx and $\ln(\beta\gamma)$ for each cell were used to fill one of these four plots according to value of W and dx for that particular cell. The W ranges were chosen by simply splitting the cell into the half closer to the APD readout ($W > 0$) and the half farther away ($W \leq 0$.) For the dx ranges, I picked a dividing line at $dx = 8$ cm by examining the distribution of dx values from my simulated steep muon track sample and from tracks in simulated genie events (shown in figure 7.4.)

To cover the phase space represented in these four dEdxLL plots, I used the reconstructed tracks from four different simulation samples. The first two samples were each 1×10^6 single 10 GeV/ c muons generated in the bottom (top) of the detector traveling up




Figure 7.2: Computed values of dE/dx vs. muon momentum (instead of $\ln(\beta\gamma)$) for simulated 2 GeV/ c muons in the upper-West quadrant of the far detector (top left) and in the bottom-East quadrant (top right.) The bottom plot shows the bin by bin asymmetry computed as $(\text{bottom-East} - \text{top-West})/(\text{bottom-East} + \text{top-West})$.

(down) at shallow angles. The last two samples were single 2.7 GeV/ c muons (also with 1×10^6 in each sample) generated in the bottom (top) of the detector traveling up (down) at steep angles. Each sample was generated with a spread of different initial angles.

An example of one of the four dEdxLL plots created from the single muon samples described above, is shown in the top plot of figure 7.5. At the low momentum end of the muon tracks, some discretization of the computed values for $\beta\gamma$ is occurring. This is evident in the vertical striping pattern that appears around $\ln(\beta\gamma) = -1.2, -0.2, +0.1$, etc. For a muon, $\ln(\beta\gamma) = -1.2$ corresponds to a momentum of about 30 MeV/ c and $\ln(\beta\gamma) = -0.2$ corresponds to about 80 MeV/ c . This difference of 50 MeV/ c is reasonable for the amount of energy a muon is expected to lose near the end of its track as it crosses one cell (assuming




Figure 7.3: Computed values of dE/dx vs. muon momentum (instead of $\ln(\beta\gamma)$) for simulated 2 GeV/ c muons in the far detector with shallow tracks (top left) and steep tracks (top right.) The bottom plot shows the bin by bin asymmetry computed as $(\text{steep} - \text{shallow})/(\text{steep} + \text{shallow})$.

a path length of 6 - 10 cm.) Given that my dE/dx calculation is done on a cell-by-cell basis, this kind of discretization is expected.

Since my intention behind using $\ln(\beta\gamma)$ on the x -axis instead of momentum was to use one set of plots for all particle types, this kind of gap can cause difficulties. Naturally, this gap is not a problem when using this dEdxLL plot to identify tracks from true muons since hits on muon tracks clearly do not occupy that region of the plot. But trying to use this plot for particles whose mass is significantly different from the muon could result in missing information. For example, a dEdxLL plot generated from a sample of 1 GeV/ c protons reconstructed with the BPF proton fit assumption is shown in the bottom plot of figure 7.5. It is clear from this plot that for protons, most of the relevant dE/dx information falls right




Figure 7.4: Distribution of computed dx values for the steep track sample (red) and a sample of tracks from 100 reconstructed GENIE events. The dividing line for the two dx ranges was chosen to be 8 cm since this broke the steep track sample up into two distinct populations and most GENIE tracks had cell path length less than this value. Both plots are normalized to unit area.

in the gap seen in the muon plot, which could result in an inaccurate dE/dx log-likelihood value.

The solution I applied to the problem of the gaps in the muon plots was to generate four more single particle samples under the exact same initial conditions as the original muon samples, using 1 GeV/ c protons in place of the muons. I decided to use 1 GeV/ c protons since they filled in the gap in the muon plot well while minimizing the overlapping information. Using lower energy protons also minimizes energy losses due to strong interactions, which are not modeled by equation 7.1.

For the final step in creating the dEdxLL plots, I added the plots created from reconstructing protons under the BPF proton fit assumption with the plots created from reconstructing muons under the BPF muon fit assumption, and normalized the results in vertical strips as described above. An example of the final dEdxLL plot for cell hits with $W > 0$ and $dx < 8$ cm is shown in figure 7.6. Some discretization is still evident in this plot, but the gap in the muon plot has been significantly filled in and the dE/dx information




Figure 7.5: Top: Plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, generated from 10 GeV/ c muon tracks reconstructed with the BPF muon fit assumption. **Bottom:** Plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, generated from 1 GeV/ c proton tracks reconstructed with the BPF proton fit assumption.

from the protons lines up reasonably well the information from the muons.

Armed with the finished set of dEdxLL plots, I compute a dE/dx log-likelihood value for each BPF track. Since the purpose is to identify muons, I use only the track reconstructed under the BPF muon fit assumption. The premise is that for true muon tracks, the muon fit assumption will give correct values for $\beta\gamma$ causing the calculated dE/dx values to line up with the dEdxLL plots, resulting in a high log-likelihood value. The log-likelihood value




Figure 7.6: Normalized plot of dE/dx vs. $\ln(\beta\gamma)$ for cell hits with $W > 0$ and $dx < 8$ cm, created by combining the muon and proton results together.

for each track is calculated as

$$LL = \frac{1}{N} \sum_N \mathcal{L}(\ln[\beta\gamma], dE/dx), \quad (7.3)$$

where N is the number of cells intersected by the track trajectory, and $\mathcal{L}(\ln[\beta\gamma], dE/dx)$ is the bin content of the appropriate dEdxLL histogram (determined for each cell by the values of W and dx .) These dEdxLL plots were generated for the far detector which obviously has much longer cells than the cells in the near detector, so special handling of the near detector case is required. In order to use the same plots, for near detector muon identification, I always assume that $W > 0$ (i.e. the near detector is assumed to be just like the far detector but with every hit “close to the APD readout.”) I will discuss the results of applying this method, including showing plots of the distribution of dE/dx log-likelihood values in the next section.

7.2.2 Input Variables to the Muon ID

The first input variable to the muon ID algorithm is the dE/dx log-likelihood value calculated with equation 7.3 according to the procedure described above. By itself, this variable

has significant discriminating power and is therefore one of the most influential inputs to the muon ID. Plots showing the distribution of dE/dx log-likelihood values for signal and background are shown in figure 7.7. For all plots in this section, “signal” refers to a reconstructed track for which the particle that contributed the most energy was a true muon (μ^- or μ^+) and “background” refers to all other tracks. All tracks have been taken from simulated GENIE events in the far detector, selecting only contained true ν_μ CC events with at least 10 cell hits in each view.

The second input variable to the muon ID is the BPF track fit total χ^2 per degree of freedom value discussed in section 6.7.1 and calculated with equation 6.18 from the BPF muon fit assumption. This variable is a measure of how well the multiple-scattering model applies to the track fit under this assumption, and should be closer to 1.0 for true muons. The distribution for signal and background tracks is shown in figure 7.7.




Figure 7.7: Left: Distribution of dE/dx log-likelihood values for muon (signal) and non-muon (background) tracks. **Right:** Distribution of BPF muon track fit total χ^2 per degree of freedom values. All distributions are normalized to unit area.

In ν_μ CC interactions, the muon tends to be the longest track since they typically carry a significant amount of the total event energy and they do not lose energy by participating in strong interactions. In fact, simply selecting the longest track in a ν_μ CC interaction correctly picks the muon roughly 80% of the time. I chose to include track length as an

input variable for this reason. The distribution of track lengths for signal and background tracks is shown in figure 7.8. For this variable, track length is the path length along the trajectory from the track fit under the BPF muon assumption.

The last input variable to the muon ID algorithm is the ratio of hits intersected by the track trajectory to hits in the Fuzzy-K prong used to make the track. I discovered in an early study that it was possible for some shower-like events (mostly photons) to have reasonably good dE/dx log-likelihood values [80]. This can occur because I am computing dE/dx by only considering the cells intersected by the track trajectory, and can therefore be inaccurate for shower-like events with energy deposited off the main track. I picked this hit ratio variable as a way of helping to distinguish track-like and shower-like events. The distribution for the hit ratio variable for signal and background tracks is shown in figure 7.8. Spikes in the background distribution are caused by the fact that this variable is a ratio of two integers. This is not as prominent in the signal distribution because muon tracks tend to have enough hits to smooth this discretization out.




Figure 7.8: **Left:** Distribution of track lengths for muon (signal) and non-muon (background) tracks. **Right:** Distribution of track/prong hit ratios. All distributions are normalized to unit area.

7.2.3 The Muon Identifier

The job of the muon ID algorithm is to select the most muon-like track from all of the reconstructed tracks within a given slice. Since ν_μ CC event selection is handled by the method described in section 8.1.2, the tracks selected by the muon ID from this sample of events are assumed to be the primary muons generated in the ν_μ CC interactions. To select the most muon-like track in a slice, I generated a “muon score” for each track in the slice, taking the track with the highest score. This “muon score” is computed using a k-nearest-neighbor (kNN) classification algorithm, generated with the TMVA package [79] in ROOT [31,32]. I chose this method for its simplicity in implementation.

A kNN classification algorithm works by building up a phase space density-map for both “signal” and “background” events, distributed according to some group of input variables. An example of such a distribution is shown in figure 7.9. A collection of signal events and background events must be fed to the algorithm to build the initial version of this density map (a process referred to as “training” the kNN.) Once the training is complete, an unknown event can be classified by searching the density map to find the “k” number of nearest neighbors. The score for that unclassified event is the relative probability that the event is of “signal type” calculated as

$$P_s = \frac{\sum_i^{N_s} W_{s,i}}{\sum_i^{N_s} W_{s,i} + \sum_j^{N_b} W_{b,j}}, \quad (7.4)$$

where N_s (N_b) is the number of signal (background) events among the nearest neighbors, W_s and W_b are weights that can be set for each individual event, and $N_s + N_b = k$. In the simple case of a training sample composed of an equal number of signal and background

events, all with weight $W = 1$,

$$P_s = \frac{N_s}{N_s + N_b} = \frac{N_s}{k}. \quad (7.5)$$

A minimum value of P_s can then be chosen to select a sample of signal events that meets some desired efficiency and purity.

Given that the input variables are likely to have different units, the variable with the widest distribution will dominate a normal Euclidean metric used to compute distances within the density-map. Therefore, a weighted Euclidean metric is used when determining the nearest neighbors to an unclassified event. This weighted distance r between an event in the density-map x and the unclassified event y is given by

$$r^2 = \sum_i^d \frac{1}{w_i^2} |x_i - y_i|^2, \quad (7.6)$$

where the index i classifies the d phase space variables, and w_i is the width of the x_i distribution for the combined sample of signal and background events. So as to not be distorted by events out in the tails of the distribution, this width w_i can be chosen to contain only a certain percentage of the events around the central region of the distribution. Overall, the kNN classification method works well for situations in which the borders between signal and background are vague or irregularly shaped.

I created the BPF muon ID kNN using version 5.34/25 of ROOT [31, 32] with version 4.2.0 of the TMVA toolkit [79]. The signal and background samples were tracks reconstructed from far detector events simulated with GENIE. The events (slices) from which the training samples were picked, were required to have no cell hits within 50 cm of a wall, at least 10 hits in each view, and to be true ν_μ CC events. Signal and background tracks from these events were divided according to whether or not the particle that contributed




Figure 7.9: An example of the phase space density-map for a kNN with 2 input variables, x_1 and x_2 (taken from [79].) The filled circles are the signal events and the open circles are the background events. The star represents an unclassified event with the “nearest neighborhood” drawn in a circle around it. In this example, the number of signal and background events within the circle is roughly the same. Therefore $P_s \approx 0.5$.

the most energy to the track was a muon (μ^\pm) or not. The number of signal (background) tracks passing all cuts in the training sample was roughly 410,000 (460,000). The default mode for the kNN classification scheme (which was used for this training) is to scale the background events so that the total sum of the event weights for both signal and background is the same. Each event in the background sample was therefore given a weight of $W_b = 0.893$. The distribution widths (w_i) of each of the input variables, was set to the width of the central 80% of the values for that variable.

To choose the number of neighbors k , I trained several kNNs using $k = 20, 40,$ and 80 . In each case, I examined the percentage of slices in which the track with the best muon ID value was a true muon for true ν_μ CC events and for true ν_μ CC events with a best muon ID value > 0.4 . For both of these samples, these percentages showed essentially no variation for the different values of k . I chose to use $k = 80$ since the percentages were slightly higher for this value in the sample of true ν_μ CC events.

The distribution of muon ID values for signal and background tracks is shown in figure 7.10. The data used to generate both plots in this figure was taken from a set of far detector events simulated with GENIE, separate from the training sample, but selected by the same criteria used to select the training events. The second plot in figure 7.10 shows the same distribution for background tracks only, broken down by particle type. The background track sample was approximately 40% protons, 40% photons, and 20% π^\pm .




Figure 7.10: Top: Distribution of BPF muon ID values for signal and background tracks (normalized to unit area.) **Bottom:** Distribution of BPF muon ID values for background tracks only, broken down into the particle that contributed the most energy to the track.

Since I am using this muon ID value to select the most “muon-like” track from a set of events selected as ν_μ CC, the best performance metric for the BPF muon ID value is how

often it correctly picks the muon in this scenario. Shown in table 7.1 are the percentage break downs by true particle type, for the most muon-like track in a slice selected by the BPF muon ID. The break downs are given for ν_μ CC events selected by truth, ν_μ CC events selected by the ReMId algorithm discussed in section 8.1.2, and true ν_μ CC events selected by ReMId (included for comparison since the ReMId selected sample does include a small number of NC events). ReMId also uses a kNN with input variables from a different tracking algorithm and is used to select events for the standard NOvA ν_μ CC disappearance analysis. I use ReMId to select the events for my ν_μ CC disappearance analysis as well. For this sample, the BPF muon ID correctly selects a muon track over 96% of the time. The misidentified non-muon tracks are roughly an equal mixture of protons, photons, and π^\pm .

particle	true ν_μ CC	ReMId selected	ReMId and true ν_μ CC
μ^\pm	87.6	96.3	97.9
γ	4.3	1.1	0.7
p	4.9	0.9	0.6
π^\pm	3.0	1.6	0.8

Table 7.1: Percentages for the particle that contributed the most energy to the track selected by the BPF muon ID to be the most “muon-like” track within a slice. Roughly one third of all events are NC. The ReMId selection algorithm is discussed in section 8.1.2.

7.3 The BPF Energy Estimator

The BPF energy estimator is the final step in the BPF analysis. It provides an energy and an estimate of the energy resolution for each neutrino interaction seen in the NOvA detectors. These reconstructed energies and their resolutions are used to generate the final sensitivity contours for my ν_μ CC disappearance analysis, which is described in detail in chapter 8. The energy estimator takes as input the output from the NOvA reconstruction steps described in chapter 6. For convenience, a summary of these steps appears below.

- **Step 1:** The first step produces “slices” (section 6.3) which are clusters of spatially

and temporally correlated hits.

- **Step 2:** The next step produces a vertex (sections 6.4 and 6.5) for each slice, representing the estimated neutrino interaction location.
- **Step 3:** Following vertex identification is the creation of “prongs” by the Fuzzy-K algorithm (section 6.6.) Prongs are clusters of hits believed to belong to a single particle, and include a start point and a direction.
- **Step 4:** The last step is to make particle tracks from the Fuzzy-K prongs. This is done by the BPF tracker (section 6.7) which creates three tracks under the three particle assumptions for each 3D Fuzzy-K prong.

7.3.1 The Energy Estimation Method

To build the BPF energy estimator, I used a multi-variate regression method trained with a large sample of far detector neutrino events simulated with GENIE. The regression method is similar to the classification method described above. It must be fed a sample of “training” events parameterized by some set of input variables, to generate some kind of phase space mapping. However unlike a classification method, a regression method takes only signal events and maps the input variables for that event to one or more “target” variables associated with the event.

For this energy estimator, each input event was a slice associated with a single neutrino interaction. Each of these events was parameterized by six input variables related to the event energy, which are described in detail below. To map these six inputs to an estimated event energy and event energy resolution, I used two target variables. The first was the true neutrino energy E_ν and the second was the true energy squared E_ν^2 . I estimated the energy resolution σ_E for each event using

$$\sigma_E^2 = \overline{E_\nu^2} - \overline{E_\nu}^2. \quad (7.7)$$

To generate the energy estimator, I turned again to the ROOT [31,32] TMVA toolkit [79]. In the early development stages, I tried out two different multi-variate techniques: a “multilayer perceptron” (MLP) and a kNN. The MLP method is an artificial neural network which uses a “feed-forward” system of layered neurons to generate a value for the target variables based on the neuron’s response to the inputs. I chose the MLP method because it worked extremely well with some of my initial toy models. However, the performance of the MLP was not as good when mapping to more than one target variable and showed some instability in the situation where the targets were complex functions of the input variables (as is the case when mapping reconstruction variables to neutrino energies.) On the other hand, the performance of the kNN regression method was overall more stable across a variety of different types of event, and the mapping to the first target variable was totally unaffected by adding a second one.

The kNN regression method works in a similar way to the kNN classification method, by building a map in the input variable phase space from a set of training events. Each of these events has a value for the target variable(s) associated with it, and can be assigned an individual weight. For a test event, the kNN will find the k nearest neighbors and compute the regression value(s) for the test event as the weighted average of the neighbor’s target value(s). That is, the regression value V for a test event is given by

$$V = \frac{\sum_i^k W_i T_i}{\sum_i^k W_i}, \quad (7.8)$$

where W_i the individual event weight and T_i is the target value for each event within the group of nearest neighbors. If multiple target values are associated with each event, then a regression value can be computed for each one.

7.3.2 Input Variables to the Energy Estimator

I chose six variables for the inputs to the BPF energy estimator. Some of them were chosen due to their direct relationship to the original neutrino energy, some were chosen because they were more closely related to our ability to estimate σ_E , and some were chosen simply because they would help the kNN match to events with similar topologies in the training sample. The six input variables are the reconstructed momentum for the muon track, the z -component of the muon track direction, the number of 3D Fuzzy-K prongs, the total summed energy in the 3D Fuzzy-K prongs, the remaining summed energy in the slice, and the number of photo-electrons (PE) summed over all hits not on the muon track. Plots showing the distribution of each of these variables are shown in figure 7.11.

The muon momentum (P_μ) and track direction ($DirZ_\mu$) variables were chosen because of their direct relationship to the neutrino energy for ν_μ CC events. In the case of a two-body QE interaction, these are the only two variables needed to calculate E_ν directly with

$$E_\nu^{QE} = \frac{M_n - \frac{1}{2}(M_n^2 + M_\mu^2 + M_p^2)}{M_n - E_\mu + \sqrt{E_\mu^2 - M_\mu^2} \cos \theta}, \quad (7.9)$$

where θ is the angle between the incoming neutrino and the outgoing muon. The muon kinematic variables remain important for non-QE events as well since the muon still carries a significant fraction of the neutrino energy which helps the kNN match to a region of the input variable phase space with similar events. The muon track was identified using the BPF muon ID algorithm described in section 7.2. The momentum for this track was calculated using the BPF path-length dE/dx integrator `Path` (described in section 6.7.2) and $DirZ_\mu$ was simply taken as the z -component of the reconstructed track direction. This z -component is measured with respect to the z -axis of the detector, which as noted in chapter 3 is about 3° different from the direction of the incoming beam neutrinos due to




Figure 7.11: Distributions for the six input variables to the BPF energy estimator. Comparisons of these variables to data are discussed in chapter 8.

the curvature of the Earth over the 810 km baseline. However this variable is not being used directly in equation 7.9, and this 3° will also be present in all of the events used to train the kNN, so no negative effects due to this offset are expected.

I divided the remaining energy in the slice (not on the track identified as the muon) into two categories based on whether or not it came from “trackable” activity. I did this on the premise that the other reconstructed tracks in the slice represented information of a

different quality than a simple collection of hits and should therefore be treated differently. The first of these categories was the sum of the energy on all 3D Fuzzy-K prongs (E_{3D}) excluding the one identified as the muon. I computed this variable for each event as the sum of the absolute energy of the cell hits in these prongs. The absolute energies are corrected for attenuation using the initial prong direction projected through the cell to compute the cell's W coordinate. Since E_{3D} is a sum over all non-muon 3D prongs, I also included the total number of 3D prongs (N_{3D}) to help distinguish events with different numbers of hadrons for which E_{3D} might be the same.

All of the remaining energy in the slice was put into the variable E_{remain} . This variable is the sum of two contributions. The first contribution is the sum of the absolute energies over all remaining cell hits in the slice. This includes hits in 2D Fuzzy-K prongs as well as all remaining slice hits, and will account for energy from things like stray photons, neutrons, and any non-reconstructable tracks. Since these hits do not belong to 3D prongs, an accurate estimate of where the hit occurred along the length of the cell is not available. To compute the absolute energy of these hits, I used the average of the x or y coordinates from the cell hits in the opposite view (computed from just the non-3D prong hits) for the cell's W coordinate.

The second contribution to the E_{remain} variable comes from hadronic energy found superimposed on the track identified as the muon, which can happen for cells at the beginning of the track near the neutrino interaction vertex. This energy does not affect my calculation of muon energy (or momentum) since the BPF algorithm performs a path-length integration using the expected energy loss. However, a significant amount of hadronic energy can go unaccounted if it overlaps with the muon track since this energy won't be included in the E_{3D} variable.

To compute the hadronic energy contamination on the muon track, I made use of an

existing algorithm within the NOvA analysis code framework called `TrackCleanUpAlg`. This algorithm works by finding a “vertex region” which appears to have energy deposited above what would be expected from a muon in the “minimum-ionization” region of the Bethe-Bloch curve (see figure 7.1.) The size of this vertex region is determined dynamically for each event by computing dE/dx on a plane by plane basis for the hits along the muon track. This is done by starting at the beginning of the track, and finding the point where the computed dE/dx values, averaged over three consecutive planes, drop below 0.0022 GeV/cm and stay there. An example of the distribution of dE/dx values along a muon track leading to a clear identification of the vertex region is shown in figure 7.12. For cell hits on the muon track in this vertex region, the expected energy deposited by the muon alone is estimated from the track path length through the cell and an assumed maximum energy loss rate of 0.0022 GeV/cm. Any energy deposited in the cell above what is expected based on this estimate is considered to be “hadronic contamination” and is added to the E_{remain} variable. An example showing the distribution of hadronic contamination energy found on true muon tracks is shown in figure 7.12. Further details on the `TrackCleanUpAlg` algorithm can be found in [81].

The last variable for the BPF energy estimator is Sum_{PE} . As the name implies, this variable is the sum of the PE from all cell hits in the slice excluding those on the track identified as the muon. It is similar to E_{3D} and E_{remain} , but as a raw measurement of PE, it does not include attenuation corrections. I included this variable to help improve the estimate of σ_E by distinguishing the amount of light seen in an event from the estimated calorimetric energy. Events that occur farther from the APD readout (for example, in the bottom-East corner of the far detector) will have on average less total PE than similar events closer to the APD readout, and should therefore have a poorer energy resolution.

Using a kNN regression method makes determining the “importance” of each of these




Figure 7.12: **Left:** Example of dE/dx values calculated by the `TrackCleanUpAlg` algorithm along the muon track used to determine the vertex region. This plot was taken from [81]. **Right:** Distribution of hadronic energy contamination computed by the `TrackCleanUpAlg` algorithm for muon tracks produced by BPF and identified by the BPF muon ID. This plot was made from a small sample of contained, true ν_μ CC events, simulated for the far detector with GENIE.

input variables difficult. However, using the RMS widths of energy fractional bias plots (like those shown in figures 7.16 and 7.17) divided by truth into QE and non-QE samples as a metric, I made the following observations. Not including the variable N_{3D} slightly improved the QE sample and made the non-QE sample worse by about 0.1% each. Not including the $DirZ_\mu$ variable made the QE sample better by about 0.1% but made the non-QE sample worse by about 0.5%. Including the variable Sum_{PE} improved the non-QE sample by about 0.5% (at the expense of a 0.1% drop for the QE sample) while making improvements to the width of the $\Delta E/\sigma_E$ plot (see section 7.3.4 for a description of this plot.) Lastly, including the hadronic energy contamination in the E_{remain} variable improved the fractional bias for both the QE and non-QE samples by almost 1% each.

7.3.3 Examining Bias in the Energy Estimator

I created the BPF energy estimator kNN using version 5.34/25 of ROOT [31, 32] with version 4.2.0 of the TMVA toolkit [79]. The input events were slices made from neutrino

interactions simulated with GENIE for the NOvA far detector. As with the events selected for the muon ID training, I required all training events to have no cell hits within 50 cm of a detector wall, at least 10 hits in each view, and to be true ν_μ CC events. After these cuts, the training sample had a total of approximately 430,000 events. To choose the optimal value of k , I examined the width of the fractional bias plots and the width of my energy resolution estimation plot (see figures 7.16 and 7.17.) The fractional bias plots showed some slight improvement for larger values of k ($k = 80$) but the energy resolution estimation plot was clearly worse. The overall shape of the reconstructed energy spectra with respect to the true energy spectra was also visibly worse for $k = 80$ due to the training sample biases discussed in this section. Therefore for the official training, I chose $k = 10$.

Given a training sample of events with some energy spectrum, matching data events to a set of nearest neighbors will naturally pull my reconstructed energy spectrum towards the most heavily populated regions of the training sample, introducing an energy bias. To explore the effects of this bias, I created three different energy estimators based on three different neutrino energy spectrum shapes. The spectrum shapes were derived from three different assumptions about the ν_μ CC disappearance oscillation parameters. Using the same set of training events, I was able to change the shape of the energy spectrum for the training sample by applying an individual weight to each event. These weights were computed according to equation 2.9 using the true neutrino energy. The first energy estimator (“kNN-1”) was trained with the assumption of no oscillations (see the true energy spectrum in figure 2.4.) The second estimator (“kNN-2”) used a “minimum” oscillation assumption, using $|\Delta m_{32}^2| = 2.4 \times 10^{-3} [eV^2]$ and $\sin^2 2\theta_{23} = 0.95$, and the third (“kNN-3”) used a “maximum” oscillation assumption with $|\Delta m_{32}^2| = 2.4 \times 10^{-3} [eV^2]$ and $\sin^2 2\theta_{23} = 1.0$. The true energy spectra for these last two assumptions can be seen in figure 2.5.

Figure 7.13 shows the true and reconstructed energy spectra from each of the three

energy estimators described above, applied to an unoscillated sample of test events. Figures 7.14 and 7.15 show the three estimators applied to the same set of test events using the “minimal” and “maximal” oscillation assumptions applied to the true event energies. From these plots, several things stand out very clearly. First, using the “right” kNN produces the least biased reconstructed energy spectrum (using kNN-1 for the unoscillated test sample, using kNN-2 for the “minimally” oscillated test sample, etc.) Second, in a similar manner using the kNN trained with the energy spectrum that is the most different from the test sample produces the most distorted reconstructed energy spectrum (using kNN-3 for the unoscillated case and vice versa.) Third, for all cases there is a spectral bias towards the most densely populated regions of the training sample. This bias is evident in the fact that the peaks in the reconstructed spectra are always taller and typically skinnier than the peaks in the true spectra.

It appears that both of the oscillated estimators (kNN-2 and kNN-3) do reasonably well with the minimally and maximally oscillated cases with only a slight bias towards the shape of their respective training samples. However, choosing the best one to use for an unknown far detector spectra can be somewhat tricky. To quantify the response of each training assumption to the three corresponding truth assumptions, I computed a total χ^2 between the true energy spectra and the reconstructed energy spectra for all 9 possibilities. The χ^2 value was based on the assumption that each bin in the energy spectra was an independent Poisson distributed number [8] using

$$\chi^2 = 2 \sum_{i=1}^N \left[\mu_i - n_i + n_i \ln \frac{n_i}{\mu_i} \right], \quad (7.10)$$

where μ_i is the model expected value (from the true energy spectra) and n_i is the “measured” value (from the reconstructed energy spectra.) These χ^2 values are listed in table 7.2 and they agree with the above statements about which training oscillation assumption best

matches each of the true oscillation assumptions. I will note however that for the case of the maximally oscillated test sample, both kNN-2 and kNN-3 produce very similar χ^2 values which is not the case for the minimally oscillated test sample. Of course when determining the best fit parameters for the data, I will be generating a comparison between an energy spectrum from data and an energy spectrum from simulations using the same energy estimator for each. So provided the data and the simulations both respond to the kNN in the same way, then some of these biases can perhaps be avoided. That being said, the numbers in table 7.2 imply that kNN-2 might produce the least bias in an unknown far detector energy spectrum. For this reason, I choose kNN-2 (with the “minimally” oscillated training sample) to use for far detector energy estimation. In section 7.4, I will show contours for a true minimal and true maximal oscillation parameter case generated with kNN-2 that demonstrate its ability to pick out the correct answer in either case. The near detector energy estimation will naturally be done with the unoscillated training sample using kNN-1.

kNN training	no osc. test sample ($\sin^2 2\theta_{23} = 0$)	min osc. test sample ($\sin^2 2\theta_{23} = 0.95$)	max osc. test sample ($\sin^2 2\theta_{23} = 1.0$)
kNN-1	48.8	24.1	51.5
kNN-2	123.7	9.1	16.6
kNN-3	333.9	20.8	15.1

Table 7.2: χ^2 values computed between the true and reconstructed energy spectra for the plots shown in figures 7.13, 7.14, and 7.15. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively.

7.3.4 Energy Estimator Performance

The plots shown in figures 7.14 and 7.15 are examples of the spectra created by the BPF energy estimator. As a test of the overall performance, I used two primary metrics. The




Figure 7.13: Reconstructed energy spectra for the three energy estimators from an unoscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.

first was the reconstructed energy fractional bias (E_{FB}) defined as:

$$E_{FB} = \frac{E_{reco} - E_{true}}{E_{true}}, \quad (7.11)$$




Figure 7.14: Reconstructed energy spectra for the three energy estimators from a “minimally” oscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.

and was obviously a direct measure of the goodness of the energy estimate. To evaluate the estimate of σ_E , I plotted $\Delta E/\sigma_E$ where $\Delta E \equiv E_{reco} - E_{true}$. If σ_E was an accurate estimate of the uncertainty on E_{reco} , then the plot of $\Delta E/\sigma_E$ should have a Gaussian shape, centered on 0.0 with an RMS of 1.0. To generate plots of E_{FB} and $\Delta E/\sigma_E$, I selected only contained far detector events (simulated with GENIE using the NuMI ν mode flux files)




Figure 7.15: Reconstructed energy spectra for the three energy estimators from a “maximally” oscillated sample of test events. kNN-1,2,3 were trained under the assumption of $\sin^2 2\theta_{23} = 0, 0.95, 1.0$ respectively. Each plot is shown with the same true neutrino energy spectra of the test sample. All three plots are scaled to the expected exposure for 3 years of nominal data taking with a 14 kiloton detector while running in the NuMI beam ν mode.

which passed our basic ν_μ CC event selection (described in section 8.1.2) and had a value of E_{reco} between 0.5 and 3.5 GeV (the region most sensitive to oscillations.) Shown in figures 7.16 and 7.17 are the plots for this sample, generated with kNN-2 for events oscillated with the minimal and maximal oscillation parameters respectively.

The mean and RMS for each of the plots shown in figures 7.16 and 7.17 are displayed in




Figure 7.16: Top: Energy fractional bias (E_{FB}) for simulated far detector events oscillated with the “minimal” parameters and reconstructed with kNN-2. The blue curve is for all events and the green curve is only the NC sample. **Bottom:** Distribution of $\Delta E/\sigma_E$ for the same events. The means and RMSs for both plots and the Gaussian fits (red curves) are shown in table 7.3.




Figure 7.17: Top: Energy fractional bias (E_{FB}) for simulated far detector events oscillated with the “maximal” parameters and reconstructed with kNN-2. The blue curve is for all events and the green curve is only the NC sample. **Bottom:** Distribution of $\Delta E/\sigma_E$ for the same events. The means and RMSs for both plots and the Gaussian fits (red curves) are shown in table 7.3.

table 7.3. Since these numbers can be skewed by the events far out in the tails, I have also included in this table the mean and RMS from a Gaussian fit to the central region (shown in red on the plots.) For the plots of E_{FB} , I have superimposed in light green the distribution for the NC events that slipped into this selected ν_μ CC event sample. As expected, the NC events are the primary component of the longer tail on the negative side.

plot	plot mean	plot RMS	fit mean	fit RMS
E_{FB} (min osc.)	-0.022	0.163	-0.0010	0.083
$\Delta E/\sigma_E$ (min osc.)	-0.248	1.504	0.0337	0.997
E_{FB} (max osc.)	-0.035	0.173	-0.0080	0.081
$\Delta E/\sigma_E$ (max osc.)	-0.378	1.536	-0.0526	0.971

Table 7.3: Shown are the numbers corresponding to the plots in figures 7.16 and 7.17. The variables “plot mean/RMS” were computed from all data in the entire plot, and “fit mean/RMS” are the parameters from the Gaussian fit to the central region. For the plots of E_{FB} , the fit range was from -0.2 to 0.2 and for the $\Delta E/\sigma_E$ plots, the fit range was from -2.0 to 2.0.

From the fractional bias plots and the numbers in table 7.3, it appears that the energy estimator performs reasonably well with no large trend towards either over or under estimation. It also appears that the performance of kNN-2 is very similar for both the “minimal” case (the one it was trained for) and the maximal case, producing an average energy resolution of roughly 10% across the selected events. Both plots of $\Delta E/\sigma_E$ have a slightly asymmetric non-Gaussian shape, but the Gaussian fit to the central region does have a mean close to 0.0 and an RMS close to 1.0, suggesting that σ_E is a reasonable estimate of the event energy resolution. But the RMS from the whole $\Delta E/\sigma_E$ distribution indicates that σ_E is typically an underestimate of the true uncertainty. However, the purpose of this variable is to separate events into groups which have similar energy resolutions to provide more power to the final fits. So even in the case that σ_E is either an over or under estimate of the true energy resolution, it can still be used to discriminate events for this purpose.

7.4 Using Energy Resolution to Improve the Final Contours

The last piece of the puzzle that is unique to my analysis is using the event-by-event energy uncertainty estimates to improve the final sensitivity contours. Without my estimate of σ_E , I could simply compute the χ^2 values between the 1D reconstructed energy spectra from data and some set of predicted spectra with different assumptions about the oscillation parameters. To take advantage of my knowledge of σ_E , I will use a 2D energy spectra, binning in σ_E/E_{reco} on the second axis. Equation 7.10 is still valid in this context if I use it to compute a χ^2 by summing over all bins in the 2D spectrum. This is equivalent to dividing my events into multiple samples, and computing a total χ^2 summed over the 1D spectrum for each sample.

To explore the best way of using a 2D energy spectra, I created a toy model for making contours using simulated far detector events. This toy model was designed to be simple and quick to allow me to easily test out different binnings for my 2D energy spectra, as well as continue to explore the idea of energy estimator bias discussed in section 7.3.3. It compares two sets of simulated far detector events, where one is treated as fake data and the other is used to create the predicted far detector spectrum. It does not include a proper extrapolation from a simulated near detector spectrum or any systematic uncertainties. All of the contours in this section were generated with this toy model. The full “proper” contour-making process with all the bells and whistles is discussed in chapter 8.

7.4.1 Determining the Energy Resolution Binning

The distribution of values for σ_E and σ_E/E_{reco} are shown in figure 7.18. These plots were generated with a set of contained far detector events simulated with GENIE, that pass our basic ν_μ CC event selection and have E_{reco} between 0.5 and 3.5 GeV. The true energy spectrum for these events was oscillated with the minimal oscillation parameters and the

event energy reconstruction was done with kNN-2.

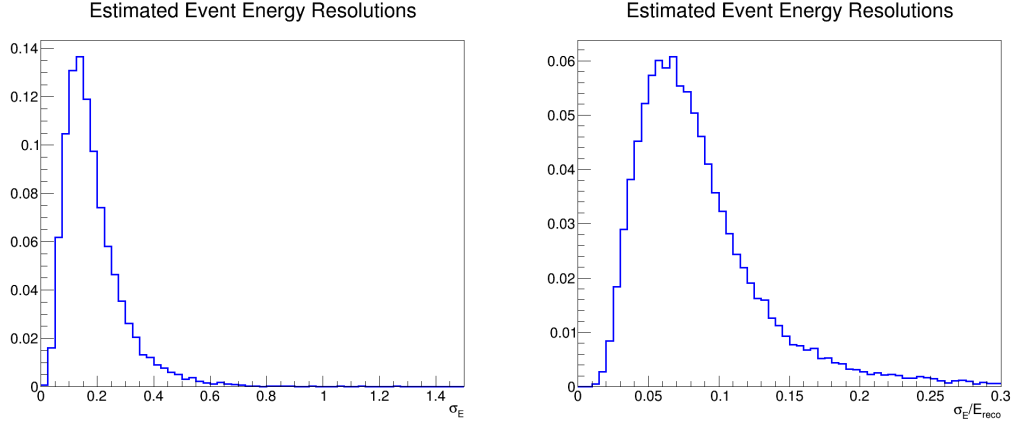


Figure 7.18: Left: Distribution of values for σ_E for a sample of simulated far detector events. **Right:** Distribution of values for σ_E/E_{reco} for a sample of simulated far detector events. All events passed a containment and basic ν_μ CC event selector cuts and had E_{reco} between 0.5 and 3.5 GeV. Both plots are normalized to unit area.

To prevent spreading my data out too thinly and reducing the statistical power, I decided that for my 2D reconstructed energy spectrum, fewer bins in σ_E/E_{reco} was better. I tried three different binnings: one bin (for a baseline comparison), three bins, and five bins. For the three and five bin cases, I picked cuts on σ_E/E_{reco} that divided the total sample into roughly equal groups. Examples of these 2D energy spectra for the selected far detector simulated events are shown in figure 7.19. The RMS of the true energy fractional bias for each of these samples is shown in table 7.4. The sensitivity contours generated from these three different binnings for σ_E/E_{reco} are discussed in the next section.

7.4.2 Generating the Sensitivity Contours

Using the toy model described above, I generated sensitivity contours in $\Delta m_{32}^2, \sin^2 2\theta_{23}$ space with the 2D energy spectra like those shown in figure 7.19. I divided my total sample of simulated events in two groups of roughly 130,000 events each (the equivalent of 70 years of data) to use one as “data” to be tested against the other, which was taken as my simulation sample. For the binning on the energy axis, I used a bin width of 0.1 GeV,

sample	σ_E/E_{reco} range	% of sample	RMS
1 bin	-	100.0	0.166
3 bins, bin 1	< 0.06	31.2	0.086
3 bins, bin 2	$0.06 \rightarrow 0.0925$	34.7	0.131
3 bins, bin 3	≥ 0.0925	34.1	0.236
5 bins, bin 1	< 0.0475	16.7	0.070
5 bins, bin 2	$0.0475 \rightarrow 0.065$	20.3	0.099
5 bins, bin 3	$0.065 \rightarrow 0.0825$	19.5	0.127
5 bins, bin 4	$0.0825 \rightarrow 0.1125$	22.1	0.171
5 bins, bin 5	≥ 0.1125	21.3	0.264

Table 7.4: Showing the energy fractional bias RMS values for the reconstructed events separated into 1, 3, and 5 samples by σ_E/E_{reco} .

which makes the bin width 5% at 2.0 GeV. I chose this width on the premise that separating events by energy resolution will only have a positive impact if the bins are of the size of what I believe to be the resolution of my “best” sample. Making bins larger than this will degrade the effectiveness of this high resolution sample.

To continue to test biases introduced by using an energy estimator trained to a specific oscillation assumption, I used three different sets of oscillation parameters for the data, always using kNN-2 to estimate the energy for all events. My three different oscillation parameter sets included the parameters used to train kNN-2 (the “minimal” parameters), the parameters used to train kNN-3 (the “maximal” parameters), and an “alternate” set not used for any training ($\Delta m_{32}^2 = 2.35 \times 10^{-3} [\text{eV}^2]$ and $\sin^2 2\theta_{23} = 0.975$.) I also generated contours for two different exposure assumptions, 18×10^{20} POT which is equivalent to 3 nominal years of running, and 2×10^{20} POT which is the estimated exposure for our first analysis data set. The toy model sensitivity contours are shown in figures 7.20, 7.21, and 7.22.

There are three things that stand out from these contours. First, it appears that kNN-2




Figure 7.19: 2D distributions of the reconstructed event energies using 1 bin, 3 bins, and 5 bins for σ_E/E_{reco} . The cuts on σ_E/E_{reco} to separate these samples are listed in table 7.4. All events passed a containment and basic ν_μ CC event selector cuts.

does not suffer from any *major* bias due to the shape of the training sample spectra. The best-fit value (indicated by the dots in figures 7.20, 7.21, and 7.22) chosen by the bin with the minimum χ^2 in each of the three oscillation assumptions, is always very close to the true value. Second, improvements in the contours exist for both of the assumed exposures. Third, using 3 or 5 bins for σ_E/E_{reco} both improve the contours over using 1 bin, but they appear to have roughly the same performance. I expect that this is because what the 5 bin case gains by having better energy resolution in the first bin, is lost by statistics when




Figure 7.20: Example sensitivity contours for the true “minimal” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the 1 σ_E/E_{reco} bin (black), 3 σ_E/E_{reco} bins (blue), and 5 σ_E/E_{reco} bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.

spreading the events out over more bins. Since the performance of the 3 bin and 5 bin cases appears to be roughly the same, I have chosen to use 3 bins for σ_E/E_{reco} with the cuts outlined in table 7.4 to be on the safe side considering the smaller number of events




Figure 7.21: Example sensitivity contours for the true “maximal” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the $1 \sigma_E/E_{reco}$ bin (black), $3 \sigma_E/E_{reco}$ bins (blue), and $5 \sigma_E/E_{reco}$ bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.

expected in the first analysis.




Figure 7.22: Example sensitivity contours generated for the true “alternate” oscillation case. The top plot assumes an exposure of 18×10^{20} POT and the bottom plot assumes 2×10^{20} POT. All plots show the 90% confidence intervals for the $1 \sigma_{E/E_{reco}}$ bin (black), $3 \sigma_{E/E_{reco}}$ bins (blue), and $5 \sigma_{E/E_{reco}}$ bins (red) cases. The green star shows the true oscillation parameters used for the MC ensemble.

7.4.3 Exploring a Set of Toy Experiments

As discussed in this chapter, one of the primary driving factors behind developing an analysis that uses the estimated event energy resolutions to improve our oscillation parameter fits is

the idea that the events near 1.5 GeV are extremely important for distinguishing the minimal and maximal oscillation cases. Seeing a few events in this region with energy resolutions that are estimated to be good could be either lucky if the energy resolution estimates are accurate and the true oscillation parameters are the minimal ones, or extremely unlucky if the energy resolution estimates are not accurate and the true oscillation parameters are the maximal ones. The contours in section 7.4.2 were generated with a large sample of simulated events scaled down to our expected exposures. Therefore, any instances of “(un)lucky” events, such as we might expect in the real data, have been smoothed out. So as a last test of the method of fitting a 2D reconstructed energy and energy resolution spectra, I generated a large sample of “toy” experiments, to explore the effects of “(un)lucky” events.

To generate my toy experiments, I took the 2D spectrum from my “fake” data sample (the middle plot in figure 7.19) using the assumption of maximal oscillations and scaled it down to an exposure of 2×10^{20} POT. A 2D energy spectrum for each toy experiment was then generated using Poisson-distributed random numbers with a mean drawn from the bin contents of the scaled down spectrum. I generated sensitivity contours for each of these toy experiments for the $1 \sigma_E/E_{reco}$ bin and $3 \sigma_E/E_{reco}$ bin cases. A sample of these sensitivity contours are shown in figure 7.23.

I used two metrics as a way of gauging the overall performance of the oscillation parameter fits across these experiments. The first was “ $\Delta \sin^2 2\theta_{23}$ Errors” which was the magnitude of the difference between the left edge of the 90% confident interval and the best fit value in $\sin^2 2\theta_{23}$ space for the 1 bin case, minus the same for the 3 bin case. This is essentially a measurement of how wide the contours are, with positive values corresponding to tighter contours using the 3 bin fit. The second metric was “ $\Delta \sin^2 2\theta_{23}$ True” which was the magnitude of the difference between the best fit value and the true value in $\sin^2 2\theta_{23}$ space for the 1 bin case, minus the same for the 3 bin case. This is a measurement of how

close the best fit was to the truth, with positive values corresponding to better performance with the 3 bin fit. Plots for these two metrics for 1000 toy experiments are shown in figure 7.24.

Next I repeated this process of generating 1000 toy experiments using the minimal oscillation parameters to oscillate the true spectrum of my “fake” data set. A sample of the contours and the corresponding metric plots are shown in figures 7.25 and 7.26. There are a couple of things to be learned from this. First, for both metrics in both the maximal and minimal oscillation cases, there are experiments that fell on both sides of zero. This implies (as expected) that there are both “lucky” and “unlucky” events that can throw off the best fits. But, with the exception of the “ $\Delta \sin^2 2\theta_{23}$ True” metric for the minimal oscillation case, the means of these plots are all positive with visible distribution skews towards the positive side. Second, it appears that the effect of getting lucky is more pronounced in the maximal oscillation case. However, I would be remiss if I didn’t point out that the true best fit for the maximal oscillation case is right up against the border of my oscillation parameter space. It might be possible that the 3 bin fit prefers solutions in the unphysical region, which could lead to the false conclusion of better performance than with the 1 bin fit. We must also be cautious in interpreting these results since the RMS of each plot is bigger than the mean. In general though, I take this to mean that for both the maximal and minimal cases, we on average get lucky more often than getting unlucky, which of course agrees with what we would expect given the contours discussed in section 7.4.2.




Figure 7.23: Examples of the 90% sensitivity contours from generating 1000 toy experiments assuming true maximal oscillation parameters. For each plot, the black curve/dot represents the $1 \sigma_E/E_{reco}$ bin fit, and the red curve/dot represents the $3 \sigma_E/E_{reco}$ bin fit. The blue dot is the true oscillation parameters.




Figure 7.24: The two metrics “ $\Delta \sin^2 2\theta_{23}$ Errors” and “ $\Delta \sin^2 2\theta_{23}$ True” for the 1000 toy experiments generated assuming true “maximal” oscillation parameters.




Figure 7.25: Examples of the 90% sensitivity contours from generating 1000 toy experiments assuming true minimal oscillation parameters. For each plot, the black curve/dot represents the $1 \sigma_{E/E_{reco}}$ bin fit, and the red curve/dot represents the $3 \sigma_{E/E_{reco}}$ bin fit. The blue dot is the true oscillation parameters.

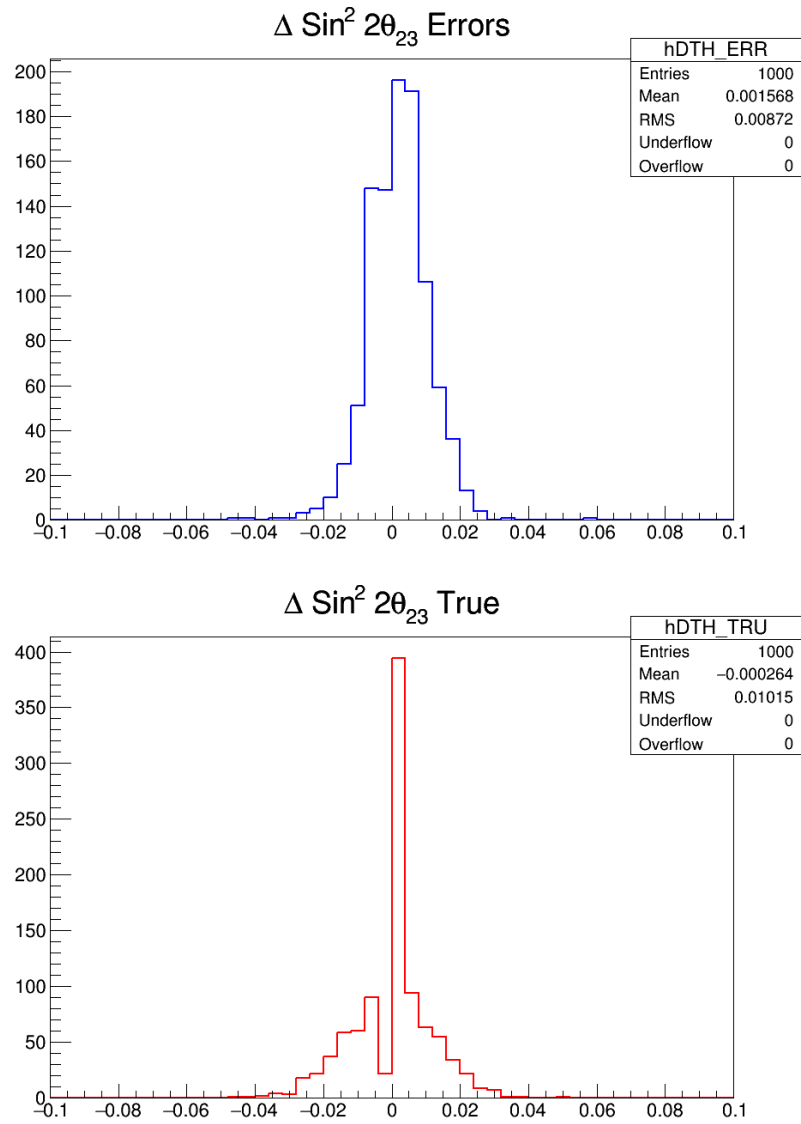


Figure 7.26: The two metrics “ $\Delta \sin^2 2\theta_{23}$ Errors” and “ $\Delta \sin^2 2\theta_{23}$ True” for the 1000 toy experiments generated assuming true “minimal” oscillation parameters.

CHAPTER 8

Results of The ν_μ Charged-Current Disappearance Analysis Using the Break Point Fitter Method

This chapter presents the final steps in producing a measurement of $|\Delta m_{32}^2|$ and $\sin^2 \theta_{23}$ through the ν_μ CC disappearance channel, including details specific to the Break Point Fitter (BPF) analysis. It includes a summary of the algorithms used to select a contained sample of ν_μ CC events for both detectors, the rejection of NC and cosmic-ray background events, a summary of the final data sets used for this analysis, and a description of the analysis framework used to produce the final results. At the end of this chapter, these final results are presented including sensitivity contours and the best fits for the oscillation parameters $|\Delta m_{32}^2|$ and $\sin^2 \theta_{23}$.

8.1 Event Selection

The official NOvA ν_μ CC disappearance analysis has already provided a means for selecting ν_μ CC signal events while rejecting backgrounds from NC events, ν_e CC events, and cosmic rays. I have opted to use this same event selection criteria for my analysis which I describe in the sections that follow. The far and near detector analysis sets are chosen by proceeding through three levels of cuts. The first level includes the basic data quality selection described in section 4.3. The second level includes selecting activity that is within the NuMI beam spill window and within the fiducial volume of the detector. Lastly, the third level selects

ν_μ CC candidate events and rejects cosmic rays that pass through the second level. The cuts and algorithms associated with the second and third levels are described below.

8.1.1 Contained Event Selection

As described in chapter 3, a NuMI beam triggered event is a 550 μsec long readout of detector activity, with the 10 μsec wide beam spill starting roughly 218 μsec after the trigger start time, T_0 . An event is required to be between $T_0 + 217 \mu\text{sec}$ and $T_0 + 229 \mu\text{sec}$ to be considered a NuMI beam event. Through an automated search for neutrino events in the far detector data, it was discovered that at least one strong neutrino candidate event occurred 64 μsec after the NuMI beam spill window [82]. This was traced to a failure mode that occurs during a reboot of the far detector timing distribution units (TDUs) which causes them to initialize their clocks offset by exactly 64 μsec . An extensive check of the log files generated by the timing system revealed strong reasons to be fully confident that this problem did not affect any data taken during and after October 2014, following the installation of a timing system calibration device. For the data taken before this time, there was no way of confirming if or when this problem occurred. It was decided that it was better to accept more background events than to reduce the already limited number of far detector signal events. Consequently, a second timing window was opened between $T_0 + 281 \mu\text{sec}$ and $T_0 + 293 \mu\text{sec}$ only for the far detector data taken before October 2014.

The next step is to choose events that are well contained within the detectors. Selecting contained events improves the overall event energy resolution by removing events with exiting tracks for which the estimate of the event energy will be poor. This selection also helps remove cosmic-ray events in the far detector and neutrino interactions that occur in the surrounding rock for the near detector. Decisions about containment are based off of variables from three reconstruction modules, Slicer4D (described in section 6.3), the cosmic

tracking algorithm (described in 4.2.1), and the tracking algorithm used for the standard NOvA ν_μ CC analysis called “Kalman tracker.” This algorithm was designed to accurately track non-showering particles such as muons, using a method based on a Kalman filter [83].

The NOvA Kalman tracker works by propagating tracks in a slice starting with the cell hits with the largest z coordinates, stepping backwards plane by plane using the current estimate of the track position and slope to estimate the location of expected track hits in the next plane. A probability score is computed for hits in the next plane and hits with good scores are included in the track. When a hit is added to a track, the track fit is corrected to include the measured information on position and slope. For muons with momenta greater than 0.6 GeV/ c , the average track completeness (defined by equation 6.5) is between 90 and 95%, and the reconstructed track length resolution is around 1%. More details on the NOvA Kalman tracker can be found in [84].

The containment cuts for the far detector were selected to optimize both the reconstructed neutrino energy resolution and the final sensitivity contours [85]. A few general cuts are applied which require the number of slice hits to be > 20 , the number of continuous planes (adjacent planes with cell hits) in the slice to be > 4 , and the number of 3D cosmic tracks to be > 0 , which ensures that reliable reconstruction is possible. The first two cuts remove very small slices with little or no useful information and very vertical slices that are much more likely to be cosmic rays. The third cut ensures that the cosmic tracking algorithm did not fail, which can happen for slices of extremely low quality. The next cuts require slices to have no hits in the two outermost cells of a plane and no hits in either the first two or the last two detector planes. A cut is also placed on the forward (backward) projected distances from the Kalman and cosmic track end (start) locations to the nearest detector edge. These distances are quantified in terms of the number of cells that these projections pass through, using the single track produced by the cosmic tracker and the

Kalman track selected by the ReMId algorithm (described in section 8.1.2.) The forward and backward projected distances for the ReMId selected Kalman track are both required to be > 10 cells. For the cosmic track, these distances must be > 0 , which simply requires the cosmic track to start and stop inside the detector. All distances measured with respect to the detector edges take into account the varying size of the far detector during early data taking which is described below in section 8.2.1.

For near detector containment, the same general cuts requiring slices with > 20 hits, > 4 continuous planes, and at least one cosmic track are applied. The basic cuts on slice containment are also the same, which require no hits in the two outermost cells and in the first and last two planes. The projected forward and backward distances for the ReMId selected Kalman track must be > 4 and 8 cells respectively. No cuts are applied to the projected distances for cosmic tracks. In addition to these cuts, several tighter requirements are placed on events that may interact with the muon catcher. The z coordinate of the start of the ReMId selected Kalman track is required to be at least 125 cm from the upstream end of the muon catcher. For tracks in the top $1/3^{\text{rd}}$ of the detector, the track must end prior to the start of the “air gap” above the muon catcher (see figure 3.8.) Lastly, there must be less than 30 MeV of “hadronic energy” (calorimetric energy not on the ReMId selected Kalman track) for cell hits in the muon catcher. More details on the near detector containment cuts can be found in [86].

8.1.2 Signal Event Selection

Identifying ν_μ CC interactions in both detectors is handled by the **Reconstructed Muon Identification** (ReMId) algorithm. This algorithm is very similar to the BPF muon identification algorithm described in section 7.2. It uses a k -nearest-neighbor classifier to compute a “muon score” for each Kalman track produced. Under the premise that a muon is the

characteristic signal of a ν_μ CC event, an event ID is assigned to each slice using the highest muon score from all of the Kalman tracks associated with that slice. The input variables to the ReMId algorithm are computed from the reconstructed Kalman tracks, and include a dE/dx log-likelihood value, a multiple scattering log-likelihood value, the track length, and the ratio of how much of the track overlapped with hadronic activity as computed by the `TrackCleanUpAlg` algorithm (described in section 7.3.) The distributions of slice event IDs for signal (true ν_μ CC events) and background (true NC and ν_e CC events) from simulated events in both detectors is shown in figure 8.1. The ν_μ CC event sample is selected by requiring a slice to have an event ID > 0.75 . This cut produces a signal selection efficiency of 82% (69%) and a sample purity of 93% (98%) for the far (near) detector. Further details on the ReMId algorithm can be found in [87, 88].

8.1.3 Rejection of Cosmic Rays in the Far Detector

Since the far detector sits on the surface, it sees a very high flux of cosmic-ray muons, approximately one every 10 μsec (the size of the NuMI beam spill window.) Given a 24 hour period of NuMI beam running, this equates to 65,000 cosmic-ray muons within the NuMI beam spill windows per day. For the full 14 kiloton detector, the expected number of contained ν_μ CC events per day is on the order of a few. In order to achieve a signal to background ratio of better than 10:1, cosmic rays must be rejected with an efficiency of greater than 99.99999% [89]. The containment cuts described above are expected to remove roughly 99% of cosmic rays in the far detector, therefore additional measures must be taken to reject as many of the remaining cosmic rays as possible.

The far detector cosmic rejection algorithm uses a multivariate classification method called a boosted decision tree (BDT) which was developed using the TMVA [79] package in ROOT [31, 32]. To maximize the performance, the BDT was trained using only signal

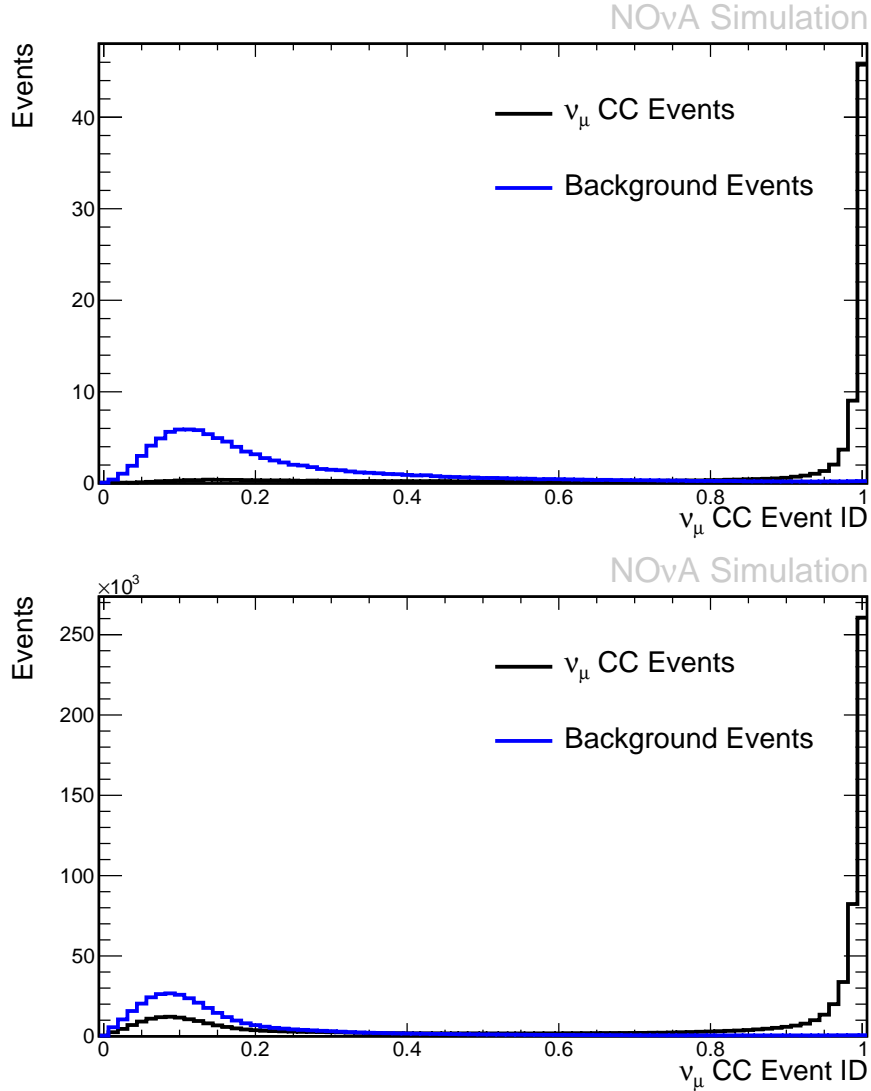


Figure 8.1: Distribution of slice event IDs computed by the ReMId algorithm for GENIE simulated events in the far (top plot) and near (bottom plot) detectors. For the far detector, the ν_μ CC events were oscillated using $|\Delta m_{32}^2| = 2.4 \times 10^{-3} [\text{eV}^2]$ and $\sin^2 \theta_{23} = 0.5$. Both distributions were normalized to the expected far and near exposures for this analysis of 3.52×10^{20} and 1.65×10^{20} POT.

(true ν_μ CC events) and background (simulated cosmic-ray muons) samples that pass the containment and ReMId selection cuts. There were 11 different input variables to the BDT which included quantities such as the y component of the direction vector for the ReMId selected Kalman track, the length, number of hits, and cosine of the angle with respect to the NuMI beam for this track ($\cos \theta_{Kal}$), and the maximum and minimum y positions of

the cell hits in the slice. The trained BDT returns a number classifying how “signal-like” an event is, based on the values for these input variables.

The choice of a cut value to use on the BDT output (referred to the “ ν_μ contained ID”) was based on defining a figure of merit $FOM \equiv S/\sqrt{B}$ where S is the number of signal events and B is the number of background events. A plot of this FOM for different cuts on the ReMId event ID value is shown in figure 8.2. The cut was chosen to be ν_μ contained ID > 0.535 , within the plateau of maximum FOM . In addition to this cut, the full far detector cosmic rejection algorithm requires the number of hits in the slice to be < 400 and $\cos\theta_{Kal}$ to be > 0.5 . More information on the design and implementation of the far detector cosmic rejection algorithm can be found in [89].




Figure 8.2: Figure of merit, defined as S/\sqrt{B} , as a function of the cut value applied to the cosmic rejection BDT output, for different cuts on the ReMId event ID value.

8.1.4 Break Point Fitter Quality Cut

In addition to the cuts described above, I imposed a “BPF quality” cut to ensure that the results of my energy estimator were sensible. This was necessary because I discovered some

extremely rare cases in which the Kalman tracker succeeded but the BPF tracker failed. These cases were typically caused by either the Elastic Arms or Fuzzy-K reconstruction algorithms producing odd results, such as vertices placed several meters away from the bulk of the slice, that prevented BPF from making any tracks. The BPF quality cut simply requires the BPF energy estimator to produce a value ≥ 0 and the best BPF muon ID value for the slice to be ≥ 0 . Applying this cut rejected an additional 0.034% of the near detector data and near detector simulation, and 0.008% of the far detector simulation. None of the ν_μ CC selected far detector data was rejected by this cut.

8.1.5 Event Selection Summary

Given the small size of the near detector, the vast majority of the unoscillated ν_μ beam events are not contained. After applying the containment and ReMIId selection cuts listed above, roughly 1.7% of the total number of slices remain. This is in agreement to within 2% of the prediction made from simulated near detector data. The simulations have shown that of the remaining events, 97.5% are ν_μ CC and 2.5% are NC, with negligible contamination ($< 0.1\%$) from rock events. The out of time data predicts that the contamination from cosmic rays in the near detector data is approximately $5 \times 10^{-5}\%$ [86].

For the far detector, the ν_μ CC signal and cosmic ray efficiencies were determined from the simulated ν_μ CC event sample (selecting only truly contained events and oscillating the ν_μ CC events using $|\Delta m_{32}^2| = 2.4 \times 10^{-3} \text{ [eV}^2\text{]}$ and $\sin^2 \theta_{23} = 0.5$), and out of time NuMI beam data respectively. For the signal events, the fractions of the original sample that survives the containment, ReMIId selection, and cosmic rejections cuts are 78%, 64%, and 41%. For the cosmic rays, these numbers are 1.0%, 0.5%, and $9 \times 10^{-6}\%$. The final sample of selected far detector events is expected to be 91.0% ν_μ CC events, 4.1% cosmic rays, and 4.9% NC events. A separate study was done to estimate the number of ν_τ CC appearance

events in the far detector event sample. The results of this study predicted 0.04 events per 10^{20} POT, which equates to a negligible contribution of < 1 event in 3 years of nominal running and ≈ 0.1 events for the data used in this analysis [90].

8.2 Summary of the Selected Data Sets

8.2.1 Far Detector Data

The far detector data taking was divided into three periods. The first period began on February 6th, 2014 and ended on September 5th, 2014 during the start of a scheduled accelerator shutdown. The second period goes from October 24th, 2014 to March 14th, 2015, and the third period goes from March 14th, 2015 to May 15th, 2015. Excluding subruns and individual data events labeled as “bad” using the data quality cuts described in section 4.3, the total accumulated POT from these three periods is 3.52×10^{20} [91].

During the first two periods, the far detector was still under construction with only part of it read out into a data stream to be used for analyses. As diblocks were constructed, they were added to this data stream one at a time after passing a series of commissioning checks. Figure 8.3 shows the growing detector as a function of time scaled to an approximate fiducial mass of 11.4 kilotons. All of the containment cuts discussed in section 8.1.1 are adjusted for the individual detector configurations of each subrun using a software package designed for this purpose described in [92]. Scaling the POT accumulated during each subrun by the fraction of the full detector mass, the equivalent POT for a 14 kiloton detector is 2.8×10^{20} or roughly 45% of a nominal year [91].

8.2.2 Near Detector Data

The near detector data set consists of data taken during the first two periods described in section 8.2.1. Unlike the far detector, the near detector was fully built and commissioned




Figure 8.3: Far detector mass and integrated exposure (mass \times POT) as a function of time. Periods of running with different configurations are evident as is the scheduled accelerator shutdown September/October 2014. This is scaled to an approximate fiducial mass of 11.4 kilotons [91].

prior to the start of the first data taking period, so no adjustments due to the growing size of the detector were required. The total number of POT seen by the near detector during this time was 1.65×10^{20} . After the ν_μ CC selections cuts described in section 8.1 are applied to the near detector data, roughly 5.1×10^5 events remain.

8.2.3 Simulated Data Sets

Chapter 5 contains the details of the NOvA simulation processes. Both the near and far detector simulated data sets were done with “run-matched” channel masks, taken from the real data. These masks hide channels labeled as bad according to the nearline output files (see section 4.2) from all of the downstream reconstruction modules. Naturally, none of the channels in the simulated data are bad, but this was done to make the simulations as close to the real data as possible. These masks were especially important for the far detector, where they ensure that the simulations reflect the changing size of the detector over time. The near detector simulations consisted of an accumulated total of 9.92×10^{20} POT, or roughly 6 times the amount of near detector data. The far detector simulations consisted of 5.76×10^{23} POT (almost 1000 years of NuMI beam!) which is roughly 1600 times the amount of data.

8.3 Data Simulation Comparisons

8.3.1 Results from the Kalman Tracker

The near detector data was used to make extensive comparisons to the simulations. Shown in figure 8.4 are the distributions of track angle with respect to the x and y detector axes and the x and y coordinates of the track start and stop for the track produced by the Kalman tracking algorithm with the highest ReMId value in each slice.

Figure 8.5 shows the distributions of track length and number of hits on the track in the slice with the highest ReMId value, for all slices in the selected ν_μ CC event sample in near detector data and simulation. The agreement in these variables is within 10% and is typically within 5% [85]. Also shown in this figure are the distributions of ν_μ CC event IDs (best ReMId value from all tracks in the slice) for contained slices. There is some disagreement outside of the signal region (for values less than 0.75) but the signal region in the data is within 10% of the simulation. The agreement for all of the plots shown in figure 8.5 indicates that the out-going muon in ν_μ CC events is modeled well.

There are however, significant differences in the reconstructed hadronic energies for the selected ν_μ CC event sample. For hadronic hits (defined as hits not on the primary muon track) there are fewer of them with on average less visible energy (energy deposited in the scintillator) in data than in the simulations. These distributions are shown in figure 8.6. Also shown in this figure are the distributions of visible on and off track hadronic energy for each slice. The off-track energy is simply the summed calorimetric energy for all hits not on the primary muon track. The on-track energy is the hadronic energy contamination on the primary muon track computed by the `TrackCleanUpAlg` algorithm described in section 7.3.

The effect of this hadronic energy difference is perhaps most striking in a plot of reconstructed neutrino energies. Shown in figure 8.7 are the reconstructed energy spectra for the




Figure 8.4: Distributions of track variables from near detector data and simulations for the track in the slice with the highest ReMID value. Only tracks from slices that passed the near detector containment cuts are shown. Plots were taken from [85].




Figure 8.5: Top Left: Track length distributions for the Kalman track with the highest ReMid value for near detector data and simulation. **Top Right:** Distribution of the number of hits on the track with the highest ReMid value in near detector data and simulation. Only tracks from slices that passed the near detector ν_μ CC event selection cuts are shown. Plots were taken from [93]. **Bottom:** Distribution of ReMid values for all contained slices in near detector data and simulation. Plot taken from [85].

selected ν_μ CC event samples in data and simulation. The energy estimator used to create these spectra is the one being used for the standard NOvA ν_μ CC disappearance analysis, and is described further in [94]. This estimator takes the neutrino energy to be the sum of the muon energy and the remaining hadronic energy, where the muon energy is determined from the track length and the hadronic energy is scaled from the visible energy to the true energy from the simulations using a spline fit. The overall effect of the hadronic energy difference is a shift to the left for the data in the neutrino energy peak of about 100-150




Figure 8.6: Distributions of variables associated with hadronic energy. Hadronic hits are defined as hits in a slice not on the primary muon track. Only tracks from slices that passed the near detector containment and ν_μ CC event selection cuts are shown. Plots were taken from [85].

MeV. I will discuss further the possible origins of this hadronic energy difference between the data and the simulations and how it will be treated in the ν_μ CC disappearance analysis in section 8.4.

8.3.2 Results from the Break Point Fitter Tracker

Figure 8.8 shows the distributions of the six input variables to the BPF energy estimator for near detector data and simulation, with zoomed versions of some of the plots in figure 8.9. The distributions of P_μ values are very similar, but do show some slight shape differences

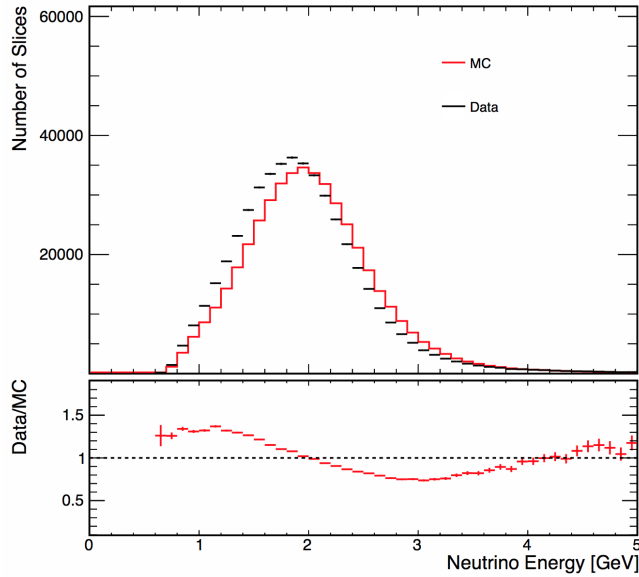


Figure 8.7: Distributions of reconstructed neutrino energies for slices passing the ν_μ CC event selection cuts in near detector data and simulation. Plot was taken from [93].

between 1 and 1.5 GeV/ c not immediately apparent in the Kalman track distributions shown in figure 8.5. However, the differences in this region are less than 10% and the agreement across the rest of the distribution is good. The N_{3D} variable shows significant differences between data and simulations. Particularly, the data shows a preference for events with fewer 3D prongs produced by the Fuzzy-K algorithm, which is in agreement with the general idea that there is less hadronic energy in the data over the simulations.

The zoomed plots shown in figure 8.9 indicate that the BPF energy estimator sees the same hadronic energy differences discussed in section 8.3.1. The E_{3D} and Sum_{PE} variables both have shapes similar to the plots of number of hadronic hits and visible off-track energy shown in figure 8.6. In an attempt to make a more direct comparison, I defined BPF “ E_{had} ” as the sum of the E_{3D} and E_{remain} variables. The distribution of values for this variable in both near detector data and simulation is shown in figure 8.10, and closely resembles the distribution of the number of hadronic hits. Also shown in this figure is the BPF reconstructed neutrino energy distribution showing the same shift of 100-150 MeV in data




Figure 8.8: Distributions of the six input variables to the BPF energy estimator in near detector data and simulations.

seen in figure 8.7, and the distribution of BPF muon ID values which like the ReMid values, indicates that muons are being modeled well.

8.4 A Discussion on Hadronic Energy

Comparing the near detector data to the simulations shows that muons in the selected ν_μ CC event sample are being modeled well but that a significant difference exists for reconstructed hadronic energy (see section 8.3.) Many possible causes for this discrepancy were considered including calibration, beam modeling, and simulations of the detector response, but none of them were able to explain the observed differences [85, 93]. Additional studies were




Figure 8.9: Zoomed versions of three of the input variables to the BPF energy estimator showing comparisons between data and simulations. The top two plots are both distributions for the variable E_{3D} .

performed looking at the differences between the data and the simulations at the slice level [95] and looking at the differences between two different neutrino interaction generators: GENIE and NuWro [96]. Both studies hinted at the idea that GENIE has a tendency to produce a few too many neutrons. Because no reasonably sized systematic can produce the hadronic energy discrepancy and because there are known hadronic modeling differences between different neutrino interaction simulators, it is assumed that the disagreements in the near detector data and simulations are products of the neutrino interaction model.

A decision was made for the standard NOvA ν_μ CC analysis to apply a simple scaling to the hadronic energy component of each event [85,93]. This analysis uses an energy estimator (described briefly in section 8.3.1) which takes the reconstructed neutrino energy to be

$$E_\nu = E_\mu + E_{had}, \quad (8.1)$$

where E_μ and E_{had} are the reconstructed muon and hadronic energy components of the




Figure 8.10: Top: Energy spectra generated by the BPF energy estimator for near detector data and simulation. **Bottom Left:** BPF “hadronic energy” taken as the sum of the E_{3D} and E_{remain} variables. **Bottom Right:** Distribution of BPF muon ID values for the BPF track in each slice with the highest BPF muon ID value. All plots include only the slices that pass the ν_μ CC selection criteria.

event. For each event in data, the value of E_{had} is scaled up by 21% leaving the value of E_μ unchanged. The decision to scale the data to the simulation was based on the assumption that we know where the peak in the energy spectrum should be since this is well understood from two-body pion decay through equations 3.1 and 3.2. The simulation lines up with this expected peak by construction, but the data does not since the energy reconstruction was tuned with the simulation. Shifting the data so that the reconstructed energy spectra lines up with the expected peak accounts for the possible miss-modeling in the simulations. This shift is applied to both the near and far detector data under the assumption that an effect from miss-modeling physics will be present in both detectors. Figure 8.11 shows the near detector hadronic energy spectra before and after shifting and the final reconstructed neutrino energy spectra as compared to the simulations. Note that the agreement in the

E_{had} variable is better but still shows some shape differences, and that the reconstructed neutrino energy spectra agrees quite well.




Figure 8.11: Top: Distributions of the E_{had} variable in data and simulations, showing the effect of scaling this variable up by 21% in the data. **Bottom:** Showing the reconstructed neutrino energies for data events with E_{had} scaled up by 21% as compared to the unaltered simulated spectra. All plots include only the slices that pass the ν_μ CC selection criteria.

Since the exact physical reasons behind the need for this correction are not yet well understood, a conservative estimate of 100% is used for the absolute uncertainty on the hadronic energy [85, 93]. However since the shape of the energy spectra are expected to be very different in both detectors, a relative uncertainty on the hadronic energy is also applied. To compute this relative uncertainty, a more complicated three parameter shift was applied to fit the simulation to the near detector data. This fit broke the simulated events into quasi-elastic (QE), resonant (RES), and deep-inelastic scattering (DIS) by truth, and allowed the

hadronic energy component of the RES and DIS events and the overall normalization of the DIS events to vary independently. The results of this three parameter fit show an excellent agreement between the data and the simulation for both the total reconstructed event energy and just the hadronic energy component [97,98]. Taking this shift as a “true” answer, the relative uncertainty was computed using the double ratio

$$R = 1 + \left(\frac{1 + \Delta X_1^{FD}}{1 + \Delta X_2^{FD}} \right) \left(\frac{1 + \Delta X_2^{ND}}{1 + \Delta X_1^{ND}} \right), \quad (8.2)$$

where $\Delta X_{1,2}^{FD,ND}$ is the percent change in either E_{had} or the overall normalization for the far (FD) and near (ND) detectors using method 1 (the simple 21% shift) and method 2 (the three parameter fit.) Using this method yields a relative uncertainty of 2% on the E_{had} scale and 1% for the normalization [93].

Since the BPF energy estimator uses a k-nearest-neighbor algorithm to produce a reconstructed neutrino energy from the weighted average of the neighbors in the input variable phase space, applying the shift described above is not a viable option. So to apply a similar shift to my energy spectra, I computed for each event the ratio of the shifted energy (from equation 8.1) to the unshifted energy, and then I multiplied the BPF reconstructed neutrino energy by this ratio. The near detector reconstructed energy spectra with this shift applied is shown in figure 8.12. Like the spectra shown in figure 8.11, the peak of the shifted BPF reconstructed energy spectra for the near detector data lines up with the expected value in the simulation, but the overall shape agreement is not quite as good. This is to be expected since the 21% hadronic energy correction was tuned to fit the energy spectrum reconstructed with equation 8.1. Because I am using a correction to the total energy based on the same 21% hadronic energy shift, I will assume the same energy systematic uncertainties described above.




Figure 8.12: Showing the reconstructed neutrino energies with the BPF energy estimator for data events scaled by the ratio of the shifted to the unshifted reconstructed energy computed by equation 8.1, as compared to the unaltered simulated spectrum. Only slices that pass the ν_μ CC selection criteria have been included.

8.5 The NOvA Analysis Framework

Most NOvA analyses are done using a set of stripped down files that contain a minimal amount of event information called the Common Analysis Format (CAF) files. These files store the relevant data from the results of the reconstruction algorithms including information about tracks, particle identification algorithm outputs, and neutrino energy estimates. The analysis tools for using these files have been gathered into one place to allow an analysis to be put together with relative ease and to prevent many unnecessary wheel reinventions. The framework that provides the tools such as oscillation calculators, near to far extrapolations, and sensitivity contour generators using the CAF files is called the CAFAna framework [99]. For simplicity, I used this framework to generate the results of my analysis.

8.5.1 Background Estimates

For the NOvA ν_μ CC disappearance analysis, the background events come primarily from NC events in the near detector and NC and cosmic rays in the far detector, with very small

contributions from ν_e and ν_τ CC events (see the numbers listed in section 8.1.5.) Since these backgrounds are expected to be small for both detectors, the simulation is used to subtract the background spectra from the contained, ReMId selected ν_μ CC event spectrum in the near detector [85]. This spectrum is then extrapolated to the far detector (described in section 8.5.2.) The predicted far detector spectrum is formed by adding the beam background spectra (NC and ν_e , assuming the ν_τ contribution to be negligible) determined from the simulation and the cosmic background spectrum determined from the data, to the extrapolated far detector spectrum. The CAFAna framework handles the appropriate subtraction and reapplication of these background spectra from the simulations. The far detector data-derived cosmic background spectrum is added manually to the predicted ν_μ CC spectrum. Since I am using the same event selection as the official ν_μ CC disappearance analysis, I have applied this strategy as well.

Figure 8.13 shows the expected background contributions to the near detector reconstructed neutrino energy spectrum determined from the simulation. The total background is small and is almost entirely NC events. The 2D spectra shows that the vast majority of the background events are recognized as having poor energy resolution and are placed into the “worst” energy resolution bin.

The expected NC and cosmic ray background spectra are shown in figure 8.14. The cosmic backgrounds were estimated by using the out of time NuMI beam data, selecting events from roughly 400 μsec of time that comfortably avoided the NuMI beam spill windows. The estimates for the cosmic-ray backgrounds were then scaled down by the ratio of the size of the beam windows to the background selection window, appropriately accounting for both windows open for beam events prior to October 2014 (see the discussion in section 8.1.1.) As with the near detector NC background, the 2D energy spectrum shows that the majority of the cosmic ray data events are placed in the “worst” energy resolution bin. There were




Figure 8.13: Top: Reconstructed neutrino energy spectra from the near detector simulations for all ν_μ CC selected events (black) and background events (blue) in the ν_μ CC sample as determined by truth. **Bottom:** Two-dimensional energy vs. energy resolution plot for only the background events in the ν_μ CC sample, showing that most background events are put into the “worst” energy resolution bin.

51 events selected from the out of time data, which scales down to an expected 1.98 cosmic ray events in the ν_μ CC analysis sample. Examples of two of these events are shown in figure 8.15.

8.5.2 Extrapolation

Extrapolating a near detector spectrum to make a prediction for a far detector spectrum can be a tricky business. Fortunately, the extrapolation machinery is included in the CAFAna framework as one of the standard analysis tools. The basic premise is that a bin by bin (in reconstructed energy) extrapolation is performed from the ratio of the near detector data




Figure 8.14: Top: Reconstructed neutrino energy spectra for the selected ν_μ CC events in the far detector. Black and red represent the expected true ν_μ CC and NC events determined from the simulations, and blue represents the expected cosmic ray spectrum determined from out of time data. The true ν_μ CC events were oscillated using $|\Delta m_{32}^2| = 2.4 \times 10^{-3}$ [eV²] and $\sin^2 \theta_{23} = 0.5$. **Bottom:** Unscaled two-dimensional energy vs. energy resolution plot for only the 51 background cosmic ray events determined from the out of time data.

spectrum to the near detector simulation spectrum where the simulation can be oscillated according to the true energies of the events in the reconstructed energy bin. This bin by bin ratio is scaled by the reconstructed far detector spectrum from the simulations to yield a predicted far detector spectra.

To perform this extrapolation, the data and simulation spectra are divided into bins of reconstructed energy, indexed by j . The simulation spectra have an additional dimension representing the true neutrino energy, indexed by i , that can be used to calculate the oscillation probabilities of the events populating the reconstructed energy bin j . The predicted




Figure 8.15: Examples of 2 of the 51 out of time events selected by the ν_μ CC event selection criteria from the far detector data. These events were used to estimate the energy spectrum of the cosmic ray background in the ν_μ CC analysis sample. Both events occur near the top of the detector and are likely caused by neutrons entering the detector from above.

value for the j^{th} reconstructed energy bin for the far detector spectra, $F^{\text{Pred}}(j)$, is given by the sum over the true energy bins applying the appropriate oscillation probabilities using, for example, equation 2.8. That is

$$F^{\text{Pred}}(j) = \sum_i F^{\text{Pred}}(i, j) \cdot P_{\nu_\mu \rightarrow \nu_\mu}(i), \quad (8.3)$$

with $F^{Pred}(i, j)$ given by

$$F^{Pred}(i, j) = \frac{N^{Pred}(i) \cdot F^{Sim}(i, j)}{N^{Sim}(i)}. \quad (8.4)$$

Here N refers to a near detector spectra and the superscript Sim indicates a spectra that came from the simulations. In each true energy bin i , the computed rate of near detector ν_μ CC events $N^{Pred}(i)$, is given by reweighting by the reconstructed energy bins using

$$N^{Pred}(i) = \sum_j \frac{N^{Data}(j) \cdot N^{Sim}(i, j)}{N^{Sim}(j)}. \quad (8.5)$$

By using different values for the oscillation parameters in equation 8.3, this extrapolation can produce a set of predicted far detector spectra that can be used for an oscillation parameter fit to the far detector data. Further details on the extrapolation machinery can be found in [100].

8.5.3 Systematic Uncertainties

The CAFAna framework provides the tools to easily handle systematic uncertainties. Given a set of expected uncertainties, the individual events can be either shifted or reweighted as they are read in from the CAF files to determine the effects on the shape of the reconstructed energy spectrum. Shifted spectra are filled for each of the systematics, which are then passed to the tools that construct the final contours so that each of the systematic uncertainties can be taken into account during the final fit.

For the ν_μ CC disappearance analysis, by far the dominant systematic comes from the hadronic energy shift described in section 8.4. For this analysis, the hadronic energy systematic was applied as a 6% smearing in $|\Delta m_{32}^2|$. All other systematics were fit as nuisance parameters. The full list of systematics accounted for in this analysis include the

following:

- **background NC:** As described in section 8.5.1, the NC contamination in both the near and far detector spectra is expected to be small and is taken directly from the simulation. Therefore, a conservative uncertainty of 100% on the number of NC events is used [85].
- **normalization:** The uncertainty on the detector mass and the beam flux were estimated to be 0.7% [101] and 0.5% [102] with the primary beam uncertainty coming from the hadron production models. They are added in quadrature for an overall uncertainty in normalization of 0.9%.
- **GENIE:** Many of the effects of nuclear modeling are expected to cancel in the extrapolation from near detector to far detector. However many parameters within GENIE's models can be varied, only 5 of which were shown to have any impact on the ν_μ CC energy spectrum [103]. These 5 parameters were the axial masses of the CC QE, CC RES, NC RES, and NC elastic, and the vector mass for the CC RES cross sections. The sizes of the uncertainties from varying these parameters are estimated for each bin in the reconstructed energy spectra, were shown to be small. Examples of the reconstructed energy spectra with error bars from the GENIE uncertainties can be seen in [103].
- **detector modeling:** Many uncertainties are folded into this category, all of which were shown to have only a small impact due to canceling effects in the near to far extrapolation [85]. The major effects include uncertainties in modeling with GEANT [104] and in the simulated light levels in the scintillator and thresholds in the electronics [105]. Together, these effects are covered with a 1% uncertainty.
- **absolute energy scale calibration:** The absolute energy scale uncertainty was estimated to be 5% by propagating the calibration systematics all the way through the reconstruction and analysis steps [106].

- **hadronic energy shift:** As described above, the absolute uncertainty on the 21% shift to E_{had} is taken to be 100%, with a 2% relative uncertainty on the E_{had} scale and 1% on the E_{had} normalization [85,93].

Many other possible effects on the reconstructed ν_μ energy spectrum were studied and shown to be negligible for this analysis [85]. These include:

- **detector alignment:** This quantifies the effect of measurements made to determine the alignment of the detector blocks [107].
- ν_τ **contamination:** The number of ν_τ events in the ν_μ CC event sample for this analysis is expected to be < 0.1 [90].
- **hadron transport modeling:** The beam uncertainties associated with tracking hadrons from the target to the decay pipe were shown to be much smaller than the uncertainties associated with hadron production (the overall beam flux), which is accounted for in the normalization systematic [102].
- **Birk’s suppression modeling:** The tuning of the Birk’s suppression parameters described in section 5.4 is not unique and introduces some uncertainty. However, since muon energies are estimated from track length, using the wrong Birk’s parameters will only affect contributions to the reconstructed neutrino energies from non-muon sources. Therefore, this was considered to be accounted for with the hadronic energy systematics [85].
- **time-dependent channel masks:** The masks applied to the channels known to be “bad” (as determined at the subrun level from the nearline metrics) have been shown to be incorrectly stored in the database tables for a small portion ($< 1\%$) of subruns [108,109]. This was studied by applying the wrong channel masks to a set of near and far detector data with little to no effect on the overall neutrino energy reconstruction [110].

8.5.4 Sensitivity Contours

The extrapolation framework described in section 8.5.2 produces a prediction for an unoscillated far detector energy spectrum. This prediction is oscillated by the true neutrino energies of the selected events in the simulated data sets, using different values for the oscillation parameters to make comparisons to the data. For a given set of oscillation parameters, a total χ^2 is computed using equation 7.10 from the sum over all of the bins in the reconstructed energy spectra, marginalizing over the list of systematic uncertainties. A sensitivity contour is constructed by computing a χ^2 for each predicted spectrum formed from a set of different oscillation parameters. Best fit values for any of the individual parameters are calculated by marginalizing over the other parameters and fitting the resulting curve to a χ^2 distribution. The sensitivity contours for the NOvA ν_μ CC disappearance analysis are generated in a two-dimensional $|\Delta m_{32}^2|$ and $\sin^2 \theta_{23}$ space.

8.6 Final Results

8.6.1 The Selected Data Events

The total number of far detector data events selected by the ν_μ CC selection criteria was 48, 4 of which were in the +64 μsec beam spill window from the data taken prior to October 2014. Of these events, 35 had reconstructed energies < 5 GeV and were used in the final oscillation parameter fit. Assuming $\sin^2 \theta_{23} = 0.5$ and $\sin^2 \theta_{23} = 0.4$ with $|\Delta m_{32}^2| = 2.4 \times 10^{-3}$ [eV²], the predicted numbers of far detector events for the ν_μ CC event sample are 31.9 and 39.6 respectively. With no oscillations, this predicted number is 196.

Examples of some of the selected events are shown in figures 8.16 through 8.20. Figures 8.16, 8.18, 8.19 appear to be very clean QE events, with a clear recoiling proton track visible in the event in figure 8.16. Figures 8.17 and 8.20 are both higher energy events with long muon tracks and well reconstructed vertices. From the entire sample of 48 events, 5 appear




Figure 8.16: Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 2.62$ GeV, $\sigma_E = 0.11$ GeV (4.2%), $P_\mu = 2.33$ GeV/c, and a BPF event ID = 0.98.

to have non-neutrino like topologies (by visual inspection) and are very likely to be cosmic rays. All 5 of these events have energies between 1 and 3 GeV and 3 of them occurred in

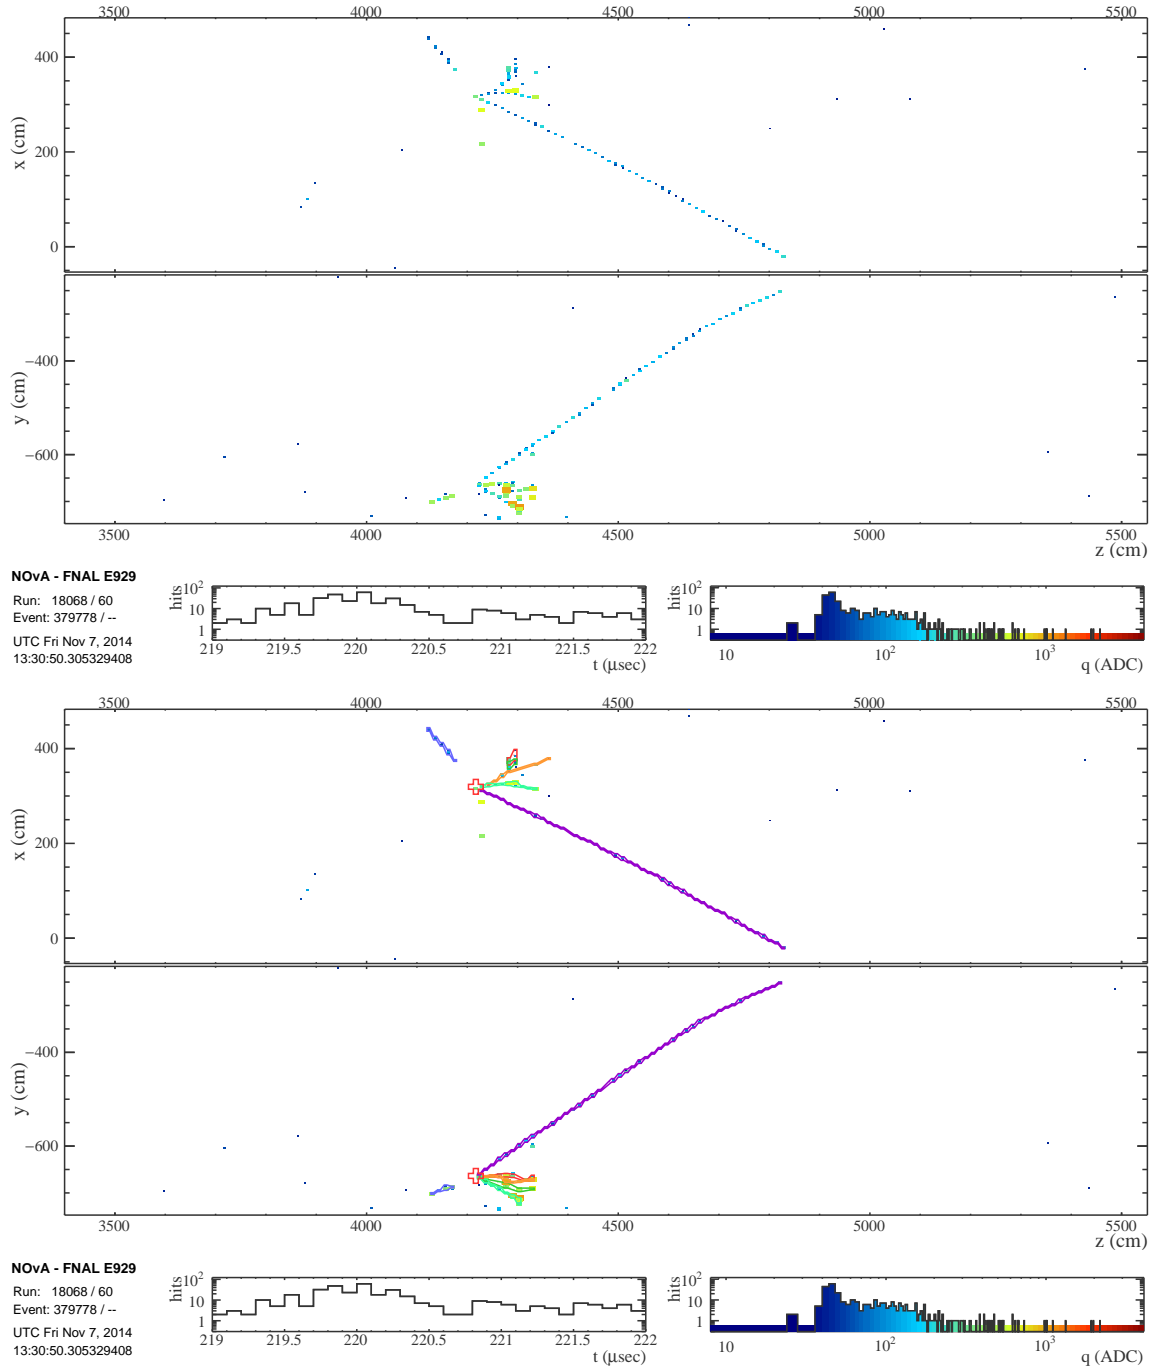


Figure 8.17: Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 4.99$ GeV, $\sigma_E = 0.4$ GeV (8.0%), $P_\mu = 1.83$ GeV/c, and a BPF event ID = 1.0.

the +64 μ sec beam spill window. Two examples of these 5 events are shown in figure 8.21.

Both examples are close to the top of the detector and appear to be entering from above.




Figure 8.18: Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 2.35$ GeV, $\sigma_E = 0.11$ GeV (4.7%), $P_\mu = 2.11$ GeV/c, and a BPF event ID = 1.0.

8.6.2 Oscillation Parameter Fits

The 2-dimensional BPF energy spectrum for the 35 selected far detector data events is shown in figure 8.22. There were 15 events in the first energy resolution bin, 12 in the second, and




Figure 8.19: Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 0.91$ GeV, $\sigma_E = 0.08$ GeV (8.7%), $P_\mu = 0.82$ GeV/c, and a BPF event ID = 0.98.

8 in the third. Performing the full near to far extrapolation with the analysis framework described in section 8.5 using the shifted near and far detector data spectra described in




Figure 8.20: Selected far detector data event from the ν_μ CC analysis sample with and without reconstruction (showing the Elastic Arms vertex and the 3D Fuzzy-K prongs.) The reconstructed parameters for this event are $E_\nu = 7.68$ GeV, $\sigma_E = 0.66$ GeV (8.6%), $P_\mu = 2.38$ GeV/c, and a BPF event ID = 0.99. This event, while well reconstructed, was not used in the final fit since $E_\nu > 5.0$ GeV.




Figure 8.21: Selected far detector data events from the ν_μ CC analysis sample likely to be cosmic rays.

section 8.4, produced the measurement contours shown in figure 8.23. Contours with and without the systematic uncertainties are drawn together in this figure. For comparison, the 90% confidence interval contours from this analysis are shown in figure 8.24 overlaid with

the most recent contours from T2K [24] and MINOS+ [25].




Figure 8.22: Top: The reconstructed energies for the 35 far detector events with the predicted spectrum in red from the near detector extrapolation using the best fit values for the oscillation parameters. **Bottom:** Two dimensional reconstructed energy spectra for the 35 selected far detector events.

For each of the two oscillation fit parameters, the other parameter was marginalized to generate the χ^2 distributions shown in figure 8.25. From these plots, the best fits for the oscillation parameters were determined to be $|\Delta m_{32}^2| = 2.49^{+0.19}_{-0.17} [\times 10^{-3} \text{eV}^2]$ and $\sin^2 \theta_{23} = 0.51 \pm 0.08$. A one dimensional version of the reconstructed energy spectrum is shown in figure 8.22 for the far detector data with the predicted spectra from the extrapolation using the oscillation parameter best fit values. With these oscillation parameters, the integral from 0 to 5 GeV of the predicted far detector spectra is 33.8 events, which includes the predicted background contributions from NC events and cosmic rays. This solution is




Figure 8.23: Final measurement contours for the BPF ν_μ CC disappearance analysis. The blue contours were made including the systematic uncertainties and the red contours without them. The dashed (solid) lines are the $1\text{-}\sigma$ (90%) confidence intervals.

consistent with maximal mixing and with the 2015 results from T2K and MINOS+.

8.6.3 Concluding Remarks

Figure 8.24 demonstrates that the NOvA results using 2.8×10^{20} POT (full detector equivalent) is already comparable to the T2K results with 6.6×10^{20} POT [24] and the MINOS+ results with 15.8×10^{20} POT [25]. With the improved statistics from the data gathered over the next 6 years, the NOvA results clearly have the potential to make significant improvements to the world measurement of $\sin^2 \theta_{23}$.

The BPF analysis performed well producing the measurements

$$|\Delta m_{32}^2| = 2.49_{-0.17}^{+0.19} [\times 10^{-3} \text{eV}^2] \quad (8.6)$$

$$\sin^2 \theta_{23} = 0.51 \pm 0.08. \quad (8.7)$$

Specifically, the technique of separating events by the estimated energy resolution helped




Figure 8.24: Final measurement contours for the BPF ν_μ CC disappearance analysis shown with the most recent results from T2K [24] and MINOS+ [25]. All contours represent 90% confidence intervals.

to minimize the impact of the background samples and emphasize the importance of well reconstructed events. My initial concerns about using an energy estimator trained under different oscillation assumptions proved to be unfounded since the data and the simulations both responded in similar ways to the training used for the energy estimator. Figure 8.26 shows final measurement contours generated with the far detector data sample using each of the three energy estimators described in section 7.3.3. While the size of the contours is different for each of the estimators, the best fit values are almost identical. Since the data appears to favor the maximal mixing solution, the energy estimator trained with maximal mixing parameters produces the best contours. Unfortunately, I had already decided to use the estimator trained with the minimal mixing parameters, but the measurement uncertainty associated with this estimator is still quite good.

Moving forward, I anticipate the NOvA analysis will benefit from many things. First, the number of background cosmic-ray events in the far detector data sample appeared to be higher than expected, which is most likely due to the necessity of opening up a second




Figure 8.25: χ^2 distributions for the two oscillation fit parameters $|\Delta m_{32}^2|$ (left) and $\sin^2 \theta_{23}$ (right.)

beam timing window for the data taken prior to October 2014. As more data is taken with the full detector using only one beam timing window, the proportionate amount of background events from cosmic rays should be reduced. Second, the systematic uncertainty associated with the hadronic energy differences discussed in section 8.4 was by far the dominant uncertainty in this analysis. With more data and further investigation, these differences can be understood leading to improved models for the simulations and tighter sensitivity contours.

There are many things on the horizon for the NOvA experiment. Investigating the hadronic energy differences seen in the near detector data has the potential to provide important insight into the nuclear models used to simulate neutrino interactions. There are also modifications that can be made to the reconstruction algorithms discussed in chapter 6 that could not only improve to the NOvA analyses, but could lead to the deployment of these reconstruction techniques to the next generation of liquid-Argon neutrino detectors. Finally, many of the unanswered questions in neutrino physics: the mass ordering, the octant of θ_{23} , the value of the CP-violating phase factor δ , have the potential of being answered (or at least having some light shed on them) by combining the ν_μ disappearance and ν_e appearance results from 3 years each of running the NuMI beam in the neutrino and anti-neutrino modes. Since the NOvA experiment is just beginning, all of these exciting




Figure 8.26: Final measurement contours using each of the three BPF energy estimators described in section 7.3.3. The green / blue / red contours were created using the unoscillated / “minimal” / “maximal” training assumption. The dashed (solid) lines are the $1\text{-}\sigma$ (90%) confidence intervals. These contours do not include systematic uncertainties.

things are hovering on the horizon, and will soon be within our grasp!

BIBLIOGRAPHY

- [1] David Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, 2nd revised edition, 2008.
- [2] F. Reines and C.L. Cowan. Detection of the free neutrino. *Phys.Rev.*, 92:830–831, 1953.
- [3] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [4] Jr. Davis, Raymond, Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys.Rev.Lett.*, 20:1205–1209, 1968.
- [5] John N. Bahcall, Neta A. Bahcall, and G. Shaviv. Present status of the theoretical predictions for the Cl-36 solar neutrino experiment. *Phys.Rev.Lett.*, 20:1209–1212, 1968.
- [6] Q.R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by 8B solar neutrinos at the Sudbury Neutrino Observatory. *Phys.Rev.Lett.*, 87:071301, 2001.
- [7] Q.R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys.Rev.Lett.*, 89:011301, 2002.
- [8] K.A. Olive et al. Review of Particle Physics. *Chin.Phys.*, C38:090001, 2014.

- [9] V.N. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Phys.Lett.*, B28:493, 1969.
- [10] G. Danby, J.M. Gaillard, Konstantin A. Goulianos, L.M. Lederman, Nari B. Mistry, et al. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys.Rev.Lett.*, 9:36–44, 1962.
- [11] M. Nakamura. Result from DONUT: Direct observation of ν/τ interaction. *Nucl.Phys.Proc.Suppl.*, 77:259–264, 1999.
- [12] K.S. Hirata et al. Observation of a small atmospheric muon-neutrino / electron-neutrino ratio in Kamiokande. *Phys.Lett.*, B280:146–152, 1992.
- [13] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti-neutrino disappearance. *Phys.Rev.Lett.*, 90:021802, 2003.
- [14] Francis Halzen and Alan Martin. *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. John Wiley and Sons, Inc., 1984.
- [15] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam. *Phys.Rev.*, D64:112007, 2001.
- [16] Hans A. Bethe and R. Wilson, James. Revival of a stalled supernova shock by neutrino heating. *Astrophys.J.*, 295:14–23, 1985.
- [17] Hiroshi Nunokawa, Stephen J. Parke, and Jose W.F. Valle. CP Violation and Neutrino Oscillations. *Prog.Part.Nucl.Phys.*, 60:338–402, 2008.
- [18] Evgeny K. Akhmedov, Robert Johansson, Manfred Lindner, Tommy Ohlsson, and Thomas Schwetz. Series expansions for three flavor neutrino oscillation probabilities in matter. *JHEP*, 0404:078, 2004.

- [19] Samoil M. Bilenky and B. Pontecorvo. Lepton Mixing and Neutrino Oscillations. *Phys.Rept.*, 41:225–261, 1978.
- [20] Animesh Chatterjee, Pomita Ghoshal, Srubabati Goswami, and Sushant K. Raut. Octant sensitivity for large $\theta(13)$ in atmospheric and long baseline neutrino experiments. *JHEP*, 1306:010, 2013.
- [21] 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.
- [22] Y. Ashie et al. A Measurement of atmospheric neutrino oscillation parameters by SUPER-KAMIOKANDE I. *Phys.Rev.*, D71:112005, 2005.
- [23] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillation. *Phys.Rev.Lett.*, 93:101801, 2004.
- [24] K. Abe et al. Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with $6.6e20$ protons on target. *Phys.Rev.*, D91(7):072010, 2015.
- [25] Alexandre Sousa. First MINOS+ Data and New Results from MINOS. *AIP Conf.Proc.*, 2015.
- [26] N. Agafonova et al. Discovery of tau neutrino appearance in the CNGS neutrino beam with the OPERA experiment. *arXiv*, 2015.
- [27] D.S. Ayres et al. The NOvA Technical Design Report. *Fermilab Publication*, 2007.
- [28] A. Waldron. NOvA internal document, docDB 9671.
- [29] A. Norman. NOvA internal document, docDB 11860.
- [30] E. Niner. NOvA internal document, docDB 12570.

- [31] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl.Instrum.Meth.*, A389:81–86, 1997.
- [32] The ROOT Team. Root - A Data Analysis Framework. <https://root.cern.ch/drupal/>.
- [33] M. Muether. NOvA internal document, docDB 10855.
- [34] J. Coelho. NOvA internal document, docDB 12733.
- [35] J. Coelho. NOvA internal document, docDB 13213.
- [36] J. Coelho and B. Chowdhury. NOvA internal document, docDB 13546.
- [37] S. Lein. NOvA internal document, docDB 13527.
- [38] X. Bu and K. Sachdev. NOvA internal document, docDB 12437.
- [39] L. Goodenough. NOvA internal document, docDB 12474.
- [40] S. Phan-Budd and L. Goodenough. NOvA internal document, docDB 13572.
- [41] Giuseppe Battistoni, S. Muraro, Paola R. Sala, Fabio Cerutti, A. Ferrari, et al. The FLUKA code: Description and benchmarking. *AIP Conf.Proc.*, 896:31–49, 2007.
- [42] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl.Instrum.Meth.*, A506:250–303, 2003.
- [43] M. Campanella, A. Ferrari, P.R. Sala, and S. Vanini. First Calorimeter Simulation with the FLUGG Prototype. *ATL-SOFT-99-004*, 1999.
- [44] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, et al. The GENIE Neutrino Monte Carlo Generator. *Nucl.Instrum.Meth.*, A614:87–104, 2010.

- [45] C. Haggmann et al. Cosmic-ray Shower Library (CRY). http://nuclear.llnl.gov/simulation/doc_cry_v1.7/cry.pdf, 2012.
- [46] C. Haggmann, D. Lange, and D. Wright. Monte carlo simulation of proton-induced cosmic-ray cascades in the atmosphere. http://nuclear.llnl.gov/simulation/doc_cry_v1.7/cry_physics.pdf, 2012.
- [47] M. Baird. NOvA internal document, docDB 9020.
- [48] G. Davies. NOvA internal document, docDB 13247.
- [49] A. Aurisano. NOvA internal document, docDB 13370.
- [50] A. Aurisano. NOvA internal document, docDB 13577.
- [51] C.N. Chou. The Nature of the Saturation Effect of Fluorescent Scintillators. *Phys.Rev.*, 87:904, 1952.
- [52] D. Pershey. NOvA internal document, docDB 13043second .
- [53] G. Davies. NOvA internal document, docDB 13597.
- [54] D. Pershey. NOvA internal document, docDB 13233.
- [55] J. Liu. NOvA internal document, docDB 7992.
- [56] L. Mualem. NOvA internal document, docDB 13008.
- [57] N. Raddatz and D. Rocco. NOvA internal document, docDB 12917.
- [58] N. Raddatz and D. Rocco. NOvA internal document, docDB 12918.
- [59] C. Backhouse. NOvA internal document, docDB 7410.
- [60] C. Backhouse. NOvA internal document, docDB 13579.

- [61] R. Toner. NOvA internal document, docDB 11691.
- [62] L. Vinton. NOvA internal document, docDB 12347.
- [63] M. Ester, H. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996.
- [64] M. Baird. NOvA internal document, docDB 9755.
- [65] M. Baird. NOvA internal document, docDB 9195.
- [66] M. Baird. NOvA internal document, docDB 12003.
- [67] L. Fernandes and M. Oliveira. Real-time line detection through an improved hough transform voting scheme. *Pattern Recognition*, 41:299 – 314, 2008.
- [68] M. Ohlsson and C. Peterson. Track finding with deformable templates - the elastic arms approach. *Computer Physics Communications*, 71:77 – 98, 1992.
- [69] M. Gyulassy and M. Harlander. Elastic tracking and neural network algorithms for complex pattern recognition. *Computer Physics Communications*, 66:31 – 46, 1991.
- [70] M. Ohlsson. Extensions and explorations of the elastic arms algorithm. *Computer Physics Communications*, 77:19 – 32, 1993.
- [71] R. Frühwirth and A. Strandlie. Track fitting with ambiguities and noise: A study of elastic tracking and non-linear filters. *Computer Physics Communications*, 120:197 – 214, 1999.
- [72] E. Niner. NOvA internal document, docDB 7648.
- [73] E. Niner. NOvA internal document, docDB 9685.
- [74] E. Niner. NOvA internal document, docDB 9864.

- [75] G. Lutz. Optimum Track Fitting in the Presence of Multiple Scattering. *Nucl.Instrum.Meth.*, A273:349, 1988.
- [76] H. Theil. A rank-invariant method of linear and polynomial regression analysis. I, II, III. In *Nederl. Akad. Wetensch., Proc*, volume 53, pages 386–392, 521–525, 1397–1412, 1950.
- [77] Donald E. Groom, Nikolai V. Mokhov, and Sergei I. Striganov. Muon stopping power and range tables 10-MeV to 100-TeV. *Atom.Data Nucl.Data Tabl.*, 78:183–356, 2001.
- [78] M. Baird. NOvA internal document, docDB 12170.
- [79] Andreas Hocker, J. Stelzer, F. Tegenfeldt, H. Voss, K. Voss, et al. TMVA - Toolkit for Multivariate Data Analysis. *PoS*, ACAT:040, 2007.
- [80] M. Baird. NOvA internal document, docDB 12694.
- [81] K. Sachdev. NOvA internal document, docDB 9729.
- [82] B. Zamorano. NOvA internal document, docDB 13441.
- [83] R. Fruhwirth. Application of Kalman filtering to track and vertex fitting. *Nucl.Instrum.Meth.*, A262:444–450, 1987.
- [84] N. Raddatz. NOvA internal document, docDB 13545.
- [85] J. Paley, G. Pawloski, K. Bays, N. Raddatz, S. Lein, and L. Suter. NOvA internal document, docDB 13641.
- [86] L. Suter. NOvA internal document, docDB 13212.
- [87] N. Raddatz. NOvA internal document, docDB 11206.
- [88] N. Raddatz. NOvA internal document, docDB 13326.

- [89] K. Bays. NOvA internal document, docDB 11205.
- [90] T. Ghosh. NOvA internal document, docDB 13409.
- [91] E. Niner. NOvA internal document, docDB 13556.
- [92] K. Bays. NOvA internal document, docDB 11470.
- [93] J. Paley, G. Pawloski, K. Bays, N. Raddatz, S. Lein, L. Suter, and D. Rocco. NOvA internal document, docDB 13703.
- [94] S. Lein. NOvA internal document, docDB 11208.
- [95] M. Baird. NOvA internal document, docDB 12834.
- [96] X. Tian, S. Mishra, H. Duyang, K. Mann, and B. Guo. NOvA internal document, docDB 12990.
- [97] K. Bays. NOvA internal document, docDB 13635.
- [98] K. Bays. NOvA internal document, docDB 13671.
- [99] C. Backhouse. NOvA internal document, docDB 9222.
- [100] J. Lozier. NOvA internal document, docDB 12563.
- [101] N. Raddatz. NOvA internal document, docDB 13237.
- [102] A. Radovic, K. Maan, R. Schroeter, R. Hatcher, H. Duyang, and S. Mishra. NOvA internal document, docDB 13584.
- [103] D. Rocco. NOvA internal document, docDB 13593.
- [104] D. Rocco. NOvA internal document, docDB 13539.
- [105] J. Hartnell. NOvA internal document, docDB 13505.

- [106] M. Tamsett. NOvA internal document, docDB 13510.
- [107] B. Zamorano. NOvA internal document, docDB 13639.
- [108] K. Sachdev. NOvA internal document, docDB 13038.
- [109] J. Paley. NOvA internal document, docDB 13163.
- [110] J. Sepulveda-Quiroz. NOvA internal document, docDB 13663.

Curriculum Vitae

- MICHAEL D. BAIRD -

Indiana University Department of Physics

Swain West, Room 332

Bloomington, IN 47405

Phone: (812) 855-1247

EDUCATION

Indiana University	Ph.D. , High Energy Physics	2011-2015
Indiana University	M.S. , Physics	2010-2011
Indiana University	M.S. , Secondary Education	2004-2005
Indiana University	Graduate Student , Astronomy	1999-2001
Grove City College	B.S. , Applied Physics	1995-1999

RESEARCH: The NOvA Experiment, 2011-2015

Reconstruction and Analysis:

- **Thesis:** *An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter*, Adviser: Dr. Mark Messier. Graduation date: August, 2015
- Developed multivariate methods for identifying muons and estimating neutrino energies.
- Developed a new clustering algorithm to separate potentially overlapping physics events

based on casual relationships between channel hits.

- Enhanced a modified Hough transform to identify major event features used as a first stage in a global vertex identification algorithm.
- Developed a particle tracking module based on a model for multiple scattering capable of fitting tracks under different particle hypotheses.
- Created an analysis package to assign particle identifications to reconstructed tracks using scattering angles, dE/dx values along the track, and reconstructed momenta.
- Performed a Monte Carlo based analysis to determine the potential sensitivity to neutrino oscillations from events entering the detector from outside.

Detector Monitoring and Commissioning:

- Developed a series of online monitoring tools that provide real-time feedback on critical detector performance. These tools display histograms of key variables and provide the ability to generate plots comparing current and previous detector states.
- Created a fully automated series of scripts for nearline processing that produce continually updated metrics posted to a webpage monitored by shifters. This information is also used for data quality purposes to determine good runs and provide bad channel masks for both of the NOvA detectors.
- Installed and maintained software across eight different computers necessary for running the online and nearline monitoring software. This included being the 24-hour on-call expert for all of the data-monitoring software running on these eight machines for over a year.
- Created the software to map a coordinate system based on electronic hardware elements to physical locations in the NOvA detectors for the data acquisition group.
- Spent one month assisting in the installation of hardware and cabling for the NOvA near detector and prototype detector.

Leadership:

- Served as the elected president of the Young NOvA group (an organization of post-docs, graduate, and undergraduate students designed to promote an increased representation of the younger members within the collaboration), August 2013 - June 2014

ADDITIONAL RESEARCH EXPERIENCE _____

Long-Baseline Neutrino Facility: 2011-2012

- Assisted with the research and development of plastic bars coated with a wavelength shifting material to be used to capture scintillation light within the LBNF detectors.

Research Assistant, Indiana University Dept. of Physics: 2010

- Created a numerical simulation of a low energy deuteron beam used to model a pyroelectric fusion device.

Research Assistant, Indiana University Dept. of Astronomy: 2000-2001

- Performed computational simulations of gravitationally interacting few-body systems. Independently developed software to analyze the statistics of the results from thousands of different simulations.

Research Assistant, Harvard-Smithsonian Center for Astrophysics: 1998

- Worked with Dr. Peter Garnavich and Dr. Eric Schlegel to analyze the type II_n supernova 1995N by processing and analyzing 25 to 30 photometric and spectroscopic digital pictures.

PRESENTATIONS

- **‘Event Reconstruction with the NOvA Far Detector’:** Poster presented at the Physics in Collision 2014 conference, September 2014, Bloomington, IN
- **‘Event Reconstruction with the NOvA Far Detector’:** Poster presented at the Neutrino 2014 conference, June 2014, Boston, MA
- **‘Expected Sensitivities from the ν_μ Disappearance Analysis Using the NOvA Detector’:** Talk given at the American Physical Society Division of Particles and Fields 2013 conference, August 2013, Santa Cruz, CA
- **‘An Analysis of the Peculiar Type II In Supernova 1995N’:** Poster presented at the 193rd meeting of the American Astronomical Society, January 1999, Austin, TX